Higher-Order Architectural Connectors[1]

Antónia Lopes

Dep. of Informatics
Faculty of Sciences
University of Lisbon
Campo Grande, 1700 Lisboa
Portugal
mal@di.fc.ul.pt

Michel Wermelinger

Dep. of Informatics
Faculty of Sciences and Technology
New University of Lisbon
2829-516 Caparica
Portugal
mw@di.fct.unl.pt

José Luiz Fiadeiro

ATX Software
Alameda António Sérgio 7, 1-C
2795-023 Linda a Velha

Portugal
jose@fiadeiro.org

Abstract – We develop a notion of higher-order connector towards supporting the systematic construction of architectural connectors for software design. A higher-order connector takes connectors as parameters and allows for services such as security protocols and fault-tolerance mechanisms to be superposed over the interactions that are handled by the connectors passed as actual arguments. The notion is first illustrated over a Unity-like parallel program design language that we have been using for formalising aspects of architectural design. A formal, algebraic semantics is then presented which is independent of any Architectural Description Language. Finally, we show how higher-order connectors can be composed.

1 Introduction

Although components have always been considered to be the fundamental building blocks of software systems, the ways the components of a system interact are determinant for establishing the global system properties, i.e. the properties that emerge from the way the individual components are interconnected. Hence, component interactions have been recognised to be first-class design entities as well, and architectural connectors have emerged as a powerful tool for supporting the design of these interactions [PW92, S93]. Although connectors are widely accepted at the conceptual level, their explicit representation at the linguistic level is not always felt to be necessary: For example, the Darwin [MKG99] Architecture Description Language (ADL) does not include connectors. However, we feel that distinct conceptual entities should correspond to distinct linguistic entities, so that they can truly become first-class and be manipulated as such. In fact, as argued in [MMP00], the current level of support that ADLs provide for connector building is still far from the one awarded to components. For instance, although considerable amounts of

---

[1] This research is partially supported by Fundação para a Ciência e Tecnologia through project POSI/32717/00 (FAST – Formal Approach to Software Architecture).

work can be found on several aspects of connectors [S95, AG97, BCK98, SG01, HUY99, MMP00], further steps are still necessary to achieve a systematic way of constructing new connectors from existing ones. Yet, the ability to manipulate connectors in a systematic and controlled way is essential for promoting reuse and incremental development, and to make it easier to address complex interactions.

At an architecture level of design, component interactions can be very simple (for instance a shared channel), but they can be very complex as well (e.g., database-accessing and networking protocols). Hence, it is very important that we have mechanisms for designing connectors in an incremental and compositional way, as well as principled ways of extending existing ones, promoting reuse. This is especially important for connectors that are used at lower levels of design because it is well known that the implementation of complex protocols is a very difficult and error prone part of system development. Furthermore, as argued in [DMT99], modularising the different kinds of services involved in interaction protocols has other advantages. It prevents interactions from being "hard-wired" across different components and makes it easier to evolve systems (possibly at run-time), because service modules may be added only when necessary, hence preventing performance penalties when such complex interactions are not required.

In this paper we take a step towards this goal by proposing a specification mechanism that allows independent aspects such as compression, fault-tolerance, security, monitoring, *etc.*, to be specified separately, and then composed and integrated with existing connectors. In this way, it becomes possible to benefit from the multiple combinations of different services, ideally chosen *à la carte*. We develop a notion of higher-order connector — a connector that takes a connector as a parameter — through which it is possible to describe the *superposition* of certain capabilities over the form of coordination that is handled by the connector that is passed as an actual argument. In this way we obtain “connector stacks” that are similar in spirit to meta-object towers [DMT99] and to network protocol stacks [OP92], where each stack layer handles a given communication or interaction protocol.

More concretely, we define a higher-order connector through a (formal) parameter declaration and a body connector that models the nature of the service that is superposed on instantiation of the formal parameter. For instance, the monitoring of messages in unidirectional communication can be captured by a higher-order connector with a parameter *Unidirectional-comm* that specifies the kind of connectors to which the service can be applied, and a body connector that describes how an actual parameter is adapted in order to transmit certain messages to a monitoring component.

A higher-order connector can be applied to any connector that instantiates the formal parameter, giving rise to a connector with the new capabilities. In the case of monitoring, the higher-order connector can be applied, for instance, to a connector that models asynchronous communication between a sender and a receiver. Higher-order connectors can also be applied to other high-order connectors. In this case, the result is also a higher-order connector. This later form of application of higher-order connectors can be defined as a parametric instantiation (the instantiation of a parameter with a parameterised entity) and models a non-commutative composition of high-order connectors through which their capabilities are superposed.

The idea of defining higher-order connectors as operators through which new connectors can be built from old ones was proposed by Garlan in [G98], arguing that, conceptually, operations on connectors allow one to factor out common properties for reuse and to better understand the relationships between different connector types. The notation and semantics of such connector operators were recognised to be among the main issues to be dealt with and were later developed by Garlan and Spitznagel [SG01] in the context of the ADL Wright.

Whereas Garlan and Spitznagel define moderately complex and specialized operations, our first attempt at systematic connector construction provided three generic and very simple connector transformations [WF98b]. Our second approach added the notion of higher-order connector, first presented in [WLF00] in a preliminary, informal, and ADL-specific way. In this paper, capitalizing on our previous work on the formal underpinning of connectors [FL97, WF98a] and using the well-known mathematical “technology” of parameterisation [G96] we present a formalisation of higher-order connectors and their composition which is not specific to any ADL. For this purpose, we use the categorical semantics of connectors presented in [FL97, FLW00]. Therein, we establish the semantics of architectural connectors, in the style defined by Allen and Garlan [AG97], independently of specific choices of design languages and behavioural models. We also make use of the characterisation of the minimal set of features that constitutes an ADL, presented in [FL99, FLW00]. As a result, the reader will be able to understand and verify the extent up to which his/her favourite ADL can support the higher-order mechanisms that we are going to present, and to extend it if necessary and desired according to the semantics that we propose.

The paper is organised as follows. Section 2 illustrates, through an example, the key ideas of the notion of higher-order connector we wish to put forward. It shows, in a communication service that involves compression of messages, how the communication service can be separated from the compression service. In section 3, following a categorical approach, we present a parallel program design language, inspired by Unity [CM88] and Interacting Processes [FF96], which is nearer to the abstractions used by conventional programming languages than the process calculi used by others [MKG99, AG97, SG01]. We then present the way higher-order connectors can be defined over this setting and, in section 4, we show that the definition is not specific to the design formalism we have adopted in section 3, as long as the chosen formalism satisfies some structural properties. In section 5, we present a notion of composition of higher-order connectors. We finish with some concluding remarks, including a comparison with related work.

2 Motivation

Software Architecture has put forward connectors as first-class entities for modelling interactions between systems components. According to [AG97], a connector is defined by a set of *roles* and a *glue* specification. Each role describes the behaviour that is expected of each of the interacting parts, i.e., it determines the obligations that they have to fulfil to become instances of the roles. The glue describes how the activities of the role instances are coordinated.

For instance, asynchronous communication through a bounded channel can be represented by a connector *Async* with two roles — *sender* and *receiver*. The glue of *Async* is a bounded *buffer* with a FIFO discipline that prevents the sender from sending a new message when there is no space and prevents the receiver from reading a new message when there are no messages.

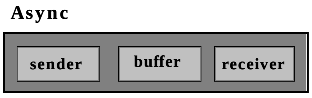

The use of a connector in the construction of a particular system consists in the instantiation of its roles with specific components of the system. The instantiation of a role with a component is possible if and only if the component fulfils the obligations the role determines. Therefore, instantiation corresponds, typically, to a form of refinement.

Let us suppose that in a given system, two components, say *A* and *B*, are connected through *Async*, *A* playing the role of *sender* and *B* playing the role of *receiver*.

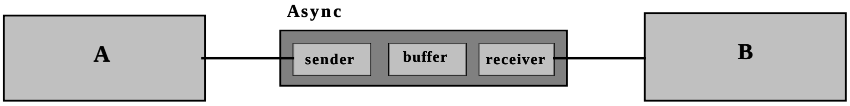

Suppose that, for some reason, the information transmitted from *A* to *B* must be compressed. Clearly, we may develop, from scratch, a new connector *C-Async* with the required functionality, possibly keeping the same roles *sender* and *receiver* but replacing the glue with a new one – *c-async*, and then replace *Async* by *C-Async* in the instantiation to *A* and *B*.

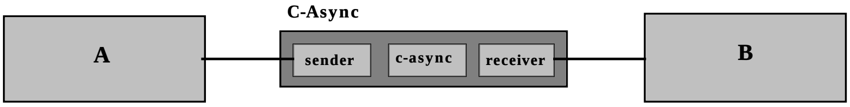

However, it would be certainly better if we could obtain the new connector by simply installing a compress/decompress service over the existing communication service as modelled through *Async*. The idea is to modify *Async* in a way that messages are compressed for transmission without intruding on the original connection, i.e. without “rewiring” the connections to the buffer. Hence, in the resulting connector, the outgoing messages should be compressed before they are put into the buffer and decompressed when they are removed from the buffer, before being delivered to the receiver. It is not difficult to realise that this form of coordination of the sender and receiver activities, embodied by the glue *C-Async*, can be obtained by instantiating the sender role of *Async* with a component *comp* that compresses messages before it transmits them, and by instantiating the *receiver* role with a component *decomp* that decompresses the messages it receives (see figure below).

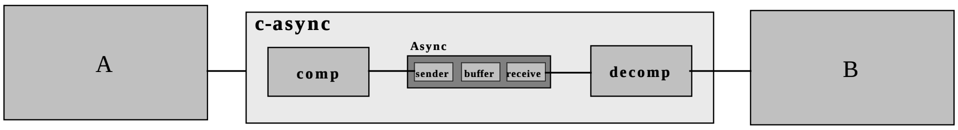

In this way, in the resulting protocol *C-Async*, with the same roles as before and *c-async* as the new glue, a message sent by *sender* is first compressed by *comp*, which then uses *Async* to transmit it to *decomp*. Finally, *decomp* decompresses the message and forwards the result to *receiver*.

It is important to realise that the procedure we described for installing the compress/decompress service over *Async* can be applied to other connectors. In fact, it is possible to give a parameterised description of the compress/decompress service such that the installation of the service over a given connector can be obtained by a suitable instantiation of the parameter.

In this paper, our aim is to introduce these parameterised entities which, as we shall see, are connectors with a distinguished formal parameter part and, hence, are called higher-order connectors. In the example of the compression, this means that we shall define a higher-order connector *Compression(Uni_comm)* whose formal parameter *Uni_comm* is a connector that models a generic unidirectional communication protocol. This formal parameter can be instantiated by several different connectors, in particular by the asynchronous message passing connector that we have been discussing.

3 Higher-Order Connectors in CommUnity

We use the parallel program design language CommUnity in order to make the ideas put forward in the previous section more concrete, and also to motivate the general categorical semantics of higher-order connectors to be given in Section 4 as an extension to our previous work on formalising architectural connectors [FL97, LF99, WF98a, WF98b, WLF00, FLW00].

CommUnity is a Unity-like design language that was initially presented in [FM95, FM97] to show how programs fit into Goguen's categorical approach to General Systems Theory [G73]. Since then, the language and the design framework have been extended in order to provide a formal platform for the architectural design of open, reactive and reconfigurable systems [L99, LF99, WLF01].

3.1 Component Designs

We start by presenting an example of a CommUnit design — *help*. This design models a box consisting of a button, a sensor and a light. Its purpose is to allow a patient to request help in case of medical emergency, with the transmission of the current value of the sensor (e.g., pulse). Pressing the button, which is modelled by the execution of *hreq*, turns on the light, which is modelled by variable *off* becoming false. The light is turned off when the help request is acknowledged. After the button is pressed, the current value of the sensor is read, which is modelled by the execution of the private action *read*, and made available for transmission in the output variable *data*. The private variable *rd* is used to distinguish between states in which the value in *data* is the value to be transmitted or not.

```
component help is
in      sensor:int
out     data:int, off: bool
prv     rd: bool
do      hreq: off → off:=false
□  prv  read: ¬rd ∧¬off →data:=sensor‖rd:=true
□       hack: rd → rd:=false‖off:=true
```

A CommUnity component design is of the form

```
component P is
in      in(V)
out     out(V)
prv     prv(V)
do        □         g: L(g), U(g) →    ||      v:∈F(g,v)
      g∈sh(Γ)                        v∈D(g)
           □    prv g: L(g), U(g) →    ||      v:∈F(g,v)
      g∈prv(Γ)                       v∈D(g)
```

where

- *V* is the set of *channels*. Channels can be declared as *input, output* or *private*. Input channels are used for reading data from the environment of the component. Output and private channels are local to the component, i.e., the values they hold are under the control of the component and cannot be modified by the environment. We use *loc(V)* to denote the set of local channels. Output channels make available data that can be read by the environment but the data available from private channels cannot. Each channel *v* is typed with a sort *sort(v)*.

- Γ is the set of *action names*. Actions can be declared either as *private* or *shared* (for simplicity, we only declare which actions are private). Private actions represent internal computations and their execution is uniquely under the control of the component. In contrast, shared actions represent interactions between the component and the environment and their execution is also under the control of the environment.
- For each action *g*:
    - *D(g)* consists of the local channels to which action *g* can write – its write frame. For every local channel *v*, we also denote by *D(v)* the set of actions that can write into *v*.
    - *L(g)* and *U(g)* are two conditions such that $U(g) \supset L(g)$. They establish an interval in which the enabling condition of *g* must lie: *L(g)* is the lower bound of the interval, i.e., it is implied by the enabling condition; *U(g)* is the upper bound, i.e., it implies the enabling condition. Therefore, the negation of *L(g)* establishes a *blocking* condition, whereas *U(g)* establishes a *progress* condition. The enabling condition is fully determined when *L(g)* and *U(g)* are equivalent, in which case we write only one condition.
    - For every channel *v* in *D(g)*, *F(g,v)* is a expression that denotes a set – the set of values that *g* can write into *v*. We abbreviate deterministic assignments of the form $v{:}\in\{t\}$ as *v:*=t. When an action has empty domain, we use the expression `skip` instead.

Notice that CommUnity supports several mechanisms for underspecification – actions may be underspecified in the sense that their enabling conditions may not be fully determined and subject to refinement by reducing the interval established by *L* and *U*, and their effects on the channels may be undetermined and also subject to refinement by replacing the assignment sets by proper subsets.

When, for every $g \in \Gamma$, *L(g)* and *U(g)* coincide, and *F(g,v)* is a singleton for every local channel *v*, then the design is called a *program*. The behaviour of a program without input channels is as follows. At each execution step, one of the actions whose enabling condition holds of the current state is selected, and its assignments are executed atomically in parallel. Furthermore, private actions that are infinitely often enabled are guaranteed to be selected infinitely often (see [LF99] for a model-theoretic semantics of CommUnity). A program with input channels is *open* in the sense that it needs to be connected to other components of the system to read data. We will explain how such connections can be established in later sections.

As a language, CommUnity is independent of the actual data types that are used and, hence, we have assumed that there are pre-defined sorts and functions given by a fixed algebraic specification $\Xi=<<S,\Omega>,\Phi>$ where S is a set of sorts, Ω is a set of operations, and Φ is a set of first-order axioms specifying the functionality of the operations [EM85]. For the purposes of examples in this paper, we consider an algebraic signature containing basic data types such as booleans (*bool*), integers (*int*), queues *(queue(_,_)), etc.*, with the usual operations.

Formally, CommUnity designs can be defined as follows.

A design signature is a tuple $<V,\Gamma,tv,ta,D>$ *where*

- *V is an S-indexed family of mutually disjoint finite sets,*
- *Γ is a finite set,*
- *$tv: V\to\{out,in,prv\}$ is a total function,*
- *$ta: \Gamma\to\{sh,prv\}$ is a total function,*
- *$D: \Gamma\to 2^{loc(V)}$ is a total function.*

A design is a pair $<\theta,\Delta>$ *where* $\theta=<V,\Gamma,tv,ta,D>$ *is a design signature and* Δ*, the body of the design, is a tuple* $<F,L,U>$ *where:*

- *F assigns to every action $g\in\Gamma$ and to every $v\in D(g)$ a set expression of sort $sort(v)$ and is undefined in all other cases,*
- *L and U assign a proposition over V to every action $g\in\Gamma$.*

In order to support higher levels of design, a CommUnity component design may also be given by a parameterised design, more concretely, by a design parameterised by an algebraic specification indicated after the name of the component (see examples in the next section). This parameter is instantiated at configuration time, i.e., when a specific component needs to be included in the configuration of the system being built, or as part of the reconfiguration of an existing system.

3.2 Connectors

In CommUnity it is also possible to describe a design as the interconnection of a number of interacting component designs by defining a configuration. Hence, a connector in CommUnity consists of a set of roles, a glue specification, and a configuration involving these designs.

Let us consider the connector *Async[t+K]* representing asynchronous communication of values of type *t* through a bounded buffer with capacity *K*. This connector has two roles — *sender[t]* and *receiver[t]* — defining the behaviour required of the components to which the connector can be applied. For the *sender* we require that it does not produce another message before the previous one has been processed. After producing a message, the *sender* should expect an acknowledgement to produce a new message. For the *receiver*, we simply require that it has an action that models the reception of a message.

```
component sender[t] is                          component receiver[t] is
out    o:t                                      in   i:t
prv    rd:bool                                  do   rec:true,false→skip
do  prv prod:¬rd,false→o:∈t ∥ rd:=true
 □      send:rd,false→rd:=false
```

In order to leave unspecified when and how many messages the sender (receiver) will send (receive), and in which situations the sender will produce a new message, the progress guards of these actions are all false. Because progress conditions establish an upper bound for enabledeness, if the progress condition of an action is false then its enabling condition can be as strong as we wish. Furthermore, we chose the least deterministic assignment $o{:}\in t$ for the production of messages in order to avoid committing to a particular discipline of production.

Both designs are parameterised by the algebraic specification $<<\{t\},\varnothing>,\varnothing>$ that consists just of a sort (without operations nor axioms), which we denoted simply by t. When the connector needs to be used in the configuration of a specific system, the sort t is instantiated with whatever sort is appropriate.

The glue of *Async[t+K]* is a bounded buffer with a FIFO discipline that prevents the sender from sending a new message when there is no space and prevents the receiver from reading a new message when there are no messages. In CommUnity this buffer can be designed as follows.

```
component buffer [Ξ^b] is
in   i:t
out  o:t
prv  rd:bool; q:queue(K,t)
do   put:¬full(q)→ q:=enqueue(i,q)
□prv next:¬empty(q)∧¬rd → o:=head(q) ∥ q:=tail(q) ∥ rd:=true
□    get: rd → rd:=false
```

This buffer can store, through the action *put*, messages of sort t, received from the environment through the input channel i, as long there is space for them. It can also make stored messages available to the environment through the output channel o and the action *next*. Naturally, this activity is possible only when there are messages in store and the current message in o has already been read by the environment, which is modelled by the action *get* and the private channel *rd*.

The type of messages as well as the capacity of the buffer are part of the parameter specification Ξ^b defined below, where we use Ξ^{nat} to represent the subspecification of Ξ that is concerned with the specification of natural numbers. We also use $t+K$ to denote Ξ^b.

```
spec Ξ^b is Ξ^nat +
sorts  t
ops    K: ->nat
```

It remains to define in which way the roles and the glue are connected. The model of interaction between components in CommUnity is based on action synchronisation and the interconnection of input channels of a component with output channels of other components. Although these are common forms of interaction, CommUnity requires interaction between components — name bindings — to be made explicit in the systems configurations. Name bindings are established as relationships between the signatures of the corresponding components and are defined with the help of additional signatures (representing the interaction points) and signature maps (morphisms).

For instance, in order to establish that messages from a *sender* component are sent (to a receiver) through a bounded channel, we consider the following configuration

$$\begin{tikzcd} & \mathtt{cable[t]} \arrow[ddl] \arrow[ddr] & \\ & \mathtt{o} \leftarrow \bullet \rightarrow \mathtt{i} & \\ & \mathtt{send} \rightarrow \bullet \leftarrow \mathtt{put} & \\ \mathtt{sig(\textbf{sender}[t])} & & \mathtt{sig(\textbf{buffer}[t+K])} \end{tikzcd}$$

where *cable[t]* is a signature that consists of an input channel of sort *t* and a shared action. The names of this channel and of this action are not relevant: they are only placeholders used to define the name bindings, and hence, we used • for both.

In this configuration, the input channel of *cable* is mapped to the output channel *o* of the *sender* and to the input channel *i* of *buffer*. This establishes an I/O interconnection between the *sender* and the *buffer*. Moreover, the actions *send* of *sender* and *put* of *buffer* are mapped to the shared action of *cable*. This defines that *sender* and *buffer* must synchronise each time either of them wants to perform the corresponding action.

The signature morphisms involved in the configurations are defined as follows.

A morphism $\sigma: \theta_1 \rightarrow \theta_2$ *between signatures* $\theta_1=<V_1,\Gamma_1,tv_1,ta_1,D_1>$ *and* $\theta_2=<V_2,\Gamma_2,tv_2,ta_2,D_2>$ *is a pair* $<\sigma_{var},\sigma_{ac}>$ *where*

- $\sigma_{var}: V_1 \rightarrow V_2$ *is a total function satisfying*
  1. $sort_2(\sigma_{var}(v))=sort_1(v)$ *for every* $v \in V_1$;
  2. $\sigma_{var}(o) \in out(V_2)$ *for every* $o \in out(V_1)$;
  3. $\sigma_{var}(i) \in out(V_2) \cup in(V_2)$ *for every* $i \in in(V_1)$;
  4. $\sigma_{var}(p) \in prv(V_2)$ *for every* $p \in prv(V_1)$.
- $\sigma_{ac}: \Gamma_2 \rightarrow \Gamma_1$ *is a partial mapping satisfying for every* $g \in \Gamma_2$ *s.t.* $\sigma_{ac}(g)$ *is defined:*

5. *if* $g \in sh(\Gamma_2)$ *then* $\sigma_{ac}(g) \in sh(\Gamma_1)$*;*

6. *if* $g \in prv(\Gamma_2)$ *then* $\sigma_{ac}(g) \in prv(\Gamma_1)$*;*

7. $\sigma_{var}(D_1(\sigma_{ac}(g))) \subseteq D_2(g)$*;*

8. $\sigma_{ac}(D_2(\sigma_{var}(v))) \subseteq D_1(v)$ *for every* $v \in loc(V_1)$*.*

A morphism σ from θ_1 to θ_2 is intended to support the identification of a way in which a component with signature θ_1 is embedded in a larger system with signature θ_2. This justifies the various constructions and constraints in the definition.

The function σ_{var} identifies for each channel of the component the corresponding channel of the system. The partial mapping σ_{ac} identifies the action of the component that is involved in each action of the system, if ever. Sorts of channels have to be preserved (condition 1) but, in terms of their classification, input channels of a component may become output channels of the system (condition 3). This is because the result of interconnecting an input channel of a component with an output channel of another component is an output channel of the system (mechanisms for internalising communication can be applied but they are not the default in a configuration). Conditions 7 and 8 on write frames imply that actions of the system in which a component is not involved cannot have local channels of the component in their write frame. That is, change within a component is completely encapsulated in the structure of actions defined for the component.

In the case of configurations involving components parameterised with data type specifications, the signature morphisms must satisfy an additionally property:

A morphism $\sigma{:}\theta_1[\Pi_1] \rightarrow \theta_2[\Pi_2]$ *between parameterised signatures is a signature morphism* $\sigma{:}\ \theta_1 \rightarrow \theta_2$ *for which* Π_1 *is a subspecification [EM85] of* Π_2*.*

Let us consider again the connector *Async[t+K]* we have been defining. The configuration depicted below completes its definition, establishing how the roles and the glue are connected.

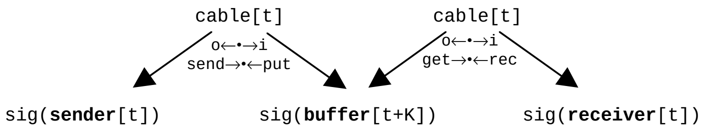

As explained previously, the left-hand side morphisms define that *sender* and *buffer* must synchronise on actions *send* and *put*, and establish the interconnection of the output channel *o* of *sender* with the input channel *i* of *buffer*. On the other hand, the right-hand side morphisms define that *buffer* and *receiver* must synchronise on actions *get* and *rec* and establish the interconnection of the output channel *o* of *buffer* with the input channel *i* of *receiver*.

Not every diagram of signatures represents a meaningful configuration in the sense that there are restrictions on the way that we can interconnect components that are not captured by the notion of morphism alone but require the whole diagram. The two following rules express the restrictions on diagrams that make them well-formed configurations:

- An output channel of a component cannot be connected (directly or indirectly through input channels) with output channels of the same or other components.
- Private channels and private actions cannot be involved in the connections.

It is important to notice that the second rule establishes the configuration semantics of private actions and channels. It supports the intuitive semantics we gave in section 2.1, namely that private channels cannot be read by the environment and that the execution of private actions is uniquely under the control of the component.

Rather than using diagrams involving signatures and signature morphisms, a more user-friendly notation may be adopted. For instance, the asynchronous communication defined above could be described as follows.

In this notation, the name bindings are still explicit but are expressed in terms of arcs that connect channels and actions directly. These configurations can be easily translated into categorical diagrams involving signatures and signature morphisms.

The semantics of connectors, and of configuration diagrams in general, relies on an extension of the notion of signature morphism that allows us to establish relationships between designs. Design morphisms capture relationships between components and the systems that they are part-of. They can be seen to provide a formalisation for a notion of superposition that is similar to those that have been used for parallel program design [CM88, K93].

A superposition morphism $\sigma:P_1\rightarrow P_2$ *consists of a signature morphism* $\sigma:\theta_1\rightarrow\theta_2$ *s.t., for every* $g\in\Gamma_2$ *s.t.* $\sigma_{ac}(g)$ *is defined:*

1. *for every* $v\in D_1(\sigma_{ac}(g)), \Phi \models (F_2(g,\sigma_{var}(v))\subseteq \sigma(F_1(\sigma_{ac}(g),v)))$;
2. $\Phi \models (L_2(g) \supset \sigma(L_1(\sigma_{ac}(g))))$;
3. $\Phi \models (U_2(g) \supset \sigma(U_1(\sigma_{ac}(g))))$;
where $\models$ *denotes validity in the first-order sense. Designs and superposition morphisms constitute the category* ***c-DSGN****.*

A morphism $\sigma:P_1\rightarrow P_2$ identifies a way in which P_1 is "augmented" to become P_2 so that P_2 can be considered as having been obtained from P_1 through the superposition of additional behaviour, namely the interconnection of one or more components.

Condition 1 corresponds to the preservation of the functionality of P_1: the effects of the actions have to be preserved or made more deterministic. Conditions 2 and 3 allow the bounds for the enabledeness of each action to be strengthened but not weakened. Strengthening of the lower bound is typical of superposition and reflects the fact that all the components that participate in the execution of a joint action have to give their permission for the action to be performed. On the other hand, it is clear that progress for a joint action can only be guaranteed when all the components involved can locally guarantee so.

Let us consider that we have established interactions between component designs $P_1,\ldots,P_n$ at the level of their signatures through a diagram ***D***. Such a diagram can be trivially lifted to a diagram ***D'*** of designs and superposition morphisms: the signature of each design is replaced by the design itself; every cable *cb* in ***D*** is replaced by ***dsgn****(cb),* the design with signature *cb* and tautological bounds for the enabledeness of each action, and the least deterministic assignment for the channels in the write frame of each action. More concretely, ***dsgn****(cb)* consists of, for each action name *g* in *cb*, the action

$$g\!: true,\ true \rightarrow \mathop{\|}_{v \in D(g)} v\!:\in sort(v)$$

Defined in this way, ***dsgn****(cb)* is a design that is "neutral" with respect to the establishment of superposition morphisms in the sense that every signature morphism $\sigma{:}cb \rightarrow \mathbf{sig}(P)$ defines a superposition morphism $\sigma{:}\mathbf{dsgn}(cb) \rightarrow P$.

On the account of this transformation, every configuration can be transformed into a single design that represents the whole system by taking the colimit of the diagram ***D'*** in the category ***c-DSGN***. We now describe the intuitive meaning of the colimit and its construction.

Consider first the simplest case, the one in which there are no interactions. This means that any two channels of two designs are different, even if they have the same name. Therefore, the channels of the resulting design are the disjoint union of the components' channels. Concerning the actions, the parallel composition contains all possible combinations of actions that involve one at most one action from each component. This is because there is no restriction on their co-occurrence. More concretely, the actions of the resulting design are the tuples of actions of the components $a_1|\ldots|a_k$, containing at most one action of each component. In this way, the colimit provides not an interleaving but a concurrent semantics for parallel composition.

In the presence of interactions, the colimit “merges” the input channels identified with an output channel into that output channel, and each tuple $a_1|\ldots|a_k$ is retained iff, for every action a_i in the tuple, every action that is required to synchronise with a_i is also in the tuple.

With or without interactions, for each action $a_1 | \dots | a_k$, the bounds for enabledeness are the conjunction of the bounds of all a_i, and its assignments are the union of the assignments of all a_i.

In the case of configurations involving parameterised designs, we need first to define the corresponding superposition morphisms.

*A superposition morphism from a parameterised design $P_1(\Pi_1)$ to a parameterised design $P_2(\Pi_2)$ is a morphism of **c-DSGN** from P_1 to P_2 for which Π_1 is a subspecification of Π_2.*

In this way, in the case of a configuration involving parameterised designs, the single design that represents the whole system is parameterised by the union of the involved data type specifications.

As an illustration of the transformation of a system configuration into a design that represents the whole system, we consider again the connector *Async[t+K]*. The lifting of the diagram presented previously, which establishes how the roles and the glue are connected at the level of their signatures, is

$$\begin{tikzcd}
 & \mathbf{dsgn}(\mathtt{cable[t]}) \arrow[dl] \arrow[dr] \arrow[d, phantom, "\substack{\mathtt{o}\leftarrow\bullet\rightarrow\mathtt{i} \\ \mathtt{send}\rightarrow\bullet\leftarrow\mathtt{put}}"] & & \mathbf{dsgn}(\mathtt{cable[t]}) \arrow[dl] \arrow[dr] \arrow[d, phantom, "\substack{\mathtt{o}\leftarrow\bullet\rightarrow\mathtt{i} \\ \mathtt{get}\rightarrow\bullet\leftarrow\mathtt{rec}}"] & \\
\mathbf{sender}\mathtt{[t]} & & \mathbf{buffer}\mathtt{[t+K]} & & \mathbf{receiver}\mathtt{[t]}
\end{tikzcd}$$

where ***dsgn****(cable[t])* is a design with an input channel of sort *t* and a shared action whose bounds for enabledeness are *true*.

The colimit of this diagram returns the design given, up to an isomorphism, by the design below. This design models the parallel composition of *sender*, *receiver* and *buffer* with the restrictions defined by the given configuration diagram and defines the semantics of the connector.

```
component async[t+K] is
out     o_s,o_b:t
prv     rd_s, rd_r:bool; q:queue(K,t)
do prv  prod:¬rd_s, false→ o_s:∈t‖rd_s:=true
□       send|put:¬full(q)∧¬rd_s, false→ q:=enqueue(o_s,q) ‖rd_s:=false
□  prv  next:¬empty(q)∧¬rd_r→ o_b:=head(q)‖ q:=tail(q)‖rd_r:=true
□       rec|get: rd_r, false → rd_r:=false
```

This design provides the means for global properties of the protocol that it defines to be derived. For instance, using the logical formalism for reasoning about CommUnity designs defined in [L99], it is possible to conclude that *async[t+K]* has the following property expressed in a branching time temporal logic:

$$\boldsymbol{A}((send \mid put \wedge o_s{=}msg) \supset \boldsymbol{F}(o_b{=}msg \wedge \langle rec \mid get\rangle true))$$

This sentence expresses that if a message *msg* is sent, eventually *msg* will be made available in the input channel of the receiver, ready to be received. In other words, at least a copy of each message is delivered. In the same way, it is possible to conclude that the correctness of the transmission/reception of data (in order message delivery) does not depend on the speed at which messages are produced and consumed.

3.3 Using Connectors in System Construction

The use of a connector in the construction of a particular system is achieved by the instantiation of its roles with specific components of the system. To model instantiation we use a different kind of design morphism that ensures that the behaviour specified by a role is satisfied by the instance. These morphisms correspond to a form of refinement and, hence, are called refinement morphisms.

A refinement morphism $\sigma{:}P_1 \rightarrow P_2$ *consists of a signature morphism* $\sigma{:}\theta_1 \rightarrow \theta_2$ *s.t.:*

1. $\sigma_{var}(inp(V_1)) \subseteq inp(V_2)$ *and* $\sigma_{var}\downarrow(out(V_1) \cup inp(V_1))$ *is injective;*

2. *For every* $g \in sh(\Gamma_1)$, $\sigma^{-1}_{ac}(g) \neq \emptyset$.

For every $g \in \Gamma_2$ *s.t.* $\sigma_{ac}(g)$ *is defined:*

3. $\Phi \models (F_2(g,\sigma_{var}(v)) \subseteq \sigma(F_1(\sigma_{ac}(g),v)))$ *for every* $v \in D_1(\sigma_{ac}(g))$;

4. $\Phi \models (L_2(g) \supset \sigma(L_1(\sigma_{ac}(g))))$.

For every $g_1 \in \Gamma_1$:

5. $\Phi \models (\sigma(U_1(g_1)) \supset \bigvee_{\sigma_{ac}(g_2)=g_1} U_2(g_2))$.

Designs and refinement morphisms constitute the category ***r-DSGN****.*

A refinement morphism supports the identification of a way in which a design P_1 is refined by another design P_2. Each channel of P_1 has a corresponding channel in P_2 and each action *g* of P_1 is implemented by the set of actions $\sigma^{-1}_{ac}(g)$ in the sense that $\sigma^{-1}_{ac}(g)$ is a menu of refinements for action *g*. The actions for which σ_{ac} is left undefined (the new actions) and the channels which are not in $\sigma_{var}(V_1)$ (the new channels) introduce more detail in the refined description of the component.

Condition 1 ensures that an input channel cannot be made local by refinement and that different channels of the interface cannot be collapsed into a single one (refinement does not alter the border between the system and its environment). Condition 2 ensures that those actions that model interaction between the design and its environment have to be implemented. Conditions 4 and 5 state that the "interval" of (allowed) non-determinism defined by the two bounds for enabledess can only be preserved or reduced by refinement. This is intuitive because refinement, pointing in the direction of an implementation, should reduce allowed non-determinism. This is also the reason why the non-determinism of assignments must be preserved or decreased (condition 3).

It is important to notice that, although refinement and superposition morphisms have some conditions in common, the two relationships are very different. As evidence of this notice the fact that in CommUnity, as in other formalisms such as CSP [H85], a design is not necessarily refined by a system of which it is a component.

The component design *help* defined previously is an example of a refinement of the design *sender(int)* — the result of the instantiation of *t* in *sender[t]* with the sort *int* of Ξ. It refines *sender(int)* through the refinement morphism $\eta{:}\, sender(int) \rightarrow help$ defined by

$\eta_{var}(o)=data,\ \eta_{var}(rd)=rd$
$\eta_{ac}(read)=prod,\ \eta_{ac}(hack)=send.$

In *help,* the production of messages to be sent is modelled by the action *read* and the messages are made available in the output channel *data*. The production of messages, that was left unspecified in *sender,* corresponds to the sensor readings.

This refinement can be represented graphically as depicted below. Notice that non-private channels and actions are placed on the boundary of the component and private ones inside.

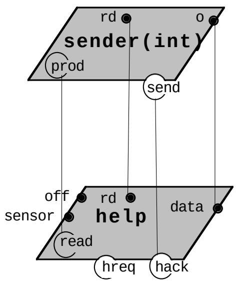

Let us suppose that *centre* is a design that models an assistance centre that refines *receiver(int)* through some refinement morphism κ. The connector *Async(int+10)* can be used to interconnect *help* and *centre*. The resulting system — a system in which the *help* component sends the help requests to the assistance *centre* through a bounded channel with capacity for ten messages— can be represented as follows.

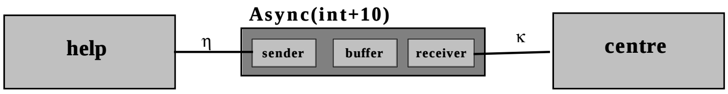

In this case, we have used the connector *Async[t+K]* with *t* instantiated with sort *int* and *K* instantiated with 10. In more abstract levels of design, it may be useful to use *Async[t+K]* for coordinating the activities of parameterised designs. In such cases, the instantiation of the connector has to be defined by refinement morphisms between parameterised designs.

A refinement morphism from a parameterised design $P_1(\Pi_1)$ *to a parameterised design* $P_2(\Pi_2)$ *is a morphism of* ***r-DSGN*** *from* P_1 *to* P_2 *for which* Π_1 *is a subspecification of* Π_2.

In the next section we shall see examples of such kind of morphisms.

3.4 Adapting Connectors

As explained before, it is important to have principled ways to adapt connectors to new situations, for instance in order to incorporate compression, fault-tolerance, security, monitoring, etc.

Let us consider *compression* once more as an example. In this case, the goal is to adapt a connector that represents a communication protocol in order to compress data for transmission in a transparent way. In order to be able to give a first-class description of this form of adaptation, the kind of communication protocol modelled by the adapted connector needs to be made more precise. We shall describe the *compression* adaptation mechanism only for connectors that model unidirectional communication protocols.

A generic unidirectional communication protocol can be modelled by the binary connector *Uni-comm[s]*

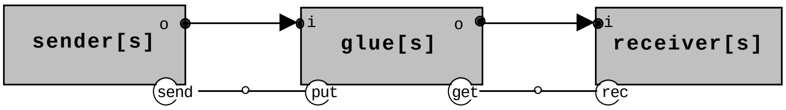

where

```
component glue [s] is
in      i:s
out     o:s
do      put: true,false → skip
□   prv prod: true,false → o:∈s
□       get: true,false → skip
```

and *sender[s]* and *receiver[s]* are defined as before. Notice that this glue leaves completely unspecified the way in which messages are processed and transmitted.

Our aim is to install a compression/decompression service over *Uni-comm*. That is to say, our aim is to apply an operator to *Uni-comm* such that, in the resulting connector, a message sent by the sender is compressed before it is transmitted through *Uni-comm* and then decompressed before it is delivered to the receiver. We shall see that such an operator can be described by a higher-order connector where the compression and decompression algorithms are taken as parameters. More concretely, it is parameterised by the algebraic specification described below.

```
spec Ξcd is Ξnat +
sorts   s,t
ops     comp:t->s        decomp:s->t
        size_s:s->nat    size_t:t->nat
axioms  decomp(comp(x))=x, for any x:t
        size_s(comp(x)) ≤ size_t(x), for any x:t
```

Sorts *t* and *s* represent the types of original and compressed messages, respectively. The operation *comp* represents the process of compression of a single message, and *decomp* the inverse process of decompression. It is required that the size of the compressed message is not greater than the size of the original message. At configuration time, these data elements must be instantiated with specific sorts and operations.

The higher-order connector itself, which we name *Compression(Uni-comm)*[Ξ^{cd}], is defined by

- the binary connector *Compression*[Ξ^{cd}]

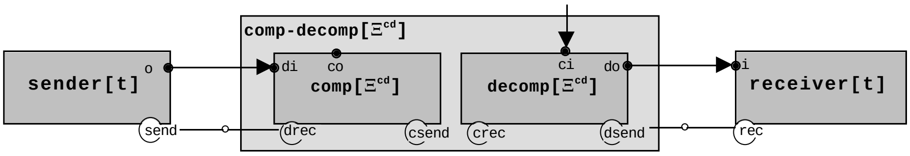

where the glue, *comp-decomp*[Ξ^{cd}], is defined in terms of a configuration with the following two components:

```
component comp[Ξcd] is
in   di:t
out  co:s
prv  v:t; rd,msg:bool
do   drec:¬msg → v:=di||msg:=true
[]prv comp:¬rd∧ msg → co:=comp(v)||rd:=true
[]    csend:rd → rd:=false||msg:=false
```

```
component decomp[Ξcd] is
in   ci:s
out  do:t
prv  v:s; rd,msg:bool
do   crec: ¬msg → v:=ci||msg:=true
[]prv dec:¬rd∧ msg → do:=decomp(v)||rd:=true
[]    dsend: rd → rd:=false||msg:=false
```

Design *comp*[Ξ^{cd}] models the compression of messages of type *t* received through *di* into messages of type *s* that are then transmitted through *co*. Design *decomp*[Ξ^{cd}] models the decompression of messages of type *s* received through *ci* into messages of type *t* that are then transmitted through *do*.

- the connector *Uni-comm[s]* — the formal parameter;
- the refinement morphisms

η_s*:sender[s]*$\rightarrow$*comp-decomp*[Ξ^{cd}] and η_r*:receiver[s]* $\rightarrow$*comp-decomp*[Ξ^{cd}]

induced, respectively, by the refinement morphisms

η^*_s*:sender[s]*$\rightarrow$*comp*[Ξ^{cd}]

$\eta^*_s(o)=co,\ \eta^*_s(rd)=rd,\ \eta^*_s(comp)=prod,\ \eta^*_s(csend)=send$

$\eta^*_r{:}receiver[s]\rightarrow decomp[\Xi^{cd}]$

$\eta^*_r(i)=ci,\ \eta^*_r(crec)=rec$

Because components *comp* and *decomp* do not interact, any component refined by one of them is also refined by their composition *comp-decomp*$[\Xi^{cd}]$. The corresponding induced morphisms have only to take into account the renaming of channels and actions that takes place in composition.

Putting the two previous pictures together we get a graphical representation of the higher-order connector *Compression(Uni-comm)*$[\Xi^{cd}]$.

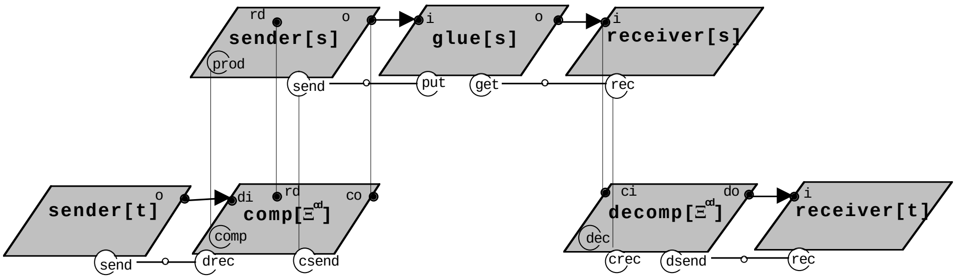

In summary, *Compression(Uni-comm)*$[\Xi^{cd}]$ has the formal parameter *Uni-comm[s]*, which restricts the actual connectors to which the service of compression/ decompression can be applied — it requires that the actual connector models a unidirectional communication protocol. The connector *Compression* describes, on the one hand, that messages sent by the actual sender are transmitted to *comp* which compresses them and, on the other hand, that *decomp* decompresses the messages it receives and delivers the result to the actual receiver. Finally, the two refinement morphisms establish the instantiation of *Uni-comm[s]* with *comp[s]* in the role of sender and *decomp[s]* in the role of receiver. In this way, it is established that the formal parameter *Uni-comm[s]* is the connector used to transmit compressed messages.

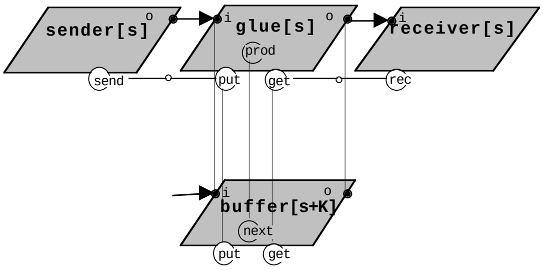

Now it remains to explain the procedure of parameter passing, i.e., how the service just described can be installed over a specific connector and how the resulting connector is obtained.

We consider again the *Async* connector. In this case, it is not difficult to realize that we may replace the formal parameter of *Compression(Uni-comm)*$[\Xi^{cd}]$ by *Async* because this connector does model a unidirectional communication protocol. More concretely, *Async* has exactly the same roles that *Uni-comm* and its glue is a refinement of *Uni-comm*'s glue.

In a more general situation, the instantiation of a higher-order connector is established by a suitable fitting morphism from the formal to the actual connector. Such a morphism formulates the correspondence between the roles and glue of the formal parameter with those of the actual parameter connector. In the next section we will present and discuss these morphisms in more detail.

The construction of a new connector from the given higher-order connector and the actual parameter connector is straightforward. We only need to compose the interconnections of the *buffer* to *sender* and *receiver* with the refinements η_s and η_r that define the instantiation of *Uni-comm* with *comp* and *decomp*, respectively. For example, channel *co* of *comp* becomes connected to the input channel *i* of *buffer* because *co* corresponds to the channel *o* of *sender* which in turn is, in *Async*, connected to *i*. The resulting configuration fully defines the connector *Compression(Async)*$[\Xi^{cd}+K]$. Its roles are *sender* and *receiver* and its glue *c-buffer-d*$[\Xi^{cd}+K]$ is defined in terms of a configuration involving *comp*, *decomp* and *buffer* as shown below.

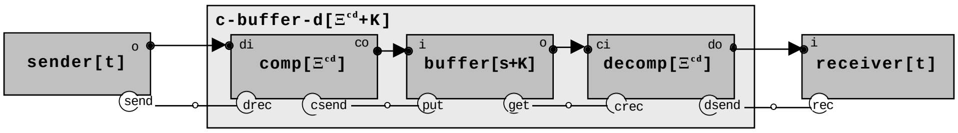

Summarising, in this section, we have described the installation of a compression-decompression service over a unidirectional communication protocol as a parameterised entity that has connectors as parameters and result and, thus, is called a higher-order connector. Then we have explained how the higher-order connector can be instantiated with a specific connector and, finally, we showed how the resulting connector is obtained.

We end this section by presenting another example of a higher-order connector — *monitoring*. The aim is to model the adaptation of a unidirectional communication protocol in order to transmit certain kind of messages (for instance, error messages) to a monitoring component.

The kind of messages that should be transmitted to the monitoring component is taken as a parameter. More concretely, we define a higher-order connector that is parameterised by the following algebraic specification:

```
spec Ξ^m is Ξ^bool+
sorts   s
ops     to_monitor:s->bool
```

Sort *s* represents the type of messages, and operation *to_monitor* identifies the special kind of messages that are to be monitored. We use Ξ^{bool} to represent the subspecification of Ξ that is concerned with the specification of booleans.

For ease of presentation, we shall consider that the communication with the monitoring component is achieved by synchronous message passing. However, it would be more appropriate to model *monitoring* with a higher-order connector with two formal parameters, both modelling a unidirectional communication protocol. One of them would be used for the normal transmission of messages and the other for the transmission to the monitoring component.

The higher-order connector itself, *Monitoring(Uni-comm)*[Ξ^m], consists of

- the connector *Monitoring*[Ξ^m] defined by

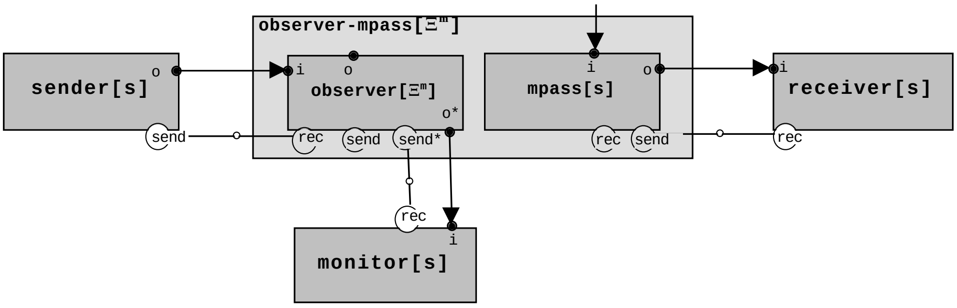

where the glue, *observer-mpass*[Ξ^m], is defined in terms of a configuration with the following two components:

```
component observer[Ξ^m] is                      component mpass[s] is
in      i:s                                     in      i:s
out     o,o*:s                                  out     o:s
prv     v:s; rd,rd*,msg:bool                    prv     rd:bool
do      rec:¬msg →v:=i||msg:=true               do      rec:¬rd→ o:=i||rd:=true
□ prv   obsv:¬rd∧¬rd*∧msg →  msg:=false||       □       send:rd→ rd:=false
          o:=v||rd:=true||o*:=v||rd*:=to_monitor(i)
□       send:rd → rd:=false
□       send*:rd*→ rd*:=false
```

Component *observer[Ξ^m]* observes the messages to be transmitted and forwards a copy of certain transmitted messages to a third component. More precisely, it sends through *o* the messages received in *i*, and sends through *o** those messages that satisfy *to_monitor*. Component *mpass[s]* just transmits through *o* the messages received in *i*.

The connector has three roles — *sender*, *receiver* and *monitor*. The role *monitor[s]* is similar to *receiver[s]*:

```
component monitor [s] is
in      i: s
do      rec: true → skip
```

Notice, however, that the progress condition of *rec* is true in *monitor* and false in *receiver*. This means that any component that acts as monitor must be always willing to read the values that are input through *i* whereas the actual receiving component may decide when and how many times it will read the values sent to it. In this way, it is ensured that the monitoring component listens to (part of) the communication between the connected components without affecting it.

- the connector *Uni-comm[s]* — the formal parameter;
- the refinement morphisms depicted below.

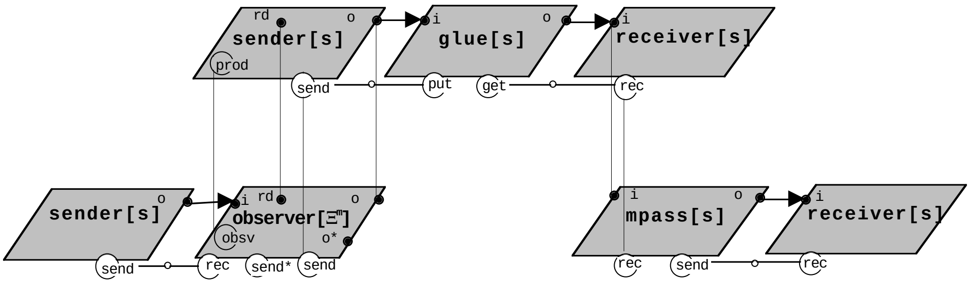

4 An ADL-independent notion of higher-order connector

The notion of higher-order connector presented for CommUnity can be generalised to other design formalisms. In this section, based on previous work [FLW00], we start by identifying the properties that such formalisms need to satisfy to support the architectural concepts and mechanisms that we have illustrated for CommUnity. Then, we shall present ADL-independent notions of connector and higher-order connector.

First, we need to fix a framework in which designs, configurations and relationships between designs, such as refinement, can be formally described. Our experience in formalising notions of structure in Computing, building on previous work of J.Goguen on General Systems Theory, suggests that, as illustrated in section 3, Category Theory provides a convenient framework for our purpose [FM96]. More concretely, we shall consider that a formalism supporting system design includes:

- a category ***c-DESC*** of component designs in which systems of interconnected components are modelled through diagrams;
- for every set *CD* of component designs, a set *Conf(CD)* consisting of all well-formed configurations that can be built from the components in *CD*. Each such configuration is a diagram in ***c-DESC*** that is guaranteed to have a colimit. Typically, *Conf* is given through a set of rules that govern the interconnection of components in the formalism.
- a category ***r-DESC*** with the same objects as ***c-DESC***, but in which morphisms model refinement, i.e., a morphism $\eta:S\rightarrow S'$ in ***r-DESC*** expresses that S' refines S, identifying the design decisions that lead from S to S'. Because the design of a composite system is given by a colimit of a diagram in ***c-DESC*** and, hence, is defined up to an isomorphism in ***c-DESC***, refinement morphisms must be such that designs that are isomorphic in ***c-DESC*** refine, and are refined exactly by, the same designs. Hence, it is required that *Isomorph(**c-DESC**)⊆Isomorph(**r-DESC**)*.

Summarising, all that we require is a notion of system design, a relationship between designs that captures components of systems, another relationship that captures refinement, and criteria for determining when a diagram of interconnected components is a well-formed configuration.

4.1 Architectural Schools

In the context of this categorical framework, we shall now present the properties of a design formalism for supporting the architectural concepts that we have illustrated for CommUnity. These properties define what we call an architectural school.

The categorical properties that a formalism needs to satisfy for supporting the notion of connector and its instantiation mechanism are identified and discussed in detail in [FLW00]. We shall summarise this characterisation and extend it in order to support higher-order connectors too.

Coordination

A key property of a formalism for supporting architectural design is that it provides a clear separation between the description of the individual behaviour of components and that of their interaction in the overall system organisation.

We shall take the separation between coordination and computation to be materialised through a functor ***sig: c-DESC→SIG*** mapping designs to signatures, forgetting their computational aspects. The fact that the computational side does not play any role in the interconnection of systems can be captured by the following properties of this functor:

- ***sig*** *is faithful;*
- ***sig*** *lifts colimits of well-formed configurations;*
- ***sig*** *has discrete structures;*

together with the following condition on the well-formed configuration criterion

- *given any pair of configuration diagrams* $\boldsymbol{dia}_1$, $\boldsymbol{dia}_2$ *s.t.* $\boldsymbol{dia}_1;\boldsymbol{sig}=\boldsymbol{dia}_2;\boldsymbol{sig}$, *either both are well-formed or both are ill-formed.*

The first condition states that morphisms of systems cannot induce more relationships than those that can be established between their underlying signatures. The second condition means that if we interconnect system components through a well-formed configuration, then any colimit of the underlying diagram of signatures establishes a signature for which a computational part exists that captures the joint behaviour of the interconnected components. The third condition implies that every signature θ has a realisation as a system component $\boldsymbol{desc}(\theta)$. In a sense, sources of morphisms in diagrams of designs are, essentially, signatures.

These three conditions ensure that any interconnection of systems can be established via their signatures, legitimising the use of signatures as channels in configuration diagrams. By requiring that any two configuration diagrams that establish the same interconnections at the level of signatures be either both well-formed or both ill-formed, the fourth property ensures that the criteria for well-formed configurations do not rely on the computational parts of designs.

In such situation, we say that *the formalism <**c-DESC**,Conf,**r-DESC**> is coordinated over **SIG** through the functor **sig***.

Compositionality

Another crucial property for supporting architectural design is in the interplay between structuring systems in architectural terms and refinement.

We start by noticing that, in order to support the refinement of an abstract description of a system, it must be possible to propagate the interactions between the components of the system when their designs are replaced by more concrete ones. This situation can be characterised by the existence, for every well-formed configuration $\boldsymbol{dia}$ involving designs $\{S_1,\ldots,S_k\}$ and every set of refinements morphisms $\{\eta_i{:}S_i\rightarrow S'_i{:}\ i\in 1..k\}$, of a well-formed configuration diagram $\boldsymbol{dia}+(\eta_i)_{i\in 1..k}$ obtained by "composing" in some way each refinement morphism η_i with the morphisms of $\boldsymbol{dia}$ whose target is S_i. The diagram $\boldsymbol{dia}+(\eta_i)_{i\in 1..k}$ describes the system obtained by replacing the designs of the components of the system (S_i) by more concrete ones (S'_i).

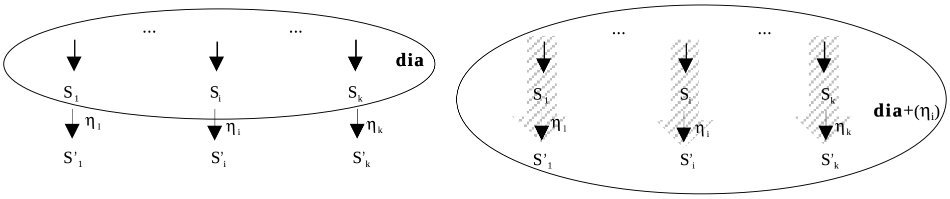

Naturally, a method of propagation of the interactions between the components of the system when their designs are replaced by more concrete ones is only significant if all decisions made previously are respected. In other words, the correctness criterion for this form of "configuration refinement" is that the colimit of ***dia***+(η_i)$_{i\in 1..k}$ provides a refinement for the colimit of ***dia***.

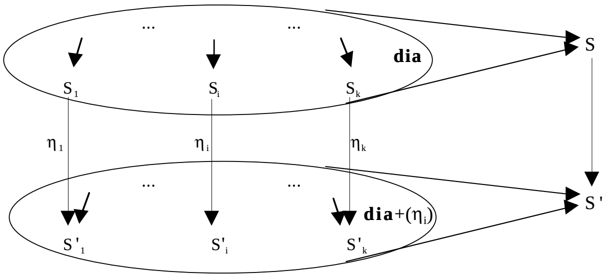

As explained in [FLW00], a formalism supports the notion of connector if it is coordinated and has a correct method of propagation of the interactions between the components of the system when their designs are replaced by more concrete ones.

To characterise the formalisms that support the notion of higher-order connector we have illustrated for CommUnity, it is necessary to know exactly which is the notion of configuration refinement of the formalism, i.e., in which situations a configuration is considered a refinement of another configuration. Given that configurations are made of components and interconnections, it is natural that a design formalism supports a notion of configuration refinement *CR* that, in addition to an operator + for the refinement of components, also allows the refinement of interconnections.

We require *CR* to be correct, i.e.,

> *For every **dia** and **dia'** s.t. **dia'** refines **dia** according to CR, the colimit of **dia'** provides a refinement for the colimit of **dia**.*

Furthermore, the configurations of the form ***dia***+(η_i) must be considered, according to *CR*, refinements of ***dia***. In this situation, we shall say that *the formalism is compositional w.r.t. CR*.

*In summary, a formalism F=<**c-DESC**,Conf,**r-DESC**> is called an architectural school over a functor **sig: c-DESC→SIG** and a configuration refinement notion CR if*

- *F is coordinated over* ***SIG*** *through* ***sig****;*
- *F is compositional w.r.t. CR.*

4.2 Connectors

Consider given an architectural school *F*=<***c-DESC***,***Conf***,***r-DESC***> over ***sig: c-DESC→SIG*** and *CR*. The generalisation of the notion of connector presented for CommUnity in section 3.2 is straightforward.

- *A connection consists of*
  - — *two designs G and R, called the glue and the role of the connection, respectively;*
  - — *a signature θ and two morphisms σ:**desc(θ)**→G,μ:**desc(θ)**→R in **c-DESC** connecting the glue and the role.*
- *A connector is a finite set of connections with the same glue that, together, constitute a well-formed configuration.*

$$\begin{tikzcd}
 & & G & & \\
\theta_1 \arrow[urr, "\sigma_1"] \arrow[d, "\mu_1"] & & \theta_i \arrow[u, "\sigma_i"] \arrow[d, "\mu_i"] & & \theta_n \arrow[ull, "\sigma_n"'] \arrow[d, "\mu_n"] \\
R_1 & & R_i & & R_n
\end{tikzcd}$$

- *The semantics of a connector is the colimit of the diagram formed by its connections.*

A connector can be applied to specific components of a system under construction, establishing the intended interactions between them, by instantiating the roles of the connector with those components. Role instantiation has to obey a compatibility requirement, which is expressed via the refinement morphisms of ***r-DESC***.

An instantiation of a connector is defined as follows:

*An instantiation of a connection with role R consists of a design P together with a refinement morphism η:R→P in **r-DESC**.*

$$\begin{tikzcd}
 & & G & & \\
\theta_1 \arrow[urr] \arrow[d, "\mu_1"] & & \theta_i \arrow[u] \arrow[d, "\mu_i"] & & \theta_n \arrow[ull] \arrow[d, "\mu_n"] \\
R_1 \arrow[d, "\eta_1"] & & R_i \arrow[d, "\eta_i"] & & R_n \arrow[d, "\eta_n"] \\
P_1 & & P_i & & P_n
\end{tikzcd}$$

An instantiation of a connector consists of an instantiation for each of its connections such that the diagram in ***c-DESC*** *connecting the role instances to the glue, obtained by composing the role morphism of each connection with its instantiation (given by* $<\sigma_i,\mu_i> +\eta_i$*), constitutes a well-formed configuration.*

The semantics of a connector instantiation is the colimit of the diagram in ***c-DESC*** *formed as described above.*

The compositionality of the design formalism ensures that the system that results from a instantiation of a connector *C* refines the semantics of *C*. In this way, the properties of connectors can be understood independently of specific contexts in which they are used.

4.3 Higher-Order Connectors

The notion of higher-order connector, as a connector that takes one connector as parameter and delivers another as a result, can be defined as follows.

- *A higher-order connector (hoc) consists of*
  - *a connector pC, called the formal parameter of the hoc; its roles, glue and connections are called, respectively, the parametric roles, the parametric glue and the parametric connections of the hoc;*
  - *a connector C – its roles and glue are also called the roles and the glue of the hoc;*
  - *an instantiation of the formal parameter connector with the glue of the hoc, i.e., a refinement morphism* η_i *from each of the parametric roles to the glue, such that the diagram in* ***c-DESC*** *obtained by composing the role morphism of each parametric connection with its instantiation*

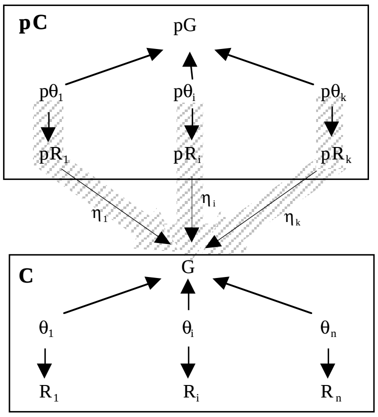

constitutes a well-formed configuration.

- *The semantics of a higher-order connector is the connector depicted below. Its roles are the roles of C and its glue is G', a design returned by the colimit of the configuration pC+(*η_i*).*

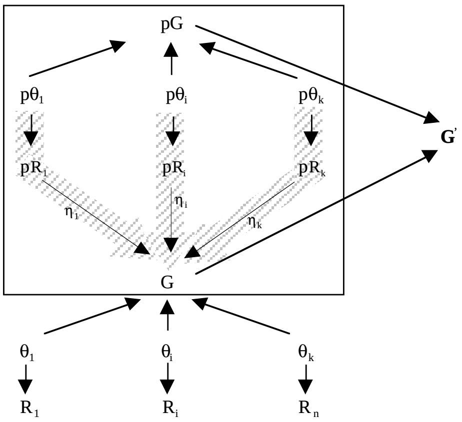

For simplicity, we have imposed one single parameter to the higher-order connector. However, the definition can be extended to the case of several parameters in a straightforward way.

Intuitively, the instantiation of the formal parameter of a higher-order connector can be regarded as the replacement of a connector (the formal parameter *pC*) that was instantiated to given components of a system (the glue of the hoc) by another connector (the actual parameter). In addition, the type of interconnection that *pC* ensures must be preserved. In other words, the design that results from the replacement must be a refinement of the design from which we started.

Like for connectors, the instantiation of the formal parameter of a higher-order connector is established via a fitting morphism from the formal to the actual parameter. These morphisms, on the one hand, formulate the correspondence between roles and glue of the formal with those of the actual parameter and, on the other hand, capture conditions under which the "functionality" of the formal parameter is preserved.

In order to be able to use, in the design of a given system, a connector *C* in place of a connector *C'*, it is obvious that the two connectors must have the same number of roles. Furthermore, *C'* has to admit to be instantiated with the same components than *C*. That is to say, every restriction on the components to which *C'* can be applied must also be a restriction imposed by *C*. In this way, fitting morphisms must require that each of the roles of *C'* is refined by the corresponding role of *C*.

For instance, we cannot replace a connector *Sync* by a *Monitoring* at once. We first have to provide the component that will play the role *monitor* in the system and then encapsulate this component by making it part of the glue (see [WF98b]).The resulting connector has only two roles, which are exactly the same of *Sync*, and hence, it can be used as a *Sync* connector.

As shown in section 3, namely with the connector *Uni-comm*, connectors may be based on glues that are not fully developed as designs (may be underspecified) and, nevertheless, the concrete commitments that have already been made determine in some extent the type of interconnection that the connector will ensure. The type of interconnection is clearly preserved if we simply consider a less unspecified glue, i.e., if we refine the glue. Hence, fitting morphisms must allow for arbitrary refinements of the glue.

Having this in mind, we arrive at the following notion of fitting morphism:

- *A fitting morphism ϕ from a connection $<\sigma_1:\mathbf{desc}(\theta_1)\rightarrow G_1,\mu_1:\mathbf{desc}(\theta_1)\rightarrow R_1>$ to a connection $<\sigma_2:\mathbf{desc}(\theta_2)\rightarrow G_2,\mu_2:\mathbf{desc}(\theta_2)\rightarrow R_2>$ consists of a pair $<\phi_G: G_1\rightarrow G_2,\phi_R: R_2\rightarrow R_1>$ of refinement morphisms in* ***r-DESC*** *s.t. the interconnection $<\sigma_1,\mu_1>+\phi_G$ of R_1 with G_2 is, according to the configuration refinement CR, refined by the interconnection $<\sigma_2,\mu_2>+\phi_R$.*

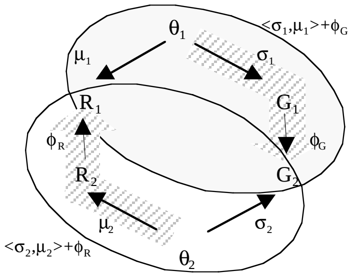

- *A fitting morphism ϕ from a connector C_1 to a connector C_2 with the same number of connections consists of a fitting morphism ϕ from each of C_1's connections to each of C_2's connections, all with the same glue refinement ϕ_G.*

If there exists a fitting morphism from a connector C_1 to a connector C_2, then we may replace each occurrence of the connector C_1 in an architectural description of a system by an occurrence of C_2. The compositionality of the design formalism w.r.t. the configuration refinement *CR* ensures that every coordination decision made previously is preserved.

Based on fitting morphisms between connectors, we define an instantiation of a higher-order connector.

- *An instantiation of a higher-order connector with formal parameter pC (figure 4.1) consists of a connector* C^A *(the actual parameter) together with a fitting morphism* $\phi: pC \rightarrow C^A$*, such that the diagram in* ***c-DESC*** *obtained by composing the role morphisms of each actual connection with the corresponding fitting component and then with the role instantiation (figure 4.2) constitutes a well-formed configuration.*

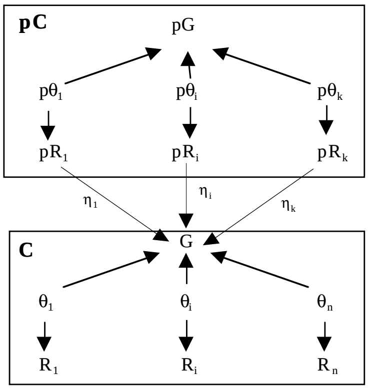

Figure 4.1

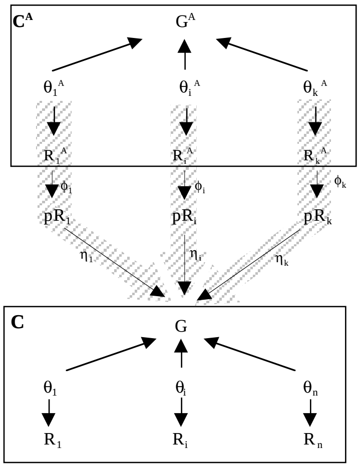

Figure 4.2

- *The semantics of a higher-order connector instantiation is the connector with the same roles as C and its glue is a design returned by the colimit of the configuration* $C^A + (\phi_i;\eta_i)$*.*

5 Composition of Higher-Order Connectors

Higher-order connectors facilitate the separation of concerns in the development of complex connectors and their compositional construction. For instance, we have seen that compression and monitoring can be modelled separately as higher-order connectors. Although we have not shown it, it is not very difficult to realize that compression can be applied to a connector that models a unidirectional communication protocol and then monitoring can be applied to the resulting connector.

An important feature of our notion of higher-order connector is that different kinds of functionality, modelled separately by different higher-order connectors, can be combined, giving rise also to a higher-order connector. In this way, it is possible to analyse the properties that such compositions exhibit, namely to investigate whether undesirable properties emerge and desirable properties are preserved.

The key idea for composition of hocs is the instantiation of a hoc with a hoc. In this section we shall present this more general form of instantiation — *parameterised instantiation*. So, for instance, *Monitoring(Uni-comm)* can be instantiated with *Compress(Uni-comm)*, giving rise to the hoc *Monitoring&Compress(Uni-comm)* which corresponds to a form of composition of *Monitoring* and *Compress* in which the messages are first observed, and possibly transmitted to the monitoring component, then are compressed and finally are transmitted via *Uni-comm*.

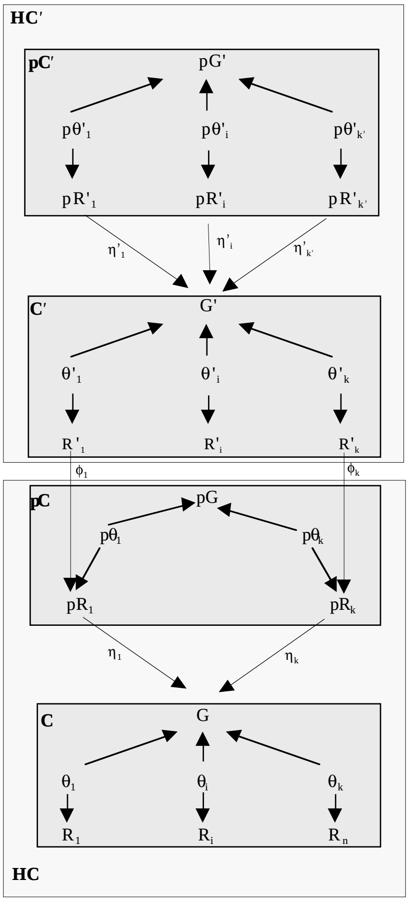

Figure 5.1

The definition of parameterised instantiation of a hoc is similar to the definition of the standard instantiation.

- *A parameterised instantiation of a higher-order connector HC with formal parameter pC consists of a higher-order connector HC' together with a fitting* *morphism $\phi:pC\rightarrow Con'$ (figure 5.1), where Con' is the connector that gives the semantics of HC', such that it is possible to extend, in a unique way, the instantiation of pC' with G' to an instantiation of pC' with the colimit of the diagram $C'+(\phi_i;\eta_i)_{i\in 1..k}$ (figure 5.2), which connects the glues of HC' and HC.*

In figure 5.2, G^{new} is the colimit of the diagram $C'+(\phi_i;\eta_i)_{i\in 1..k}$ and we have used dotted lines for the refinement morphisms whose existence we are requiring.

- *The semantics of a parameterised instantiation is the higher-order connector depicted in figure 5.3.*

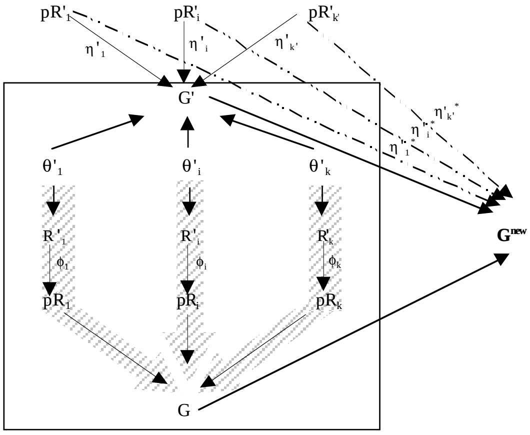

Figure 5.2

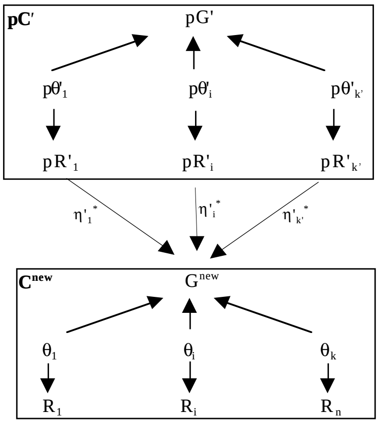

Figure 5.3

For instance, consider that we want to combine the service of compression of messages with monitoring. If we consider the parameterised instantiation of *Monitoring(Uni-comm)*$[\Xi^m]$ with *Compress(Uni-comm)*$[\Xi^{cd}]$ defined by the fitting morphism

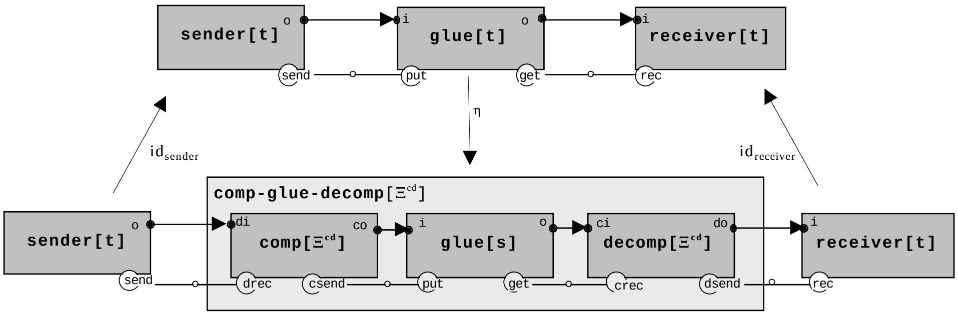

where the refinement morphism

$$\eta{:}glue[t]\rightarrow comp\text{-}glue\text{-}decomp[\Xi^{cd}]$$

is defined by

$\eta(i)=\sigma_c(di)$, $\eta(o)=\sigma_d(do)$,

$\eta(\sigma_c^{-1}(drec))=put$, $\eta(\sigma_d^{-1}(dec))=prod$, $\eta(\sigma_d^{-1}(dsend))=get$

where σ_c: *comp→comp-glue-decomp*$[\Xi^{cd}]$ and σ_d: *decomp→comp-glue-decomp*$[\Xi^{cd}]$ are the morphisms in ***c-DSGN*** returned by the colimit of the diagram

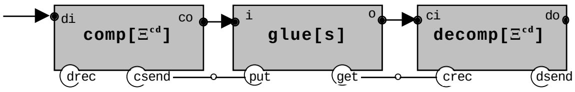

This composition gives rise to the hoc *Monitoring&Compress(Uni-comm)*$[\Xi^{cd}+\Xi^m]$ that is constituted by

- the connector *Monitoring&Compress*$[\Xi^{cd}+\Xi^m]$ defined in figure 5.4;
- the connector *Uni-comm[s]* — the formal parameter;
- the refinement morphisms

  η_s*:sender[s]→obs-c-d-mpass*$[\Xi^{cd}+\Xi^m]$ and η_r*:receiver[s]→obs-c-d-mpass*$[\Xi^{cd}+\Xi^m]$

  obtained by composing, at the level of signatures, the morphisms *sender[s]→comp*$[\Xi^{cd}]$ and *receiver[s]→decomp*$[\Xi^m]$ of *Compression(Uni-comm)*$[\Xi^{cd}]$ with the morphisms, respectively, *comp*$[\Xi^{cd}]$*→obs-c-d-mpass*$[\Xi^{cd}+\Xi^m]$ and *decomp*$[\Xi^{cd}]$*→obs-c-d-mpass*$[\Xi^{cd}+\Xi^m]$ which are given by the colimit construction (for the sake of space, we omit the proof of the fact that these signatures morphisms do define refinement morphisms).

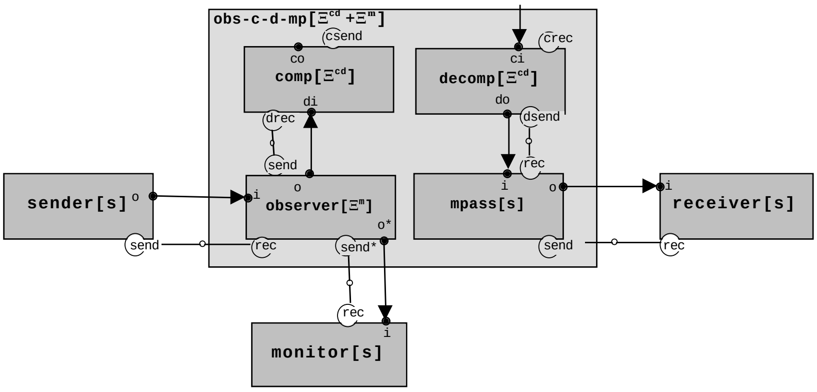

Figure 5.4

It is not difficult to realize this hoc works as described before: first messages are observed and possibly transmitted to the monitoring component, then are compressed and finally are transmitted via *Uni-comm*.

6 Related Work

Our definitions agree with Garlan's original proposal [G98] of a hoc as an operator over connectors for supporting connector construction through incremental transformation, hence allowing one to define more complex interactions in a more systematic way. More concretely, Garlan and Spitznagel [SG01] propose that a connector transformation be modelled as a function — from one or more connectors to a new connector — defined in terms of its inputs, preconditions on its application and postconditions on its result. They formalise these ideas in the context of a particular ADL, namely Wright, relying on the specific language and semantics of CSP.

Adaptation of connectors vs. adaptation of components. As explained, hocs may be used to represent connector adaptation and, in particular, the installation of additional services, such as security or fault-tolerance. For the same purpose, a radically different approach is to apply component adaptation, for instance by using wrappers or packaging (e.g., [K93, B99, DMT99]).

There are several reasons to point out in favour of our approach. For instance, as argued in [G98], it is not always possible to adapt components to work with the existing connectors. Even in those cases where it is feasible, a better alternative might be to modify the connectors because usually there are fewer connector types than components types. Despite the differences, in our setting, the two approaches are equivalent in a sense that we shall explain.

Consider again the compression example. As explained before, the instantiation of the hoc *Compression* with connector *Async* gives rise to the following configuration

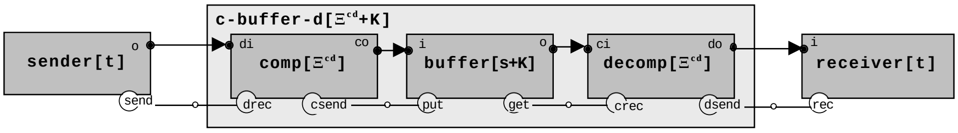

which is equivalent to the next configuration. This is because, from a categorical point of view, the two configurations correspond to two different ways of calculating the colimit of the same diagram and, hence, they give rise to equivalent designs (i.e., designs considered isomorphic in ***c-DESC***).

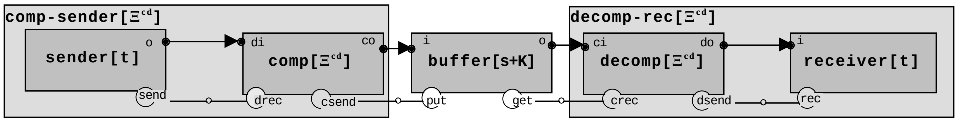

Clearly, this configuration represents a system where the compression service was installed by means of adding (different) wrappers to the sender and receiver and the connector *Async* is again used to connect the two components.

Integrating extra-functional requirements in connectors. The notion of architectural connector in the style defined by Allen and Garlan is also the basis for a completely different approach to the specification of extra-functional properties, such as security and fault-tolerance, of software architectures descriptions. Issarny and colleagues [IBS98] propose an extension of architectural connectors with a set of first-order formulas specifying the extra-functional properties offered by the connector. This extension relies on the specific language and semantics of CSP and is based on the fact that the behaviour of the protocol can be formally defined in terms of the predicates according to Hoare's logic. These predicates essentially characterize the coordination actions that are carried out by the protocol. This approach is mainly tailored to analyse whether a given architecture has some desired extra-functional properties whereas our approach is geared towards system construction.

Behavioural analysis. Many existing ADLs have some associated technique or tool to analyse the resulting behaviour of the designed system. For example, ADLs that use process calculi to specify components and connectors typically use some model-checking tool: Darwin uses LTSA [MKG99] and Wright uses FDR [AG97], just to mention two better known examples. Our current support for behavioural analysis is twofold.

On the one hand, the CommUnity Workbench [WO02] is being developed (in Java) as a proof of concept of the theoretical framework, hiding the underlying "mathematical machinery" from the user. Currently, the tool provides a graphical integrated development environment to write CommUnity programs (with fixed data types), define connectors, draw an architecture, calculate automatically its colimit, and run it. The workbench prevents the creation of ill-formed configurations (like binding output channels with each other) and gives great flexibility in testing CommUnity programs (like channel initialisation, choice of which actions and channels to trace, and verification of invariants during execution). Higher-order connectors and some advanced features of CommUnity are not yet supported.

On the other hand, as mentioned in the end of Section 3.2, we defined a logic formalism for expressing the properties of CommUnity designs, namely the co-operation properties of a design in regard to its environment. We also developed a proof method for reasoning about CommUnity designs in a compositional way, i.e., that allows us to reasoning about a system described through a configuration diagram, without requiring the calculation of the colimit design. Hence, it can be applied to the description of connectors and hocs.

In future work we would like to include an implementation of the logic, possibly using an existing theorem prover, into the CommUnity Workbench.

7 Concluding Remarks

In this paper we continued our previous work on providing formal support to the definition and use of architectural techniques for software development. Building on the categorical semantics for the notion of architectural connector that we have developed in past papers (e.g. [FL97]), we formalised a notion of higher-order connector that can be used for defining new connectors from existing ones by superposing aspects like security, fault-tolerance, etc.

We showed that a transformation of a connector can be modelled by a parameterised entity that is essentially constituted by two connectors. One of these connectors is the formal parameter that defines the kind of connectors the transformation can be applied to. The other connector — the body of the hoc — concerns the transformation itself. Owing to the formal semantics of hocs, the transformation can be understood and analysed. By adopting a categorical framework, we achieved independence relative to the ADL. Although we first used the CommUnity language to illustrate our ideas, we then presented a generalisation of these ideas which is applicable to any language that supports architectural design (in a sense that was made precise in section 4.1). Moreover, the categorical framework led naturally to a notion of composition of higher-order connectors that turned out to be useful for combining orthogonal properties.

As explained in [FM96], the use of Category Theory in the paper was not an end in itself but, rather, a means of characterising the proposed concepts and techniques in a way that, on the one hand, is independent of the way it can be offered to users in specific ADLs and, on the other hand, is amenable to formal analysis. As a result, we are able to compare how different ADLs support these notions and, if they do not, to suggest ways in which they can be extended.

On a more practical side, we are now in the process of transposing these results to the architectural development environment that we have built for Java [GKAF01]. This environment is based on a micro-architecture that transposes, into Java, the separation between Computation and Coordination that we formalised with the notion of architectural school [AFGLW00]. The idea is to support the construction of new coordination contracts through the higher-order mechanisms that we have characterised.

For simplicity, we defined hocs with one parameter only, but the extension to several parameters is straightforward. However, this fact limited the kind of examples we have used throughout the paper, namely it prevented us from showing that hocs can be used to model operators that represent more than an adaptation of connectors. For instance, it is necessary to consider an *n*-ary hoc in order to model an operation on *n* connectors so that the output of the first connector goes into the input of the second connector, etc. Another operation that requires more parameters is the aggregation of *n* connectors [SG01], a combination of *n* connectors with a controller that determines which connector is active at a time. Such an operation can be useful for systems with transient interactions, e.g., due to mobility of the components. As mentioned before, the individual specification of independent aspects such as compression and fault-tolerance as higher-order connectors makes it easier to evolve systems at run-time. Through run-time reconfiguration of the system architecture, namely through the replacement of connectors, such services may be added only when necessary, hence preventing performance penalties when such complex interactions are not required. Previous work addressed the support that is required for an architectural-driven process of reconfiguration in which connectors, as well as components, can be replaced, added or deleted. In particular, we have developed a high-level language for specifying runtime architectural changes [WLF01], based on the semantic domain of graph rewriting.

References

[AG97] R.Allen and D.Garlan, "A Formal Basis for Architectural Connectors", *ACM TOSEM*, 6(3):213-249, July 1997.

[AFGLW00] L.Andrade, J.Fiadeiro, J.Gouveia, A.Lopes and M.Wermelinger, “Patterns for Coordination", in *Proc. COORDINATION'00*, G.Catalin-Roman and A.Porto (eds), Lecture Notes in Computer Science 1906, pp. 317-322, Springer-Verlag 2000

[B99] J.Bosch, "Superimposition: A Component Adaptation Technique", *Information and Software Technology* 1999.

[BCK98] L. Bass, P.Clements and R.Kasman, *Software Architecture in Practice*, Addison Wesley 1998.

[CM88] K.Chandy and J.Misra, *Parallel Program Design - A Foundation*, Addison-Wesley 1988.

[DMT99] G.Denker, J.Meseguer and C.Talcott, "Rewriting semantics of meta-objects and composable distributed services", Internal report, Computer Science Laboratory, SRI International, 1999.

[EM85] H.Ehrig and B.Mahr, *Fundamentals of Algebraic Specification I: Equations and Initial Semantics*, Springer-Verlag, 1985.

[FF96] N.Francez and I.Forman, *Interacting Processes*, Addison-Wesley 1996.

[FL97] J.L.Fiadeiro and A.Lopes, "Semantics of Architectural Connectors", in *TAPSOFT'97*, LNCS 1214, pp. 505-519, Springer-Verlag 1997.

[FL99] J.L.Fiadeiro and A.Lopes, "Algebraic Semantics of Coordination, or what is in a signature?", in *AMAST'98*, A.Haeberer (ed), Springer-Verlag 1999.

[FLW00] J.L.Fiadeiro, A.Lopes and M.Wermelinger, "A Mathematical Semantics for Architectural Connectors". Submitted for publication.

[FM95] J.L.Fiadeiro and T.Maibaum, "Interconnecting Formalisms: supporting modularity, reuse and incrementality", in G.E.Kaiser (ed) *Proc. 3rd Symp. on Foundations of Software Engineering*, pp. 72-80, ACM Press 1995.

[FM96] J.L.Fiadeiro and T.Maibaum, "A Mathematical Toolbox for the Software Architect", in J.Kramer and A.Wolf (eds) *Proc. 8th International Workshop on Software Specification and Design*, pp. 46-55, IEEE Computer Society Press 1996.

[FM97] J.L.Fiadeiro and T.Maibaum, "Categorical Semantics of Parallel Program Design", *Science of Computer Programming* 28:111-138, 1997.

[G73] J.Goguen, "Categorical Foundations for General Systems Theory", in F.Pichler and R.Trappl (eds) *Advances in Cybernetics and Systems Research*, pp. 121-130, Transcripta Books 1973.

[G89] J.Goguen, "Principles of Parametrised Programming", in Biggerstaff and Perlis (eds) *Software Reusability*, pp. 159-225, Addison-Wesley 1989.

[G96] J.Goguen, "Parametrised Programming and Software Architecture", in *Symposium on Software Reusability*, IEEE 1996.

[G98] D.Garlan, "Higher-order connectors", Position paper for the Workshop on Compositional Software Architectures, January 1998.

[GC92] D.Gelernter and N.Carriero, "Coordination Languages and their Significance", *Communications ACM* 35(2):97-107, 1992.

[GKAF01] J.Gouveia, G.Koutsoukos, L.Andrade and J.Fiadeiro, “Tool Support for Coordination-Based Software Evolution", in *Technology of Object-Oriented Languages and Systems – TOOLS 38*, W.Pree (ed), pp. 184-196, IEEE Computer Society Press 2001

[H85] C.A.R.Hoare, *Communicating Sequential Processes*. Prentice-Hall, 1985.

[HUY99] D.Hirsch, S.Uchitel and D.Yankelevich, "Towards a periodic table of connectors", Proc. of Simposio en Tecnología de Software, Buenos Aires, 1999.

[IBS98] V.Issarny, C.Bidan and T.Saridakis, “Characterizing coordination architectures according to their non-functional execution properties”, Proc. 31st Annual Hawaii International Conference on System Sciences, pp 275-283, Jan 1998.

[K93] S.Katz, "A Superimposition Control Construct for Distributed Systems", *ACM TOPLAS* 15(2):337-356, 1993.

[L99] A.Lopes, "Não-determinismo e Composicionalidade na Especificação de Sistemas Reactivos", PhD Thesis (in Portuguese), Universidade de Lisboa, Jan. 1999.

[LF99] A.Lopes and J. L. Fiadeiro, "Using explicit state to describe architectures", in E. Astesiano (ed), *FASE'99*, LNCS 1577, pp. 144–160, Springer-Verlag 1999.

[MKG99] J.Magee, J.Kramer and D.Giannakopoulou, "Behaviour analysis of software architectures", *Software Architecture*, pp. 35-50, Kluwer Academic Publishers 1999.

[MMP00] N.Mehta, N.Medvidovic and S.Phadke, "Towards a taxonomy of software connectors", Proc. of the 22nd Intl. Conf. on Software Engineering, pp. 178-187, ACM Press, 2000.

[OP92] S. W. O'Malley and L. L. Peterson, "A Dynamic Network Architecture", *ACM Transactions on Computer Systems*, 10(2):110-143, May 1992.

[PW92] D.Perry and A.Wolf, "Foundations for the Study of Software Architectures", *ACM SIGSOFT Software Engineering Notes* 17(4):40-52, 1992.

[S93] M.Shaw, "Procedure calls are the assembly language of system interconnection: Connectors deserve first-class status", *Proc. of the Workshop on Studies of Software Design*, May 1993.

[S95] M.Shaw, R.DeLine, D.V.Klein, T.L.Ross, D.M.Young, and G.Zelesnik, "Abstractions for software architecture and tools to support them", *IEEE Trans. on Software Engineering*, 21(4):314-335, April 1995.

[SG01] B.Spitznagel and D.Garlan, "A Compositional Approach for Constructing Connectors", The Working IEEE/IFIF Conference on Software Architecture (WICSA’01), Royal Netherlands Academy of Arts and Sciences Amsterdam, The Netherlands, August , 2001.

[WF98a] M.Wermelinger and J.L.Fiadeiro, "Connectors for Mobile Programs”, *IEEE Trans. on Software Engineering* 24(5):331-341, 1998.

[WF98b] M.Wermelinger and J.L.Fiadeiro, "Towards an Algebra of Architectural Connectors: a Case Study on Synchronisation for Mobility", in *Proc. 9th International Workshop on Software Specification and Design*, pp. 135-142, IEEE Computer Society Press 1998.

[WLF00] M.Wermelinger, A.Lopes and J.L.Fiadeiro, "Superposing Connectors", in *Proc. 10th International Workshop on Software Specification and Design*, pp. 87-94, IEEE Computer Society Press 2000.

[WLF01] M.Wermelinger, A. Lopes, and J.L.Fiadeiro, "A Graph Based Architectural (Re)configuration Language", in *Proc. of ESEC/FSE'01*, pp. 21-32, ACM Press 2001.

[WO02] M. Wermelinger and C. Oliveira. "The CommUnity Workbench". Proc. of the 24th Intl. Conf. on Software Engineering, p. 713, ACM Press, May 2002.