



Combining Data Sharing with the Master-Worker Paradigm in the Common Component Architecture

Gabriel Antoniu, Hinde Lilia Bouziane, Mathieu Jan, Christian Pérez, Thierry Priol

► To cite this version:

Gabriel Antoniu, Hinde Lilia Bouziane, Mathieu Jan, Christian Pérez, Thierry Priol. Combining Data Sharing with the Master-Worker Paradigm in the Common Component Architecture. Joint Workshop on HPC Grid programming Environments and COmponents and Component and Framework Technology in High-Performance and Scientific Computing (HPC-GECO/CompFrame 2006), Jun 2006, Paris, France. pp.10-18, 2006. <inria-00101365>

HAL Id: inria-00101365

<https://hal.inria.fr/inria-00101365>

Submitted on 26 Sep 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Combining Data Sharing with the Master-Worker Paradigm in the Common Component Architecture

Gabriel Antoniu, Hinde Lilia Bouziane, Mathieu Jan, Christian Pérez, and Thierry Priol
IRISA/INRIA, Campus de Beaulieu, 35042 Rennes cedex, France
Contact: Christian.Perez@irisa.fr

Abstract—Software component technologies are being accepted as an adequate solution for handling the complexity of applications. However, existing software component models tend to be specialized to some types of resource architectures (e.g. in-process, distributed environments, etc) and/or do not provide a very high level of abstraction. This paper focuses on handling *data sharing* on operation invocations between components as a solution allowing applications to be efficiently executed on all kinds of resources. In particular, the data sharing pattern appears in master-worker applications, when workers need to access only a part of a large piece of data, either in read or write mode. This approach is applied to the Common Component Architecture model. Its benefits are discussed using an image rendering application.

I. INTRODUCTION

e-Science application programmers are facing a serious challenge for the following years, due to both the intricacy of these applications and the increasing variety of computing resources. For instance, resources may be multicore multithread multiprocessors, clusters, or grid computers. To cope with these two concerns, it is necessary to provide programming models able to make the programming of these applications simple, independent of the target computing infrastructures, while preserving high-performance. Software component models aim to handle the increasing complexity of today's applications [1]. However, component models need to be adapted for addressing high performance *independently* of the resource infrastructure. In previous work, we have addressed several shortcomings of these models: encapsulation of a parallel code into a component [2], support for the master-worker programming paradigm [3], and support for data sharing among components [4].

Communication between components is to a large extent based on the exchange of messages, which may be associated with control transfers (RPC or RMI) or with data transfers (events, streams, etc). For many applications, especially e-Science ones, this communication model is not satisfactory because it requires to explicitly compute which pieces of data need to be exchanged. When data structures are complex and sparse in memory, or when data access patterns are irregular, this communication model is not suitable nor efficient: a shared memory communication model seems more appropriate. The challenge is thus to support this communication model while keeping the good properties of software components, that is to say: the composition, the port-based communications, the deployment unit, etc.

This paper aims at showing how an existing component model, the Common Component Architecture (CCA) [5], can be extended to support shared data in general, and focuses more particularly on the master-worker paradigm. To this end, it fills the gap between data sharing and operation invocations. The proposed approach is illustrated with an image rendering application that has been chosen as it requires to share data and exhibits a master-worker pattern.

Section II describes the state-of-the art with software component models in general, and briefly summarizes our previous works on the use of the transparent data access paradigm and on the master-worker paradigm within component models. Section III presents CCA and proposes extensions to the CCA model, in order to support the two concepts, data sharing and master-worker paradigm. Section IV describes how data sharing can be enabled on operation invocation, first independently of any component model, then in the CCA model. Its benefits with respect to an image rendering application are discussed in Section V. Section VI concludes the paper and outlines some future works.

II. SOFTWARE COMPONENT MODELS

A software component according to Szyperski [1] *is a unit of composition with contractually specified interfaces and explicit context dependencies only. A software component can be deployed independently and is subject to composition by third parties.* This definition considers a component as a black box able to be composed with other components. This composition is possible thanks to well-defined ports that allow components to interact. A port may express the fact (1) that a component needs (uses) some functionality provided by another component, or (2) that it provides itself some functionality. A component-based application is then built by composing multiple components through an assembly process. Its execution can be launched after installing binary codes of all components on a given set of resources. During the deployment process, some properties associated to a component should be considered to perform a convenient resource choice. Such properties may be operating systems, processors and amount of memory requirements.

Recently, several component models (e.g. CCM [6], CCA [5], FRACTAL [7], GRID.IT [8], DARWIN [9]), etc.) have been proposed. Their aim is to facilitate the design of applications and reduce the complexity of their building process. However, to reach this goal, component models should be

able to support most of distributed application paradigms in an easy way, which is currently not the case. This section recalls previous work about the enhancement of component models with the support of two paradigms: (1) the *data sharing* paradigm and (2) the *master-worker* paradigm.

A. Transparent data access model and component models

1) *Enabling transparent access to data*: an attractive programming paradigm allowing data to be shared by multiple concurrent entities is the *shared memory* paradigm. Its advantage relies on the ease of programming: multiple processes can read/write data in a global space without any need to explicitly handle data location. This concept has been successfully applied in several contexts: (1) multithreading within the same process, (2) data segment sharing among multiple processes running on the same host, or (3) global data sharing across a cluster of workstations through Distributed Shared Memory (DSM) systems. However, it has not really been exploited in grid environments. Currently, the most widely-used approach to manage data on distributed environments (and on grid testbeds in particular) relies on the *explicit data access model*, where clients have to move data to computing servers. In order to achieve a real virtualization of the management of large-scale distributed data, a step forward has been made by enabling a *transparent data access model* through the concept of *grid data-sharing service* [10]. Such a service transparently manages data localization and persistence in a dynamic, large-scale, distributed environment. The data sharing service concept is based on a hybrid approach inspired by DSM systems and peer-to-peer (P2P) systems. The service provides data consistency guarantees based on fault-tolerant protocols [11].

The concept of data-sharing service is illustrated by the JUXMEM software experimental platform, described in detail in [10]. Its implementation relies on the JXTA [12] generic P2P framework. The JUXMEM API provides to users classical functions to allocate and map/unmap memory blocks in a globally shared space: `juxmem_malloc`, etc. The memory allocation operation returns a global data ID. This ID can be used by other nodes in order to access existing data through the use of the `juxmem_mmap` function. To obtain read and/or write access on a data, a process that uses JUXMEM should acquire the lock associated to the data through either `juxmem_acquire` or `juxmem_acquire_read`. This allows the implementation to apply consistency guarantees according to the consistency protocol specified by the user at allocation time. The choice of a C-style malloc interface for the API of JUXMEM is motivated by the targeted e-Science applications of JUXMEM: they mainly share arrays of data.

2) *Data sharing and component models*: in current component models, *ports* are defined based on the assumption of an explicit communication operation between two components. As such, components are only able to deal with data as a part of a message actually exchanged between two components. As explained earlier, sharing data among multiple components may be more appropriate for some applications where data structures are more complex and access patterns are more

```
interface AccessPort {
    float* get_pointer();
    long get_size();
    // Synchronization primitives
    void acquire();
    void acquire_read();
    void release();
}
```

Fig. 1. An interface offered to the programmer by data ports. The shared data is an array of float.

irregular. Implementing such a functionality using “classical” ports inside a component-based application is possible. In this case, the data is physically located into a component that provides it. However, with such a centralized approach, the component storing the data can easily produce a bottleneck as the number of concurrent accesses increases. Another possibility is to have a copy of the shared data on each component that uses it. In such a case, the functional code of a component would have to maintain a consistent state of all copies, each time the data is updated. For example, this can be achieved through the use of a consensus algorithm. Consequently, the management of synchronizations and concurrent accesses to data would be handled within the functional code of components, leading to an unnecessary increase in the complexity of applications.

To summarize, existing software component models do not efficiently support a *transparent data access* model. We claim that this is a limitation for current component models as data persistence, consistency and fault-tolerance are not handled.

3) *Data port model*: in a previous work [4], a family of ports named *data ports* was proposed to logically attach a shared data to a component. Two kinds of data ports were defined: a *shares* port to give an access to a shared data and an *accesses* port to enable a component to access a data exported through a *shares* port. Such ports relies on the *transparent data access model* described in Section II-A.1.

In order to access a data, an interface named *AccessPort* is implicitly associated to a data port. The API of this interface is shown in Figure 1. This interface is available through *accesses* ports as well as through *shares* ports. Indeed, a component that *shares* some data may also need to access the data. The interface provides `get_pointer/get_size` primitives to respectively retrieve a pointer to the shared data and its size. This interface is currently enough for scientific codes, mainly written in FORTRAN, that typically handle arrays. It also provides synchronization primitives, like `acquire` and `release`. The `acquire_read` primitive sets a lock in read-only mode so that multiple readers can simultaneously access a given data, whereas `acquire` sets a lock in exclusive mode.

A component which aims to access a data through an *accesses* port should have this port connected to a *shares* one. Connecting the two ports implies passing the reference of the shared data from the *shares* component to the *accesses* component. The shared data was previously allocated and associated to the *shares* port by an appropriate interface provided only on the *shares* port side.

Let us stress that a data port allows a data to be shared between components without worrying the user with the mechanism used to share the data. Such a mechanism is expected to be the memory between components collocated within the same process, shared memory segment for components in two different processes but on the same host, DSM for cluster and grid data-sharing service like JUXMEM for grids. It is the responsibility of the component model framework not of the component implementation.

We projected this model on the CORBA Component Model and we realized a prototype implementation as a proof of the concept and of the facilities offered to the programmer. JUXMEM was in particular used to share data between components located on distinct clusters. An application example using data ports can be found in [4].

B. Master-worker paradigm and component models

1) The importance of the master-worker paradigm:

the master-worker programming paradigm is widely used in distributed applications. Parametric applications for instance, are based on such a paradigm where, several workers execute simultaneously a same code but with different parameter values. Numerous research activities are dealing with the design of *master-worker* software grid-enabled environments such as for global computing systems (SETI@Home [13], XtremWeb [14], BOINC [15], etc.) or for network-enabled server environments (DIET [16], NetSolve [17], Nimrod/G [18]). However, these grid-enabled environments only focus on supporting the master-worker paradigm. Therefore, such environments seem not to be convenient to support an application partially based on the master-worker paradigm. If such an environment is not used, a programmer has the burden of managing workers. Consequently, the code complexity is increased by these non functional concerns. Moreover, it is further increased since a programmer is also required to implement request transport policies from masters to workers. These policies may be very complex as they may depend on the underlying execution environments. These drawbacks remain valid for components based applications.

2) *Master-worker paradigm and component models*: previous work [3] aimed at extending software component models to increase the abstraction level for *master-worker* applications. The proposal only requires a designer to specify a set of *worker* instances to which a *master* component is connected. Request transport concerns are handled separately while advanced transport policies are possible. Ideally, existing master-worker environments may be utilized. Last, the number of *workers* is handled as a non functional property and thus may be delegated to an adaptivity service.

Figure 2 presents an overview of the different elements of the proposal. The application designer specifies a *collection* of *worker* components. A collection is a set of *exposed* ports, bound to some internal component ports. It is described with an *abstract architecture description*. Independently, request transport *patterns* are defined by some experts. They represent




Fig. 2. An overview of a *Master-Worker* application model.

request delivery policies that may be used between *master* and *worker* components. They should be based on software components, even though existing *master-worker* environments such as DIET [16] may be used.

Once the deployment environment is known, an initial number of worker components and a suitable requests transport policy can be decided. From these choices, the abstract architecture description is converted into a concrete ADL description during a *transformation* process. In the example of Figure 2, the selected pattern is a hierarchical random scheduling policy implemented by a tree. The concrete ADL is a standard ADL, typically the ADL of the component model.

This generic model was projected to two specific component models, CCM and FRACTAL [3].

III. EXTENDING CCA

This section applies the generic models of both data sharing and master-worker paradigms to CCA. The projection is based on an extension of the CCA specifications (version 0.7.8). Data ports, collections and request delivery policy patterns are introduced by adding new operations in the framework service APIs. Before exposing these extensions, let us first provide an overview of the CCA component model and specifications.

A. Overview of CCA

The CCA component model [5] is a set of standards defined by a group of researchers from US national laboratories and academic institutions. The goal of the group is to develop a common architecture for building large scale scientific applications based on well tested software components.

A CCA component can define *uses* or *provides* ports. The specification of such ports is done by using the Scientific IDL. Unlike many component models, the assembly model of CCA is only dynamic. This means that there is not any Architecture Description Language (ADL) to describe components or component compositions. CCA relies on run-time calls, as ports are dynamically added or removed to components.

The specifications of CCA define standard SIDL interfaces that should be provided by any CCA compliant framework implementation. Three of them are of particular interests for this paper. First, the `Port` interface is an empty interface which all

ports must derive from. Second, the `BuilderService` interface deals with component creation as well as composition. For example, a user can create a component instance thanks to the `createInstance` operation and connect two ports via the `connect` operation. It may also obtain the list of ports available for a given component through introspection mechanisms available in this `BuilderService` interface. Third, the `Services` interface deals with port management with respect to a component implementation. A unique instance of this interface is given to each component instance. Through this interface, a component can for example declare a *provides* port and a *uses* port by respectively calling the `addProvidesPort` and `registerUsesPort` operations. The component can subsequently obtain a reference to a declared port as a result of a `getPort` invocation before being able to invoke an operation on this port.

B. Transparent data access model in CCA

Section II-A has described a generic proposal about transparent data sharing between components based on data ports. This section presents how data ports can be specified and implemented in CCA. It is first described from a user point of view and then from a framework implementer point of view.

1) *User view*: data ports are distinct ports from *provides* and *uses* ports. Hence, as shown in Figure 3, two new operations need to be added to the `Services` interface to handle them: the `createSharesPort` and the `createAccessPort` operations. They respectively creates a *shares* data port and *accesses* data port for a given data type.

These ports behaves differently from classical *provides/uses* ports: data ports come with their own pre-defined interfaces. For an *accesses* port (resp. a *shares* port), a programmer has access to the `AccessPort` interface (resp. a `SharesPort` interface). As shown in Figure 3, the API of the `AccessPort` contains operations to acquire a pointer to data, to read/write them and to handle their consistency. The `SharesPort` interface contains the same operations as the `AccessPort` interface, since it inherits `AccessPort`. A programmer may indeed access a data it shares. However, it also offers two operations to deal with the association of a piece of data with the port. The `associate` operation provides the ability for a memory area to be attached to a data port, whereas `disassociate` is the opposite operation. A component implementer is still responsible for allocating/freeing the memory. Note that the actual memory area attached to a *shares* port may dynamically change¹. However, it is transparent for all *accesses* port connected to it. Similarly, it is possible to allow a user to specify the address to which an *accesses* port maps the data. This can be done by extending the `AccessPort` interface as well as the properties of the `createAccessPort`.

The programmer can classically obtain a reference to an object providing an `AccessPort` or `SharesPort` interface through the `getPort` operation of the `Services` interface.

¹`SuspendPort/ResumePort` operations need to be added to the `Services` interface do it atomically.

```
interface AccessPort : Port {
    opaque get_pointer();
    long get_size();
    void acquire ();
    void release();
}
interface SharesPort : AccessPort {
    void associate(in opaque ptr, in long size);
    void disassociate();
}
interface ExtendedServices : Services {
    void createAccessPort(in string portName,
                        in string typeDataName,
                        in TypeMap properties);
    void createSharesPort(in string portName,
                        in string typeDataName,
                        in TypeMap properties);
}
```

Fig. 3. SIDL specifications for the CCA projection of data ports.

The connection process is also unchanged from the programmer point of view: connection between a *shares* and an *accesses* ports is done with the `connect` operation. However, if the data type of both data ports are not compatible, an exception is raised. Moreover, programmers do not have to worry about the underlying mechanism used to share data between components: it is the responsibility of the framework. Finally, a complete specification will require to introduce relevant introspection operations in the `BuilderService` interface. Such operations may for instance allow the programmer to obtain the full list of *shares* ports of a given component.

2) *Framework implementer view*: to provide the previously described view to programmers, the framework should internally be able to distinguish between classical and data ports. This distinction enables the framework to perform appropriate actions. For instance, the framework has to create and attach an instance of `AccessPort` or `SharesPort` when a data port is created. Also, when connecting data ports the data type needs to be checked in order to ensure the compatibility of ports. Another modification to the framework is required for the configuration process of the `AccessPort` object. A local reference of the data exported by the `SharesPort` needs to be given to the `AccessPort`. As the interface as to be independent of any data types, we propose to use the opaque SIDL type that represents a pointer to the local memory. Depending on the localization of the components, several technologies may be use to actually implement the data sharing. If the data is located in the same process, there is nothing to do. For components located into distinct processes but on the same node, shared memory segment provided by most operating systems can be used. On a cluster, a distributed shared memory system may be used. Finally, on grid environments a grid data-sharing service, such as JUXMEM, can be used. Advanced framework may dynamically decide of the data sharing mechanism to used depending on their availability and the localization of components. In [4], we show that the interface with respect to the framework may be generic and contains operations like data allocation and freeing, reading, writing and synchronization.

```

interface CollServices : Services {
    CollectionID getCollectionID();
}
interface CollectionID {
    string getCollectionName();
}
interface BindingID {
    CollectionID getCollectionID();
    string getCollectionPortName();
    ComponentID GetComponentID();
    string GetComponentPortName();
}
interface CollBuilderService : BuilderService {
    CollectionID createInstanceCollection(
        in string instanceCollName,
        in string className,
        in TypeMap properties);
    BindingID bind(in CollectionID collID,
        in string collPortName,
        in ComponentID elem,
        in string elemPortName);
    array<BindingID> getBindingIDs(
        in CollectionID collID,
        in string portName);
}

```

Fig. 4. Collection and binding related specifications for CCA.

To summarize, the proposed extension of the CCA specifications to support data ports implies the extension of the `Services` interface and the re-implementation of the `BuilderService` interface in the framework. Operations like `connect`, `getPort` and operations related to introspection interests are also concerned by these changes. With the experience gained with the implementation of data ports performed in CCM [4], these modifications seem to be limited and mostly straightforward to implement.

C. Master-worker paradigm in CCA

Our proposal for handling the master-worker paradigm, presented in Section II-B, was applied to component models that provide an ADL language. However, the aim of this model is to be generic and therefore being possibly used in any component models offering a similar level of abstraction. This section presents a projection of this model on CCA, a component model without an ADL language. The user view is first described, followed by the framework implementer view. As for data ports, this projection is also performed by extending the CCA specifications.

1) *User view*: the master-worker paradigm is seen as the particular case of the connection of a component (the master) that uses a port provided by a collection of components (the workers). Hence, three sets of interfaces are of interests for the user. The first set deals with the collection creation and port bindings. The second set, which is optional, is about the handling of request transport policy patterns. The third one, which is also optional, is related to the management of the dynamic variation in the number of elements inside a collection.

As CCA does not provide any ADL language, a collection needs to be a concrete entity for the programmer. Our proposal consists in the encapsulation of a collection

```

interface PatternInstantiation : Port {
    void instantiatePattern (in CollectionID collID,
        in string collPortName,
        in ComponentID user,
        in string usingPortName);
}
interface PatternBuilderService :
    CollBuilderService {
    void setPatternPort (in CollectionID collID,
        in string portName,
        in PatternInstantiation p);
}

```

Fig. 5. Pattern related specifications for CCA.

description in a distinct component. As illustrated in Figure 4, we extend the `BuilderService` with operations related to collection creation and port bindings. For the sake of simplicity, we decide to provide an operation very similar to the `createInstance` operation. The only difference is that the created component receives a reference of type `CollServices` instead of `Services`. Then, the component may create as many ports and (logically internal) components as it wants. Internal collection component ports need to be connected to the collection port. However, this connection is done via the `bind` operation instead of the classical `connect` operation as these ports are of same kind. For example, for a collection of workers, both the collection and the workers have declared a *provides* port. Last, only ports of type *provides*, *uses* or *accesses* can be bound. It appears meaningless for *shares* ports.

The next step, which is optional, is to associate a pattern with a *provides* port of a collection thanks to the `setPatternPort` operation described in Figure 5. This step is optional as the framework should provide a default pattern as explained in the framework implementer view. We choose to represent a pattern by an interface (`PatternInstantiation`). The `instantiatePattern` operation of this interface is called whenever a *uses* port is connected to a *provides* port of a collection. The implementation of such an operation has to actually connect the external *uses* port to some internal *provides* port. It may insert components to achieve more or less sophisticated request transport policies like round-robin policies or hierarchical policies.

The pattern is only applied on *provides* ports because it is meaningless for *uses* ports. When connecting a component with a *provides* port to a *uses* port of a collection, *uses* ports are directly connected to the *provides* port of the component. Note it is possible to connect a *provides* port of a collection to a *uses* port of another collections. In this case, it is the responsibility of the pattern of the *provides* port to correctly connect all *uses* ports of the other collection. It is worth to note that pattern interfaces are expected to be implemented by scheduling experts, not by the end user.

The third step, which is also optional, is dedicated to the dynamic management of a collection. For instance, the number of elements of a collection may be adapted accordingly to its load. This adaptation may be done by

```

interface CollectionManagement : Port {
    bool addElement(in CollectionID collID);
    bool subElement(in CollectionID collID);
}
interface DynCollServices : CollServices {
    void setCollManagementPort(
        in string portName,
        in collectionManagement port);
}
interface DynCollBuilderService :
    CollBuilderService {
    bool addElementToColl(in CollectionID collID);
    bool subElementToColl(in CollectionID collID);
}

```

Fig. 6. Dynamic collection management related specifications for CCA.

the collection itself or by an external component. Hence, the `DynCollBuilderService` interface, shown in Figure 6, provides operations to add or remove components to/from a collection. The role of this interface is to notify the collection of such events. Hence, a collection willing to receive such events has to register a port implementing the `CollectionManagement` interface. The registration is done thanks to the `setCollManagementPort` operation of the `DynCollServices`: on its creation, a collection in fact receives a reference to a `DynCollServices`. Addition or removal operations return a boolean to indicate either the success or the failure of the request (no more resources, components in use, etc).

2) *Framework implementer view*: interfaces presented in the previous section have a limited impact on the framework, except for the `connect` operation. The framework needs to identify if a *provides* port belongs to a collection. It is our main motivation for the introduction of the `createCollectionInstance` operation in the `CollBuilderService` interface. It is also possible to define an operation to turn an existing component into a collection. However, due to space constraints, it is not presented.

Whenever the `connect` operation has identified that a *provides* port belongs to a collection, it invokes its associated `instantiatePattern` operation. Note that the framework should behave correctly even if no pattern is associated to a *provides* port of a collection. Hence, a default pattern should be implemented by the framework. It may be up to the framework to decide what is the default pattern: connection of the first/random element of the collection, round-robin, etc.

If the port of the collection is of type *uses* (resp. *accesses*), the `connect` operation has to connect all *uses* (resp. *accesses*) ports to the external *provides* (resp. *shares*) port.

The dynamic collection management has a very low impact on the framework. The related interfaces only enable a delegation pattern.

IV. ENABLING DATA SHARING ON OPERATION INVOCATION

Previous Section has dealt with the issues of sharing a data between components and of supporting the master-worker paradigm in CCA. This section is about the issue of passing a shared data as a parameter of an operation invocation. This

```

interface compute {
    void inverse(in matrix& m1,
                out matrix& m2);
}

```

Fig. 7. Pseudo-IDL for the inverse operation.

```

// Create a shares data port of type matrix
SharesPort* dp1 = createSharesPort("p1", matrix)
// Associate the data to the data port
dp1->associate(ptr, size);
// Create an access data port of type matrix
AccessPort* dp2 = createAccessPort("p2", matrix);
// Invoke the method
to_server->inverse(dp1, dp2);

```

Fig. 8. Steps for invoking `inverse` operation from the component `Client` using shared data as parameters.

scenario was not previously handled and typically appears when a master-worker paradigm would like to have a shared data as a parameter of an operation. This section first presents a general model and then instantiates it as an extension of the CCA model. In the remainder of the section, we consider the example based on two components: a `Client` connected to a `Server` which provides an `compute` interface. The `Client` component invokes the `inverse` operation.

A. Illustration of the generic model

Let us start with component models that provide an IDL language such as CCA SIDL or OMG IDL. In order to share a parameter while invoking an operation, a notation needs to be introduced in the IDL language. This new notation should express that a parameter is passed by reference instead of value. For this purpose, the ampersand character (&) appears to be an obvious choice as it already fulfills this meaning in the C++ language. The modification of IDL languages is thus immediate and could be done with this conventional reference notation. Figure 7 shows a simple example.

Component models enforce that all incoming and outgoing communications go through some well defined ports. Therefore, the introduced notation needs to be translated into some ports. The idea is to straightforwardly map such shared parameters to the data ports introduced in Section II-A: at an operation invocation, a data port has to be associated to each parameter passed by reference. Note that there may be a difference in the implementation of the caller and the callee of the operation. On the caller side, data ports need to be explicitly created by the developer of the component to make the data available to callees, in particular to support multiple simultaneous invocations on possibly distinct shared data. On the callee side, data ports can be automatically generated by an IDL-based compiler if the compiler already generates skeleton code for handling incoming calls. If the component model does not provide any IDL language, operation prototypes are those defined hereinbefore: they contain data port parameters.

Let us now illustrate the use of this model through the previously introduced example. Figure 7 shows the operation prototype while Figure 8 shows how the component `Client` invokes the operation `inverse`. First, a data port has to


```

inverse_impl(AccessPort& dpm,
             SharesPort& dpn) {
    // Retrieve pointer & size of the data
    ptr = dpm.get_pointer();
    size = dpm.get_size();
    // Inverse the matrix with a F77 function
    res = f77_inverse(ptr, size);
    // Associate result with shares data port
    dpn->associate(res, size);
};

```

Fig. 9. Implementation of the inverse operation on the Server component with shared data as parameters.

be created for each shared parameter. For the `in` parameter, a shares data port of type `matrix` is created and then associated to the data. This permits to make the data available from outside the component. For the `out` parameter, an `accesses` data port needs to be created so as to receive the reference of the remote shared data returned by the `inverse` operation.

The code of the callee is shown on Figure 9. The `in` parameter of the `inverse` operation has been converted to an `accesses` port: the implementation of the operation will access an already created data. For the `out` parameter, the produced result needs to be associated to the `shares` data port. It also mainly applies for `inout` parameters.

Let us stress that the data reference extension is different from parameter modes. Classically, e.g. in CCA SIDL or OMG IDL, parameter modes determine the owner of data. For an `in` mode, the callee can not reallocate the data while it is possible for `inout` mode. For `out` mode, the callee is responsible to allocate the data. Our reference notation specifies that the data is *shared*. Hence, it is orthogonal with parameter modes. For `in` modes, the semantic is the following: the caller provides an access to an already created data. Therefore, the callee can access the data but without the right to reallocate it. For the `out` mode, the caller receives a reference to a shared data allocated by the callee. However, for the `inout` mode, the callee may reallocate the input data. More precisely, as the data is being shared, and thus possibly accessed by several components, it is not possible to simply deallocate an `inout` parameter of an operation. Instead, the data reference must first be dissociated from the `shares` data port, then another data reference can be associated to the port.

B. Applying the generic model to CCA

To enable data sharing on operation invocations in CCA, the projection is based on the proposed extension presented in Section III-B. As the used annotations for data port interfaces are similar to those used in the generic model, the projection is straightforward. With respect to the user view, the only changes are in the SIDL language, whose is enriched with the ampersand character, and in the mapping of SIDL operation with shared parameters as explained in Section IV-A. With respect to the implementer view, the CCA specifications are those introduced in Section III-B. No specialization is needed. However, SIDL compilers like Babel need to be modified to behave accordingly. All the work to manage shared data on




Fig. 10. A parallel component-based ray-tracer using data ports.

```

interface OADScnfig extends cca.Port {
    model3D& get_model3D();
    oads& get_oads();
}
interface Rendering extends cca.Port {
    void render (in model3D& m3D, in oads& o,
                in ray_t aRay, in frameBuffer& FB);
}

```

Fig. 11. Pseudo SIDL interface examples related to ports of the ray-tracer.

invocation, like port creation and connection, is mainly done in stubs and skeletons of *provides* and *uses* ports.

V. DATA SHARING WITH A COMPONENT-BASED RAY-TRACER

A. A ray tracing application

To illustrate the concept of data ports, we selected a parallel rendering algorithm based on ray-tracing. This algorithm follows a master-worker paradigm to distribute the computation of pixels among a set of machines and requires to share both the geometrical database and the frame-buffer. The ray tracing algorithm is used in computer graphics to render high quality images. It is based on simple optical laws which take effects such as shading, reflection and refraction into account. It acts as a light probe, following light rays in the reverse direction. The basic operation consists in tracing a ray from an origin point towards a direction in order to evaluate a light contribution. Computing realistic images requires the evaluation of several millions light contributions to a scene described by millions of objects. Tracing a ray requires to go through an Object Acceleration Data Structure (OADS) to discover whether a ray intersects an object (such as a polygon) without testing all objects of the 3D model. Parallelization of the ray-tracing [19] is very simple if the 3D model and the OADS is replicated or shared. For a complex 3D model with several millions of objects, sharing is more suitable than replication.

B. Component model

We split a ray-tracing algorithm into several components. Its architecture is shown in Figure 10 while the main interfaces are shown in Figure 11. The OADS Builder component creates the OADS from the 3D model of the scene to be rendered. Its `OADScnfig` port enables to retrieve references to them. The master component sends rays to Ray-Tracer components through the `render` operation. The 3D model, the OADS

and the frame-buffer are passed as shared parameters of the render operation. The OADS is in read-only access mode by the Ray-Tracer components while the frame-buffer is in write access mode. Therefore, several Ray-Tracer components concurrently update the frame-buffer by storing pixel values. The master component manages a set of frame-buffers for 3D animation that are accessed by a specific component (Mpeg) that produces an encoded video from these frame-buffers.

C. Discussion

Being able to actually share data like the 3D model and the OADS brings several advantages. First, the sequential code can be directly reused even though components are instantiated into distinct processes or machines. Second, it avoids to compute data distributions for the 3D model and OADS. Third, it transparently supports very large 3D models and OADS. The pressure is on the data sharing middleware, not on the components. Fourth, a modification to the 3D model, typically for animation purposes, is automatically propagated as the data sharing middleware implements the consistency model.

Being able to have shared parameters enables the decoupling of rendering operations from configuration operation: a worker component does not need to be explicitly connected to a 3D model and to OADS data. The *accesses* port creation and connection are automatically done on the operation invocation. Hence, it is very easy for a master component to simultaneously launch the computation of several frame-buffers depending or not on the same 3D model. The rendering workers act more like a service: any component having an image to render may use it.

The collection concept enables to have a master-worker relationship between only some components of the application. The handling of the collection size depending on the number of requests may be handled outside the master component. It was not represented here.

VI. CONCLUSION

Software components is a very promising technology to build complex applications. However, current component models do not support data sharing between components and poorly support the master-worker paradigm. Both are important: data sharing is useful for applications dealing with complex and large pieces of data, as it eases a lot their management; the master-worker paradigm is very common. First, the paper applies them to CCA. Second, it provides a generic model allowing shared data to be used as parameters of operation invocations. Such a model is mapped to CCA and its benefits are illustrated with a ray tracing algorithm.

Future work is twofold. First, we are implementing such an application to actually evaluate its performance. Second, we are investigating the impact of such a programming model on the deployment process. Being able to describe an application independently of the resources puts the burden on the resource selection algorithm. It has to take into account both data sharing and master-worker relationships between components when selecting resources to execute the application.

REFERENCES

- [1] C. Szyperski, *Component Software - Beyond Object-Oriented Programming*. Addison-Wesley / ACM Press, 1998.
- [2] C. Pérez, T. Priol, and A. Ribes, "A parallel corba component model for numerical code coupling," *The International Journal of High Performance Computing Applications*, vol. 17, no. 4, pp. 417–429, 2003.
- [3] H. L. Bouziane, C. Pérez, and T. Priol, "Modeling and executing master-worker applications in component models," in *11th International Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS)*, Rhodes Island, Greece, April 2006, to appear.
- [4] G. Antoniu, H. L. Bouziane, L. Breuil, M. Jan, and C. Pérez, "Enabling transparent data sharing in component models," in *6th IEEE International Symposium on Cluster Computing and the Grid (CCGRID)*, Singapore, May 2006, to appear.
- [5] D. E. Bernholdt, B. A. Allan, R. Armstrong, F. Bertrand, K. Chiu, T. L. Dahlgren, K. Damevski, W. R. Elwasif, T. G. W. Epperly, M. Govindaraju, D. S. Katz, J. A. Kohl, M. Krishnan, G. Kumfert, J. W. Larson, S. Lefantzi, M. J. Lewis, A. D. Malony, L. C. McInnes, J. Nieplocha, B. Norris, S. G. Parker, J. Ray, S. Shende, T. L. Windus, and S. Zhou, "A component architecture for high-performance scientific computing," *International Journal of High Performance Computing Applications*, nov 2005, aCTS Collection special issue.
- [6] Open Management Group (OMG), "CORBA components, version 3," Document formal/02-06-65, June 2002.
- [7] E. Bruneton, T. Coupaye, and J. Stefani, "Recursive and dynamic software composition with sharing," in *Seventh International Workshop on Component-Oriented Programming*, Malaga, Spain, jun 2002.
- [8] M. Aldinucci, S. Campa, M. Coppola, M. Danelutto, D. Laforenza, D. Puppini, L. Scarponi, M. Vanneschi, and C. Zoccolo, "Components for high performance Grid programming in the Grid.it project," in *Proc. of the Workshop on Component Models and Systems for Grid Applications (June 2004, Saint Malo, France)*, V. Getov and T. Kielmann, Eds. Springer, Jan. 2005.
- [9] J. Magee, N. Dulay, and J. Kramer, "A Constructive Development Environment for Parallel and Distributed Programs," in *Proceedings of the International Workshop on Configurable Distributed Systems*, Pittsburgh, US, Mar. 1994, pp. 4–14.
- [10] G. Antoniu, L. Bougé, and M. Jan, "JuxMem: An adaptive supportive platform for data sharing on the grid," *Scalable Computing: Practice and Experience*, vol. 6, no. 3, pp. 45–55, Nov. 2005.
- [11] G. Antoniu, J.-F. Deverge, and S. Monnet, "How to bring together fault tolerance and data consistency to enable grid data sharing," *Concurrency and Computation: Practice and Experience*, no. 17, pp. 1–19, Sept. 2005. [Online]. Available: <http://www.interscience.wiley.com>
- [12] "The JXTA project," <http://www.jxta.org/>, 2001.
- [13] D. Anderson, S. Bowyer, J. Cobb, D. Gebye, W. Sullivan, and D. Werthimer, "A new major SETI project based on Project SERENDIP data and 100,000 personal computers," *Conference Paper, Astronomical and Biochemical Origins and the Search for Life in the Universe, IAU Colloquium 161, Publisher: Bologna, Italy*, p. 729, 1997.
- [14] C. Germain, V. Néri, G. Fedak, and F. Cappello, "XtremWeb: building an experimental platform for Global Computing," in *Grid'2000*, December 2000.
- [15] "Berkeley Open Infrastructure for Network Computing," <http://boinc.berkeley.edu/>, 2002.
- [16] E. Caron, F. Desprez, F. Lombard, J. Nicod, M. Quinson, and F. Suter, "A Scalable Approach to Network Enabled Servers," in *Proceedings of the 8th International EuroPar Conference*, ser. Lecture Notes in Computer Science, B. Monien and R. Feldmann, Eds., vol. 2400. Paderborn, Germany: Springer-Verlag, August 2002, pp. 907–910.
- [17] H. Casanova and J. Dongarra, "NetSolve: A Network-Enabled Server for Solving Computational Science Problems," *The International Journal of Supercomputer Applications and High Performance Computing*, vol. 11, no. 3, pp. 212–223, 1997.
- [18] R. Buyya, D. Abramson, and J. Giddy, "Nimrod/G: An Architecture for a Resource Management and Scheduling System in a Global Computational Grid," *High-Performance Computing, ASIA, China, IEEE CS Press, USA*, vol. 01, no. 1, p. 283, 2000.
- [19] D. Badouel, K. Bouatouch, and T. Priol, "Ray tracing on distributed memory parallel computers: Strategies for distributing computation and data," *IEEE Computer Graphics and Application*, vol. 14, no. 4, pp. 69–77, July 1994.