



Chapter 1

Integrating Component-Based Scientific Computing Software

*Steven G. Parker, Keming Zhang, Kostadin Damevski,
Chris R. Johnson*

1.1 Abstract

In recent years, software component technology has been a successful methodology for large-scale commercial software development. Component technology combines a set of frequently used functions in an easily reusable form and makes the implementation transparent to the users or other components. Developers create new software applications by connecting together a group of components. Component technology is becoming increasingly popular for large-scale scientific computing for helping to tame the software complexity required in coupling multiple disciplines, multiple scales, and/or multiple physical phenomena. In the chapter, we discuss the successes achieved through the SCIRun Problem Solving Environment (PSE), and the biomedically oriented BioPSE. Notwithstanding these successes, component software systems for scientific computing provide a limited form of interoperability, typically working only with other components that employ the same underlying component model. As such, we propose a next generation concept of meta-components, where software components can be manipulated in a more abstract manner, providing a plug-in architecture for component models, and bridges between them, allowing for interoperability between different component models. These abstract, meta-components are manipulated and managed by a new software framework, called SCIRun2, while concrete, standard component models perform the actual work. This facility allows components implemented with disparate component models to be orchestrated together. As an example of a multi-component system, we have used this system to connect components from SCIRun, the Insight Toolkit (ITK), the Visualization Toolkit (Vtk), and the DOE Common Component



Architecture (CCA) into a single application.

1.2 Introduction

In recent years, component technology has been a successful methodology for large-scale commercial software development. Component technology combines a set of frequently used functions in a component and makes the implementation transparent to users. Software application developers typically connect a group of components from a component repository, connecting them to create a single application.

SCIRun¹ is a scientific Problem Solving Environment (PSE) that allows the interactive construction and steering of large-scale scientific computations [21, 20, 22, 12, 23, 26, 11]. A scientific application is constructed by connecting computational, modeling, and visualization elements [9]. This application may contain several computational elements as well as several visualization elements, all of which work together in orchestrating a solution to a scientific problem. Geometric inputs and computational parameters may be changed interactively, and the results of these changes provide immediate feedback to the investigator.

Problem solving environments, such as SCIRun, often employ component technology to bring a variety of computational tools to an engineer or scientist for solving a computational problem. In this scenario, the tools should be readily available and simple to combine to create an application. However, these PSEs typically use a single component model (such as Java Beans, Microsoft COM, CORBA, or CCA), or employ one of their own design. As a result, components designed for one PSE cannot be easily reused in another PSE or in a standalone program. Software developers must *buy in* to a particular component model and produce components for one particular system. Users must typically select a single system or face the challenges of manually managing the data transfer between multiple (usually) incompatible systems.



SCIRun2 [27], currently under development, addresses these shortcomings through a *meta-component model*, allowing support for disparate component-based systems to be incorporated into a single environment and managed through a common user-centric visual interface.

In this chapter, Section 1.3 discusses the SCIRun and BioPSE problem solving environments. Other scientific computing component models are discussed in Section 1.4. The remainder of the chapter discusses the design of SCIRun2, including a discussion of meta-components 1.5, support for distributed computing 1.6, and parallel components 1.7. We present conclusions and future work in Section 1.8.

1.3 SCIRun and BioPSE

SCIRun is a scientific PSE that allows the interactive construction and steering of large-scale scientific computations [21, 20, 22, 12, 23]. A scientific application is constructed by connecting computational elements (modules) to form a program

¹Pronounced “ski-run.” SCIRun derives its name from the Scientific Computing and Imaging (SCI) Institute at the University of Utah.



(network), as shown in Figure 1.1. The program may contain several computational elements as well as several visualization elements, all of which work together in orchestrating a solution to a scientific problem. Geometric inputs and computational parameters may be changed interactively, and the results of these changes provide immediate feedback to the investigator. SCIRun is designed to facilitate large-scale scientific computation and visualization on a wide range of architectures from the desktop to large supercomputers.

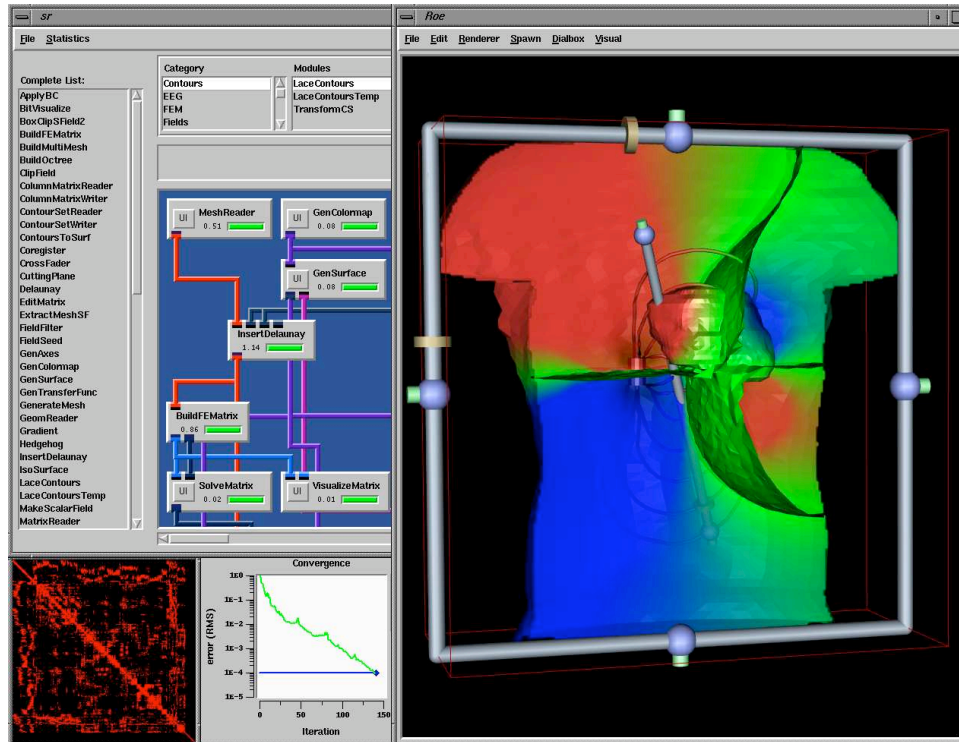




Figure 1.1. *The SCIRun PSE, illustrating a 3D finite element simulation of an implantable cardiac defibrillator.*

BioPSE [11] is a PSE based on SCIRun that is targeted to biomedical computing research problems. BioPSE provides the modeling, simulation, and visualization modules; data structures; data converters; and application networks for bioelectric field applications. With these components, researchers can investigate biomedical research issues, such as bioelectric fields produced from epileptic seizures in the brain to atrial arrhythmias in the heart.

An example electroencephalography (EEG) neural source localization application is shown in Figures 1.2 and 1.3. Figure 1.2 contains the dataflow network that implements the inverse EEG application. At the top of the network, the input data files are loaded; these include the finite element mesh that defines the geometry and electrical conductivity properties of the model and a precomputed lead-field



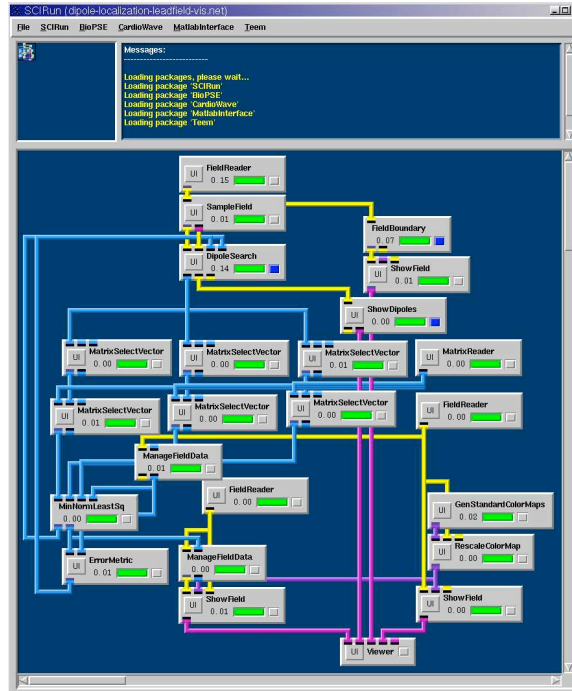




Figure 1.2. *BioPSE neural source localization network. The optimal dipole source is recovered using a multi-start optimization algorithm.*

matrix that encodes the relationship between electric sources in the domain and the resulting potentials that would be measured at the electrodes. Further down in the network, we have a set of modules that optimize the dipole location in order to minimize the misfit between the measured potentials from the electrodes and the simulated potentials due to the dipole. Finally, we have visualization and rendering modules, which provide interactive feedback to the user. The visualization that accompanies this network is shown in Figure 1.3. The potentials that were measured at the electrodes on the scalp are rendered as pseudo-colored disks; the potentials originating from the simulated dipole source are shown as pseudo-colored spheres embedded within the disks. The rainbow colors of the disks and spheres correspond to voltages, with red mapping to positive potentials, blue mapping to negative potentials, and green mapping to ground. The difference in the color of the sphere and the color of the disk at any particular electrode indicates the misfit between the measured and simulated potentials at that site. The dipoles that are iteratively approaching the true dipole location are shown as gray and blue arrows, and the outside of the head model has been rendered with wire-frame cylinders.

PowerApps Historically, one of the major hurdles to SCIRun becoming a tool for the scientist, as well as, the engineer has been SCIRun’s dataflow interface.








Figure 1.3. *Visualization of the iterative source localization. The voltages of the true solution (disks) and the computed solution (spheres) are qualitatively compared at the electrode positions as the optimization (shown as arrows) converges on a neural source location. The solution misfit can be qualitatively interpreted by pseudo-colored voltages at each electrode.*

While visual programming is natural for computer scientists and engineers, who are accustomed to writing software and building algorithmic pipelines, it can be overly cumbersome for many application scientists. Even when a dataflow network implements a specific application (such as the forward bioelectric field simulation network provided with BioPSE and detailed in the BioPSE Tutorial), the user interface (UI) components of the network are presented to the user in separate UI windows, without any semantic context for their settings. For example, SCIRun provides file browser user interfaces for reading in data. However, on the dataflow network all of the file browsers have the same generic presentation. Historically, there has not been a way to present the filename entries in their semantic context, for example to indicate that one entry should identify the electrodes input file and another should identify the finite element mesh file.

A recent release of BioPSE/SCIRun (in October 2003) addressed this short-comings by introducing *PowerApps*. A PowerApp is a customized interface built atop a dataflow application network. The dataflow network controls the execution and synchronization of the modules that comprise the application, but the generic user interface windows are replaced with entries that are placed in the context of a single application-specific interface window.

BioPSE contains a PowerApp called BioFEM. BioFEM has been built atop the forward finite element network, and provides a useful example for demonstrating the differences between the dataflow and PowerApp views of the same functionality. In Figure 1.4, the dataflow version of the application is shown: the user has separate



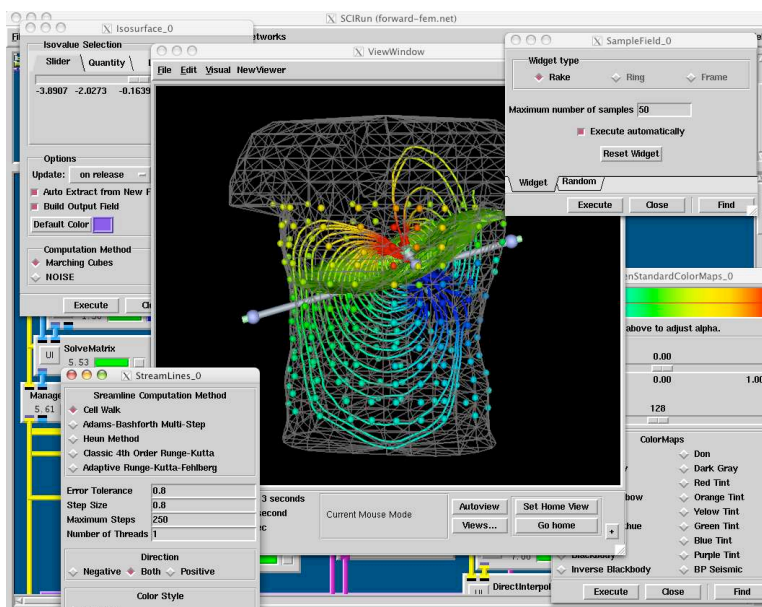


Figure 1.4. *BioPSE* dataflow interface to a forward bioelectric field application. The underlying dataflow network implements the application with modular inter-connected components called modules. Data are passed between the modules as input and output parameters to the algorithms. While this is a useful interface for prototyping, it can be nonintuitive for end-users; it is confusing to have a separate user interface window to control the settings for each module. Moreover, the entries in the user interface windows fail to provide semantic context for their settings. For example, the text-entry field on the *SampleField* user interface that is labeled “Maximum number of samples” is controlling the number of electric field streamlines that are produced for the visualization.

interface windows for controlling different aspects of the simulation and visualization. In contrast, the *PowerApp* version is shown in Figure 1.5: here, the application has been wrapped up into a single interface window, with logically arranged and semantically-labeled user interface elements composed within panels and notetabs.

In addition to bioelectric field problems, the *BioPSE* system can also be used to investigate other biomedical applications. For example, we have wrapped the tensor and raster data processing functionality of the *Teem* toolkit into the *Teem* package of *BioPSE*, and we have used that increased functionality to develop the *BioTensor PowerApp*, as seen in Figure 1.6. *BioTensor* presents a customized interface to a 140 module dataflow network. With *BioTensor* the user can visualize diffusion weighted imaging (DWI) datasets in order to investigate the anisotropic structure of biological tissues. The application supports the import of DICOM and Analyze datasets, and implements the latest diffusion tensor visualization techniques, including superquadric glyphs[13] and tensorlines[25] (both shown).

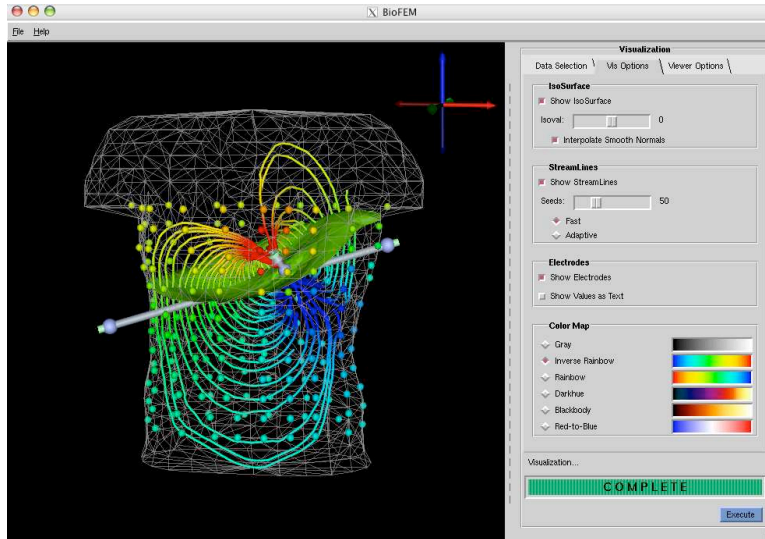




Figure 1.5. *The BioFEM custom interface. Though the application is functionality equivalent to the dataflow version shown in Figure 1.4, this PowerApp version provides an easier-to-use custom interface. Everything is contained within a single window; the user is lead through the steps of loading and visualizing the data with the tabs on the right; and generic control settings have been replaced with contextually appropriate labels; and application-specific tooltips (not shown) appear when the user places the cursor over any user interface element.*

1.4 Components for Scientific Computing

A number of component models have been developed for a wide range of software applications. Java Beans[2], a component model from Sun, is a platform-neutral architecture for the Java application environment. However, it requires a Java Virtual Machine as the intermediate platform and the components must be written in Java. Microsoft has developed the Component Object Model (COM)[17], a software architecture that allows applications to be built from binary software components on the Windows platform. The Object Management Group (OMG) developed the Common Object Request Broker Architecture (CORBA)[18], which is an open, vendor-independent architecture and infrastructure that computer applications can use to work together in a distributed environment.

Many problem solving environments, such as SCIRun, employ these component models, or one of their own. As an example, SCIRun provides a dataflow-based component model. The Common Component Architecture (CCA) Forum, a group of researchers from several DOE national laboratories and academic institutions, has defined a standard component architecture[1] for high performance parallel computing. The CCA forum has defined a minimal set of standard interfaces that a high-performance component framework should provide to implement high-performance components. This standard promotes interoperability between








Figure 1.6. *The BioTensor PowerApp. Just as with BioFEM, we have wrapped up a complicated dataflow network into a custom application. In the left panel, the user is guided through the stages of loading the data, co-registering MRI diffusion weighted images, and constructing diffusion tensors. On the right panel, the user has controls for setting the visualization options. In the rendering window in the middle, the user can render and interact with the dataset.*

components developed by different teams across different institutions. However, CCA has not yet fully addressed the architecture of parallel components combined with distributed computation.

CCA is discussed in more detail in Chapter ??, but we present an overview here. The CCA model consists of a framework and an expandable set of components. The framework is a workbench for building, connecting and running the components. A component is the basic unit of an application. A CCA component consists of one or more ports, and a port is a group of method-call based interfaces. There are two types of ports: **uses** port and **provides** ports. A provides port (or callee) implements its interfaces and waits for other ports to call them. A uses port (or caller) issues method calls that can be fulfilled by a type-compatible provides port on a different component.

A CCA port is represented by an interface, while interfaces are specified through a Scientific Interfaces Definition Language (SIDL). A compiler is usually used to compile a SIDL interface description file into specific language bindings. Generally, component language binding can be provided for many different languages such as C/C++, Java, Fortran, Python etc. The Babel [14] compiler group is working on creating this support for different languages within CCA.

SCIRun2 is a new software framework that combines CCA compatibility with connections to other commercial and academic component models. SCIRun2 is based on the SCIRun[10] infrastructure and the CCA specification. It utilizes parallel-to-parallel remote method invocation to connect components in a dis-



tributed memory environment, and is multi-threaded to facilitate shared memory programming. It also has an optional visual-programming interface.



Although SCIRun2 is designed to be fully compatible with CCA. It aims to combine CCA compatibility with the strength of other component models. A few of the design goals of SCIRun2 are:

1. SCIRun2 is fully CCA compatible, thus any CCA components can be used in SCIRun2 and CCA components developed from SCIRun2 can also be used in other CCA frameworks.
2. SCIRun2 accommodates several useful component models. In addition to CCA components and SCIRun Dataflow Modules, CORBA components, Microsoft COM components, ITK and Vtk[24] modules will be supported in SCIRun2.
3. SCIRun2 builds bridges between different component models, so that we can combine a disparate array of computational tools to create powerful applications with cooperative components from different sources.
4. SCIRun2 supports distributed computing. Components created on different computers can work together through a network and build high performance applications.
5. SCIRun2 supports parallel components in a variety of ways for maximum flexibility. This is not constrained to only CCA components, because SCIRun2 employs a M process to N process method invocation and data redistribution (MxN) library [5] that potentially can be used by many component models.

Overall, SCIRun2 provides a broad approach that will allow scientists to combine a variety of tools for solving a particular computational problem. The overarching design goal of SCIRun2 is to provide the ability for a computational scientist to use the right tool for the right job, a goal motivated by the needs of our biomedical and other scientific users.

1.5 Metacomponent Model

Systems such as Java Beans, COM, CORBA, CCA, and other successfully employ a component-based architecture to allow users to rapidly assembly computational tools in a single environment. However, these systems typically do not interact with one another in a straightforward manner, and it is difficult to take components developed for one system and re-deploy them in another. Software developers must *buy in* to a particular model and produce components for one particular system. Users must typically select a single system or face the challenges of manually managing the data transfer between multiple (usually) incompatible systems. SCIRun2 addresses these shortcomings through the meta-component model, allowing support for disparate component-based systems to be incorporated into a single environment and managed through a common user-centric visual interface. Furthermore, many



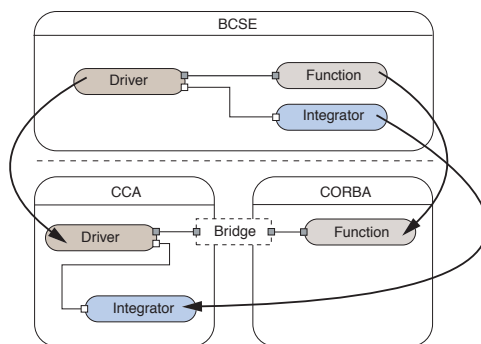




Figure 1.7. *Components of different models cooperate in SCIRun2*

systems that are not traditionally thought of as component models, but that have well-designed, regular structures, can be mapped to a component model and manipulated dynamically. SCIRun2 combines support for CCA components, "old-style" SCIRun dataflow components, and we are planning support for CORBA, COM and VTK. As a result, SCIRun2 can utilize SCIRun components, CCA components and other software components in the same simulation.

The meta component model operates by providing a plug-in architecture for component models. Abstract components are manipulated and managed by the SCIRun2 framework, while concrete component models perform the actual work. This facility allows components implemented with disparate component models to be orchestrated together.

Figure 1.7 demonstrates a simple example of how SCIRun2 handles different component models. Two CCA components, **Driver** and **Integrator**, and one CORBA component, **Function**, are created in the SCIRun2 framework. In this simple example, the Driver is connected to both the Function and Integrator. Inside SCIRun2, two frameworks are hidden: the CCA framework and the CORBA Object Request Broker (ORB). The CCA framework creates the CCA components, Driver and Integrator. The CORBA framework creates the CORBA component, Function. The two CCA components can be connected in a straightforward manner through the CCA component model. However, the components Driver and Function cannot be connected directly, because neither CCA nor CORBA allow a connection from a component of a different model. Instead, a bridge component is created. Bridges belong to a special internal component model that is used to build a connection between components of different component models. In this example, Bridge has two ports: one CCA port and one CORBA port. In this way it can be connected to both CCA component and CORBA component. The CORBA invocation is converted to request to the CCA port inside the bridge component.

Bridge components can be manually or automatically generated. In situations where interfaces are easily mapped between one interface and another, automatically generated bridges can facilitate interoperability in a straightforward



way. More complex component interactions may require manually generated bridge components. Bridge components may implement heavy-weight transformations between component models, and therefore have the potential to introduce performance bottlenecks. For the few scenarios that require maximum performance, reimplementing of both components in a common, performance-oriented component model may be required. However, for rapid prototyping, or for components that are not performance-critical, this is completely acceptable.

To automatically generate a bridge component that translates a given pair of components, a generalized translation must be completed between the component models. A software engineer designs how two particular component models will interact. This task can require creating methods of data and control translation between the two models and can be quite difficult in some scenarios. The software engineer expresses the translation into a compiler plugin, which is used as a specification of the translation process. A plugin abstractly represents the entire translation between the two component models. It is specified by an eRuby (embedded Ruby) template document. eRuby templates are text files that can be augmented by Ruby [15] scripts. The Ruby scripts are useful for situations where the translation requires more sophistication than regular text (such as control structures or additional parsing). This provides us with better flexibility and more power inside the plugin, with the end goal of being able to support the translation of a wider range of component models.



The only other source of information is the interface of the ports we want to bridge (usually expressed in an IDL file). The bridge compiler accepts commands that specify a mapping between incompatible interfaces, where the interfaces between the components differ in various names or types but not functionality. Using a combination of the plugin and the interface augmented with mapping commands, the compiler is able to generate the specific bridge component. This component is automatically connected and ready to broker the translation between the two components of different models.

Figure 1.8 shows a more complex example that is motivated by the needs of a biological application. This example works very much like the last: the framework manages components from several different component models through the meta-model interface. Components from the same model interact with each other natively, and interact with components in other models through bridges. Allowing components to communicate with each other through their native mechanisms ensures that no performance bottlenecks are introduced and that the original semantics are preserved.

1.6 Distributed Computing

SCIRun2 provides support for Remote Method Invocation (RMI) based distributed objects. This support is utilized in the core of the SCIRun framework in addition to distributed components. This section describes the design of the distributed object subsystem.

A distributed object is a set of interfaces defined by SIDL that can be refer-



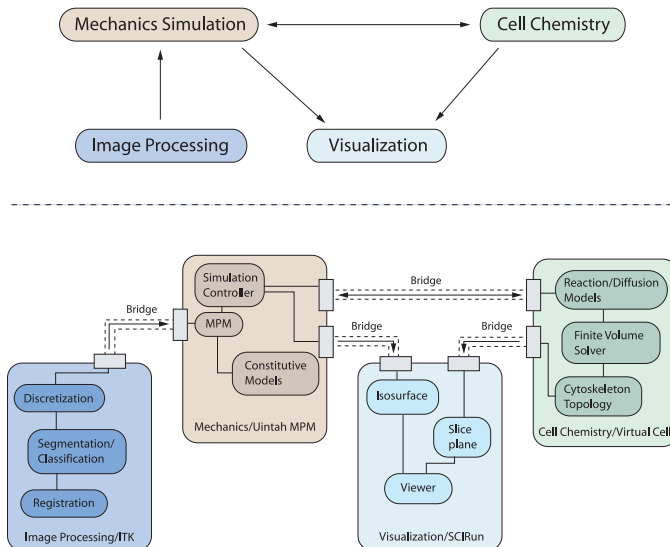




Figure 1.8. A more intricate example of how components of different models cooperate in SCIRun2. The application and components shown are from a realistic (albeit incomplete) scenario.

enced over network. The distributed object is similar to the C++ object, it utilizes similar inheritance rules and all objects share the same code. However only methods (interfaces) can be referenced, and the interfaces must be defined in SIDL. Using the SIDL language, we implemented a straightforward distributed object system. We extend the SIDL language and build upon this system for implementing parallel to parallel component connections, as discussed in the next section.

A distributed object is implemented by a concrete C++ class, and referenced by a proxy class. The proxy class is a machine-generated class that associates the user-made method calls to a call by the concrete object. The proxy classes are described in a SIDL file, and a compiler compiles the SIDL file and creates the proxy classes. The proxy classes define the abstract classes with a set of pure virtual functions. The concrete classes extends those abstract proxy classes and implement each virtual functions.

There are two types of object proxies. One is called server proxy, the other is called client proxy. The server proxy (or skeleton) is the object proxy created in the same memory address space as the concrete object. When the concrete object is created, the server proxy starts and works as a server, waiting for any local or remote methods invocations. The client proxy (or stub) is the proxy created on a different memory address space. When a method is called through the client proxy, the client proxy will package the calling arguments into a single message, and send the message to the server proxy, and then wait for the server proxy to invoke the methods and return the result and argument changes.

We created Data Transmitter, a separate layer, that is used by the generated



proxy code for handling messaging. We also employ the concept of a Data Transmission Point (DTP), which is similar to the start point and end points used in Nexus [8]. A DTP is a data structure that contains an object pointer pointing to the context of a concrete class. Each memory address space has only one Data Transmitter, and each Data Transmitter uses three communication ports (sockets): one listening port, one receiving port and one sending port. All the DTPs in the same address space share the same Data Transmitter. A Data Transmitter is identified by its universal resource identifier (URI): IP address + listening port. A DTP is identified by its memory address together with the Data Transmitter URI, because DTP addresses are unique in the same memory address space. Optionally, we could use other type of object identifiers.

The proxy objects package method calls into messages by marshaling objects and then waiting for a reply. Non-pointer arguments, such as integers, fixed sized arrays and strings (character arrays), are marshaled by the proxy into a message in the order that they are presented in the method. After the server proxy receives the message, it unmarshals the arguments in the same order. An array size is marshaled in the beginning of an array argument, so the proxy knows how to allocate memory for the array. SIDL supports a special opaque data type that can be used to marshal pointers if the two objects are in the same address space. Distributed object references are marshaled by packaging the DTP URI (Data Transmitter URI and object ID). The DTP URI is actually marshaled as a string and when it is unmarshaled, a new proxy of the appropriate type is created based on the DTP URI.



C++ exceptions are handled as special distributed objects. In a remote method invocation, the server proxy tries to catch an exception (also a distributed object) before it returns. If it catches one, the exception pointer is marshaled to the returned message. Upon receiving the message, the client proxy unmarshals the message and obtains the exception. The exception is then re-thrown by the proxy.

1.7 Parallel Components

This section introduces the CCA parallel component design and discusses issues of the implementation. Our design goal is to make the parallelism transparent to the component users. In most cases, the component users can use a parallel component as the way they use sequential component without knowing that a component is actually parallel component.

Parallel CCA Component (PCom) is a set of similar components that run in a set of processes respectively. When the number of process is one, the PCom is equivalent to a sequential component. We call each component in a PCom a *member component*. Member components typically communicate internally with MPI [16] or an equivalent message passing library.

PComs communicate with each other through CCA-style RMI ports. We developed a prototype parallel component infrastructure [4, 3] that facilitates connection of parallel components in a distributed environment. This model supports two types of methods calls: *independent* and *collective*, and as such our port model



supports both independent and collective ports.

An independent port is created by a single component member, and it contains only independent interfaces. A collective port is created and owned by all component members in a PCom, and one or more of its methods are collective. Collective methods require that all member components participate in the collective calls in the same order.

As an example of how parallel components interact, let pA be a uses port of component A, and pB be a provides port of component B, Both pA and pB have the same port type, which defines the interface. If pB is a collective port, and has the following interface,

```
collective int foo(inout int arg);
```



Then `getPort("pA")` returns a collective pointer that points to the collective port pB. If pB is an independent port, `getPort("pA")` returns a pointer that points to an independent port.

Component A can have one or more members, so each member might obtain a (collective/independent) pointer to a provides port. The component developer can decide what subset (one, many, or all components) participate in a method call `foo(arg)`. When any member component register a uses port, all other members can share the same uses port. But for a collective provides port, each member must call `addProvidesPort` to register each member port.

The *MxN* library takes care of the collective method invocation and data distribution. We repeat only the essentials here, one can reference [5] for details. If a M-member PCom A obtains a pointer ptr pointing to a N-member PCom's B collective port pB. Then `ptr->foo(args)` is a collective method invocation. The *MxN* library index PCom members with rank 0,1,...,M-1 for A and 0,1,...,N-1 for B. If $M = N$, then the i-th member component of A call `foo(args)` on the i-th component of B. But if $M < N$, then we "extend" the A's to 0,1,2,...,M, 0, 1,2,...,M, ... N-1 and they call `foo(args)` on each member component of B like the $M=N$ case, but only the first M calls request returns. The left panel of Figure 1.9 shows an example of this case with $M=3$ and $N=5$. If $M > N$, we "extends" component B's set to 0, 1, ..., N, 0, 1,...,N, ...,M-1 and only the first N member components of B are actually called, and the rest are not called but simply return the result. We rely on collective semantics from the components to ensure consistency without requiring global synchronization. The right panel of Figure 1.9 shows an example of this case with $M=5$ and $N=3$.

The *MxN* library also does most of the work for the data redistribution. An multi-dimensional array can be defined as a distributed array that associates a distribution scheduler with the real data. Both callers and callees define the distribution schedule before the remote method invocation, using an first-stride-last representation for each dimension of the array. The SIDL compiler creates the scheduler and scheduling is done in the background.

With independent ports and collective ports, we cover the two extremes. Ports that require communication among a subset of the member components present a greater challenge. Instead, we utilize a sub-setting capability in the *MxN* system to produce ports that are associated with a subset of the member components, and






Figure 1.9. $M \times N$ method invocation, with the caller on the left and the callee on the right. In the left scenario, the number of callers is fewer than the numbers of callees, so some callers make multiple method calls. In the right, the number of callees is fewer, so some callees send multiple return values.



then utilize them as collective ports.

SCIRun2 provides the mechanism to start a parallel component either on shared memory multi-processors computers, or clusters. SCIRun2 consists of a main framework and a set of Parallel Component Loaders (PCLs). A PCL can be started with ssh on a cluster, where it gathers and reports its local component repository and registers to the main framework. The PCL on a N-node cluster is essentially a set of loaders, each running on a node. When the user requests to create a parallel component, the PCL instantiate a parallel component on its processes (or nodes), and passes a distributed pointer to the SCIRun2 framework. PCLs are responsible to create and destroy components running on their nodes, but they do not maintain the port connections. The SCIRun2 framework maintains all component status and port connections.

Supporting threads and MPI together can be difficult. MPI provides a convenient communication among the processes in a cluster. However, if any process has more than one thread and the MPI calls are made in those threads, the MPI communication may break because MPI distinguish only processes, not threads. The MPI interface allows an implementation to support threads but does not require it. Most MPI implementations are not threadsafe. We provide support for both threadsafe and non-threadsafe MPI implementations so that users can choose any available MPI.

A straightforward way to support non-threadsafe MPI is to globally order the MPI calls such that no two MPI calls are executed at the same time. We implemented a distributed lock, which has two interfaces:

```
PRMI::lock()
PRMI::unlock()
```



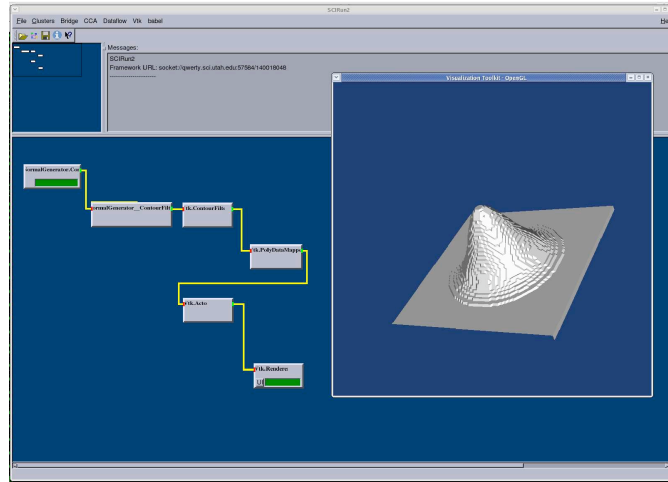




Figure 1.10. *Components of different models cooperate in SCIRun2*

The distributed lock is just like a mutex, but it is collective with respect to all MPI processes in a cluster. The critical section between `PRMI::lock()` and `PRMI::unlock()` can be obtained by only one set of threads in different MPI processes. The users must call `PRMI::lock()` before any MPI calls and call `PRMI::unlock()` after to release the lock. More than one MPI calls can be made in the critical section. In this way only one set of threads (each from a MPI process) can make MPI calls at one time. Additionally, the overhead of acquiring and releasing this lock is very high because it requires a global synchronization. However, in some cases this approach is necessary for supporting the multi-threaded software framework in an environment where a thread-safe MPI is no available.

It is fairly easier to support threadsafe MPI. Our approach is to create a distinct MPI communicator for the threads that communicate with each other, and restrict that those threads can use only that communicator for MPI communication. The special communicators are created by the PCL and can be obtained through a framework service interface. The threadsafe MPI allows multiple MPI calls executed safely at the same time, and the designated communicators help to identify the group of threads which initiated the MPI calls.

An efficient mechanism allows parallel components to efficient coordinate around error conditions [6].

Figure 1.10 shows a SCIRun2 application that uses bridging to Vtk visualization components. SCIRun2 is currently under development, but we expect a public release in the near future.



1.8 Conclusions and Future Work



We presented the SCIRun, BioPSE, and SCIRun2 problem solving environments for scientific computing. These systems all employ software components to encapsulate computational functionality into a reusable unit. SCIRun and BioPSE are open source and have biannual public releases and are used by a number of end-users for a variety of different computational applications.



Additionally, we presented an overview of the new SCIRun2 component framework. SCIRun2 integrates multiple component models into a single visual problem solving environment and builds bridges between components of different component models. In this way, a number of tools can be combined into a single environment without requiring global adoption of a common underlying component model. We have also described a parallel component architecture utilizing the common component architecture, combined with distributed objects and parallel $M \times N$ array redistribution that can be used in SCIRun2.

A prototype of the SCIRun2 framework has been developed, and we are using this framework for a number of applications in order to demonstrate the SCIRun2 features. Future applications will rely more on the system, and will facilitate joining many powerful tools, such as the SCI Institutes' interactive ray-tracing system [19] and the Uintah [7] parallel, multi-physics system. Additional large scale computational applications are under construction and are beginning to take advantage of the capabilities of SCIRun2. Support for additional component-models, such as Vtk, CORBA, and possibly others, will be added in the future.

1.9 Acknowledgments



The authors gratefully acknowledge support from NIH NCRR, NSF and the DOE ASCI and SciDAC programs. The authors would also like to acknowledge contributions from David Weinstein. SCIRun and BioPSE are available as open source at www.sci.utah.edu.







Bibliography

- [1] R. Armstrong, D. Gannon, A. Geist, K. Keahey, S. Kohn, L. McInnes, S. Parker, and B. Smolinski. Toward a Common Component Architecture for High-Performance Scientific Computing. In *Proceedings of the 8th IEEE International Symposium on High Performance Distributed Computing*, 1999.
- [2] Enterprise Java Beans. <http://java.sun.com/products/javabeans>, 2003.
- [3] F. Bertrand, R. Bramley, K. Damevski, D. Bernholdt, J. Kohl, J. Larson, and A. Sussman. Data redistribution and remote method invocation in parallel component architectures. In *Proceedings of The 19th International Parallel and Distributed Processing Symposium*, 2005. (Accepted, Best Paper Award).
- [4] K. Damevski. Parallel component interaction with an interface definition language compiler. Master's thesis, University of Utah, 2003.
- [5] K. Damevski and S. Parker. Parallel remote method invocation and m-by-n data redistribution. In *Proceedings of the 4th Los Alamos Computer Science Institute Symposium*, 2003.
- [6] Kostadin Damevski and Steven Parker. Imprecise exceptions in distributed parallel components. In *Proceedings of the 10th International Euro-Par Conference*, volume 3149 of *Lecture Notes in Computer Science*, pages 108–116, August/September 2004.
- [7] J. Davison de St. Germain, John McCorquodale, Steven G. Parker, and Christopher R. Johnson. Uintah: A Massively Parallel Problem Solving Environment. In *Proceedings of the Ninth IEEE International Symposium on High Performance and Distributed Computing*, August 2000.
- [8] I. Foster, C. Kesselman, and S. Tuecke. The Nexus approach to integrating multithreading and communication. *Journal of Parallel and Distributed Computing*, 37:70–82, 1996.
- [9] C.D. Hansen and C.R. Johnson, editors. *The Visualization Handbook*. Elsevier, 2005.
- [10] C. Johnson and S. Parker. The SCIRun Parallel Scientific Computing Problem Solving Environment. In *Proceedings of the 9th SIAM Conference on Parallel Processing for Scientific Computing*, 1999.



- [11] C.R. Johnson, R.S. MacLeod, S.G. Parker, and D.M. Weinstein. Biomedical computing and visualization software environments. *Comm. ACM*, 47(11):64–71, 2004.
- [12] C.R. Johnson and S.G. Parker. Applications in computational medicine using SCIRun: A computational steering programming environment. In H.W. Meuer, editor, *Supercomputer '95*, pages 2–19. Springer-Verlag, 1995.
- [13] G. Kindlmann. Superquadric tensor glyphs. In *The Joint Eurographics – IEEE TCVG Symposium on Visualization 2004*, pages 147–154, May 2004.
- [14] S. Kohn, G. Kumpfert, J. Painter, and C. Ribbens. Divorcing language dependencies from a scientific software library. In *Proceedings of the 10th SIAM Conference on Parallel Processing*, Portsmouth, VA, March 2001.
- [15] The Ruby Language. <http://www.ruby-lang.org/en>, 2004.
- [16] Message Passing Interface Forum. *MPI: A Message-Passing Interface Standard*, June 1995.
- [17] Component Object Model. <http://www.microsoft.com/com/tech/com.asp>, 2003.
- [18] OMG. *The Common Object Request Broker: Architecture and Specification. Revision 2.0*, June 1995.
- [19] S. Parker, M. Parker, Y. Livnat, P. Sloan, and P. Shirley. Interactive ray tracing for volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, July-September 1999.
- [20] S. G. Parker. *The SCIRun Problem Solving Environment and Computational Steering Software System*. PhD thesis, University of Utah, 1999.
- [21] S.G. Parker, D.M. Beazley, and C.R. Johnson. Computational steering software systems and strategies. *IEEE Computational Science and Engineering*, 4(4):50–59, 1997.
- [22] S.G. Parker and C.R. Johnson. SCIRun: A scientific programming environment for computational steering. In *Supercomputing '95*. IEEE Press, 1995.
- [23] S.G. Parker, D.M. Weinstein, and C.R. Johnson. The SCIRun computational steering software system. In E. Arge, A.M. Bruaset, and H.P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 1–44. Birkhauser Press, 1997.
- [24] W. Schroeder, K. Martin, and B. Lorensen. *The Visualization Toolkit, An Object-Oriented Approach to 3-D Graphics*. Prentice Hall PTR, 2nd edition, 2003.



- [25] D. Weinstein, O. Potniaguine, and L. Zhukov. A comparison of dipolar and focused inversion for EEG source imaging. In *3rd International Symposium on Noninvasive Functional Source Imaging (NFSI)*, pages 121–123, Innsbruck, Austria, September 2001.
- [26] D.M. Weinstein, S.G. Parker, J. Simpson, K. Zimmerman, and G. Jones. Visualization in the scirun problem-solving environment. In C.D. Hansen and C.R. Johnson, editors, *The Visualization Handbook*, pages 615–632. Elsevier, 2005.
- [27] K. Zhang, K. Damevski, V. Venkatachalapathy, and S.G. Parker. SCIRun2: A CCA framework for high performance computing. In *Proceedings of The 9th International Workshop on High-Level Parallel Programming Models and Supportive Environments*, April 2004.

