A Component-Based Approach to Integrated Modeling in the Geosciences: The Design of CSDMS

Scott Peckham, Eric Hutton

CSDMS, University of Colorado, 1560 30th Street, UCB 450, Boulder, CO 80309, USA

Boyana Norris

Mathematics and Computer Science Division, Argonne National Laboratory, 9700 S. Cass Ave., Argonne, IL 60439, USA

Abstract

The development of scientific modeling software increasingly requires the coupling of multiple independently developed models. Component-based software engineering enables the integration of plug-and-play components, but significant additional challenges must be addressed in any specific domain in order to produce a usable development and simulation environment that is also going to encourage contributions and adoption by entire communities. In this paper we describe the challenges in creating a coupling environment for Earth-surface process modeling and how we approach them in our integration efforts at the Community Surface Dynamics Modeling System.

Keywords:

component software, CCA, CSDMS, modeling, code generation

Preprint submitted to Computers and Geosciences: Modeling for Environmental ChangeMarch 16, 2011

Email addresses: Scott.Peckham@colorado.edu (Scott Peckham), Eric.Hutton@colorado.edu (Eric Hutton), norris@mcs.anl.gov (Boyana Norris)

1 1. Introduction

The Community Surface Dynamics Modeling System (CSDMS) project [12] 2 is an NSF-funded, international effort to develop a suite of modular numerical 3 models able to simulate a wide variety of Earth-surface processes, on time 4 scales ranging from individual events to many millions of years. CSDMS 5 maintains a large, searchable inventory of contributed models and promotes 6 the sharing, reuse, and integration of open-source modeling software. It has adopted a component-based software development model and has created 8 a suite of tools that make the creation of *pluq-and-play* components from 9 stand-alone models as automated and effortless as possible. Models or pro-10 cess modules that have been converted to component form are much more 11 flexible and can be rapidly assembled into new configurations to solve a wider 12 variety of scientific problems. The ease with which one component can be re-13 placed by another also makes it easy to experiment with different approaches 14 to providing a particular type of functionality. The CSDMS project also has a 15 mandate from the NSF to provide a migration pathway for surface dynamics 16 modelers toward high-performance computing (HPC) and provides a 720-17 core supercomputer for use by its members. In addition, CSDMS provides 18 educational infrastructure related to physically based modeling. 19

The main purpose of this paper is to present in some detail the key issues and design criteria for a component-based, integrated modeling system and then describe the design choices adopted by the CSDMS project to address these criteria. CSDMS was not developed in isolation: it builds on and extends proven, open-source technology. The CSDMS project also maintains close collaborations with several other integrated modeling projects and seeks to evaluate different approaches in pursuit of those that are optimal. As with
any design problem, myriad factors must be considered in determining what
is optimal, including how various choices affect users and developers. Other
key factors are performance, ease of maintenance, ease of use, flexibility,
portability, stability, encapsulation, and future proofing.

31 1.1. Component Programming Concepts

Component-based programming is all about bringing the advantages of 32 "plug and play" technology into the realm of software. When one buys a 33 new peripheral for a computer, such as a mouse or printer, the goal is to 34 be able to simply plug it into the right kind of port (e.g., a USB, serial, 35 or parallel port) and have it work, right out of the box. For this situation 36 to be possible, however, some kind of published standard is needed that 37 the makers of peripheral devices can design against. For example, most 38 computers have universal serial bus (USB) ports, and the USB standard is 39 well documented. A computer's USB port can always be expected to provide 40 certain capabilities, such as the ability to transmit data at a particular speed 41 and the ability to provide a 5-volt supply of power with a maximum current 42 of 500 mA. The result of this standardization is that one can usually buy a 43 new device, plug it into a computer's USB port, and start using it. Software 44 "plug-ins" work in a similar manner, relying on interfaces (ports) that have 45 well-documented structure or capabilities. In software, as in hardware, the 46 term *component* refers to a unit that delivers a particular type of functionality 47 and that can be "plugged in." 48

⁴⁹ Component programming build on the fundamental concepts of object-⁵⁰ oriented programming, with the main difference being the introduction or presence of a runtime *framework*. Components are generally implemented as
classes in an object-oriented language, and are essentially "black boxes" that
encapsulate some useful bit of functionality.

The purpose of a framework is to provide an environment in which com-54 ponents can be linked together to form applications. The framework provides 55 a number of *services* that are accessible to all components, such as the linking 56 mechanism itself. Often, a framework will also provide a uniform method of 57 trapping or handling exceptions (i.e., errors), keeping in mind that each com-58 ponent will throw exceptions according to the rules of the language that it is 59 written in. In some frameworks (e.g., CCA's Ccaffeine [1]), there is a mech-60 anism by which any component can be promoted to a framework service, as 61 explained in a later section. 62

One feature that often distinguishes components from ordinary subrou-63 tines, software modules, or classes is that they are able to communicate with 64 other components that may be written in a different programming language. 65 This capability is referred to as *language interoperability*. In order for this 66 to be possible, the framework must provide a language interoperability tool 67 that can create the necessary "glue code" between the components. For a 68 CCA-compliant framework, that tool is Babel [14, 29], and the supported 69 languages are C, C++, Fortran (77-2003), Java, and Python. Babel is de-70 scribed in more detail in a later section. For Microsoft's .NET framework [33], 71 that tool is CLR (Common Language Runtime), which is an implementation 72 of an open standard called CLI (Common Language Infrastructure), also 73 developed by Microsoft. Some of the supported languages are C# (a spin-74 off of Java), Visual Basic, C++/CLI, IronLisp, IronPython, and IronRuby. 75

⁷⁶ CLR runs a form of bytecode called CIL (Common Intermediate Language).

⁷⁷ Note that CLI does not support Fortran, Java, standard C++, or standard
⁷⁸ Python.

The Java-based frameworks used by Sun Microsystems are JavaBeans and
Enterprise JavaBeans (EJB) [17]. In the words of Armstrong et al. [3]:

Neither JavaBeans nor EJB directly addresses the issue of lan-81 guage interoperability, and therefore neither is appropriate for 82 the scientific computing environment. Both JavaBeans and EJB 83 assume that all components are written in the Java language. Al-84 though the Java Native Interface library supports interoperabil-85 ity with C and C++, using the Java virtual machine to mediate 86 communication between components would incur an intolerable 87 performance penalty on every inter-component function call. 88

While in recent years the performance of Java codes has improved steadily through just-in-time (JIT) compilation into native code, Java is not yet available on key high-performance platforms such as the IBM Blue Gene/L and Blue Gene/P supercomputers.

⁹³ Key advantages of component-based programming include the following.

• Components can be written in different languages and still communicate (via language interoperability).

Components can be replaced, added to, or deleted from an application
 at runtime via dynamic linking (as precompiled units).

5

- Components can easily be moved to a remote location (different address space) without recompiling other parts of the application (via RMI/RPC support).
- Components can have multiple different interfaces.
- Components can be "stateful"; that is, data encapsulated in the component is retained between method calls over its lifetime.
- Components can be customized at runtime with configuration parameters.
- Components provide a clear specification of inputs needed from other 107 components in the system.
- Components allow multicasting calls that do not need return values (i.e., send data to multiple components simultaneously).
- Components provide clean separation of functionality (for components,
 this is mandatory vs. optional).
- Components facilitate code reuse and rapid comparison of different implementations.
- Components facilitate efficient cooperation between groups, each doing what it does best.
- Components promote economy of scale through development of com munity standards.

¹¹⁸ 2. Background

We briefly overview the component methodology used in CSDMS and the associated tools that support component development and application execution.

122 2.1. The Common Component Architecture

The Common Component Architecture (CCA) [3] is a component ar-123 chitecture standard adopted by federal agencies (largely the Department of 124 Energy and its national laboratories) and academics to allow software com-125 ponents to be combined and integrated for enhanced functionality on high-126 performance computing systems. The CCA Forum is a grassroots organiza-127 tion that started in 1998 to promote component technology standards (and 128 code reuse) for HPC. CCA defines standards necessary for the interopera-129 tion of components developed in different frameworks. Software components 130 that adhere to these standards can be ported with relative ease to another 131 CCA-compliant framework. While a variety of other component architecture 132 standards exist in the commercial sector (e.g., CORBA, COM, .Net, and Jav-133 aBeans), CCA was created to fulfill the needs of scientific, high-performance, 134 open-source computing that are unmet by these other standards. For ex-135 ample, scientific software needs full support for complex numbers, dynam-136 ically dimensioned multidimensional arrays, Fortran (and other languages), 137 and multiple processor systems. Armstrong et al. [3] explain the motivation 138 for creating CCA by discussing the pros and cons of other component-based 139 frameworks with regard to scientific, high-performance computing. A number 140 of DOE projects, many associated with the Scientific Discovery through Ad-141

vanced Computing (SciDAC) [46] program, are devoted to the development
of component technology for high-performance computing systems. Several
of these are heavily invested in the CCA standard (or are moving toward
it) and involve computer scientists and applied mathematicians. Examples
include the following.

• TASCS: The Center for Technology for Advanced Scientific Computing 147 Software, which focused on CCA and its associated tools [9]. 148 • CASC: Center for Applied Scientific Computing, which is home to 149 CCA's Babel tool [29]. 150 • ITAPS: The Interoperable Technologies for Advanced Petascale Simu-151 lation [16], which focuses on meshing and discretization components, 152 formerly TSTT. 153 • PERI: Performance Engineering Research Institute, which focuses on 154 HPC quality of service and performance issues [30]. 155 • TOPS: Terascale Optimal PDE Solvers, which focuses on PDE solver 156 components [24]. 157 • PETSc: Portable, Extensible Toolkit for Scientific Computation, which 158 focuses on linear and nonlinear PDE solvers for HPC, using MPI [6, 7, 159 8]. 160

A variety of different frameworks, such as Ccaffeine [1], CCAT/XCAT [25],
 SciRUN [15] and Decaf [26], adhere to the CCA component architecture stan dard. A framework can be CCA-compliant and still be tailored to the needs of

a particular computing environment. For example, Ccaffeine was designed to 164 support parallel computing, and XCAT was designed to support distributed 165 computing. Decaf [26] was designed by the developers of Babel primarily as 166 a means of studying the technical aspects of the CCA standard itself. The 167 important point is that each of these frameworks adheres to the same stan-168 dard, thus facilitating reuse of a (CCA) component in another computational 169 setting. The key idea is to isolate the components themselves, as much as 170 possible, from the details of the computational environment in which they 171 are deployed. If this is not done, then we fail to achieve one of the main goals 172 of component programming: code reuse. 173

CCA has been shown to be interoperable with Earth System Modeling 174 Framework (ESMF) [20] and Model Coupling Toolkit (MCT) [27, 28, 36, 175 43]. CSDMS has also demonstrated that it is interoperable with a Java 176 version of Open Modeling Interface (OpenMI) [44]. Many of the papers in 177 our cited references have been written by CCA Forum members and are 178 helpful for learning more about CCA. The CCA Forum has also prepared 170 a set of tutorials called "A Hands-On Guide to the Common Component 180 Architecture" [11]. 181

182 2.2. Language Interoperability with Babel

Babel [29, 14] is an open-source, language interoperability tool (consisting of a compiler and runtime) that automatically generates the "glue code" necessary for components written in different computer languages to communicate. As illustrated in Fig. 1, Babel currently supports C, C++, Fortran (77, 90, 95, and 2003), Java and Python. Babel is much more than a "least common denominator" solution; it even enables passing of variables with



Figure 1: Language interoperability provided by Babel.

data types that may not normally be supported by the target language (e.g., objects and complex numbers). Babel was designed to support *scientific*, *high-performance* computing and is one of the key tools in the CCA tool chain. It won an R&D 100 design award in 2006 for "The world's most rapid communication among many programming languages in a single application." It has been shown to outperform similar technologies such as CORBA and Microsoft's COM and .NET.

In order to create the glue code needed for two components written in different programming languages to exchange information, Babel needs to know only about the interfaces of the two components. It does not need any implementation details. Babel was therefore designed so that it can ingest a description of an interface in either of two fairly "language-neutral" forms, XML (eXtensible Markup Language) or SIDL (Scientific Interface

Definition Language). The SIDL language (somewhat similar to CORBA's 202 IDL) was developed for the Babel project. Its sole purpose is to provide a 203 concise description of a scientific software component interface. This inter-204 face description includes complete information about a component's inter-205 face, such as the data types of all arguments and return values for each of 206 the component's methods (or member functions). SIDL has a complete set 207 of fundamental data types to support scientific computing, from Booleans 208 to double-precision complex numbers. It also supports more sophisticated 209 data types such as enumerations, strings, objects, structs, and dynamic multi-210 dimensional arrays. The syntax of SIDL is similar to that of Java. A com-211 plete description of SIDL syntax and grammar can be found in "Appendix 212 B: SIDL Grammar" in the Babel User's Guide [14]. Complete details on how 213 to represent a SIDL interface in XML are given in "Appendix C: Extensible 214 Markup Language (XML)" of the same guide. 215

216 2.3. The Ccaffeine Framework

Ccaffeine [1] is the most widely used CCA framework, providing the run-217 time environment for sequential or parallel components applications. Us-218 ing Ccaffeine, component-based applications can run on diverse platforms, 219 including laptops, desktops, clusters, and leadership-class supercomputers. 220 Ccaffeine provides some rudimentary MPI communicator services, although 221 individual components are responsible for managing parallelism internally 222 (e.g., communicating data to and from other distributed components). A 223 CCA framework provides *services*, which include component instantiation 224 and destruction, connecting and disconnecting of ports, handling of input 225 parameters, and control of MPI communicators. Ccaffeine was designed pri-226

marily to support the single-component multiple-data (SCMD) programming 227 style, although it can support multiple-component multiple-data (MCMD) 228 applications that implement more dynamic management of parallel resources. 229 The CCA specification also includes an event service description, but it is 230 not fully implemented in Ccaffeine vet. Multiple interfaces to configuring 231 and executing component applications within the Ccaffeine framework exist, 232 including a simple scripting language, a graphical user interface, and the abil-233 ity to take over some of the operations normally handled by the frameworks, 234 such as component instantiation and port connections. 235

A typical CCA component's execution consists of the following steps:

- The framework loads the dynamic library for the component. Static linking options are also available.
- The component is instantiated. The framework calls the setServices
 method on the component, passing a handle to itself as an argument.
- User-specified connections to other components' ports are established
 by the framework.
- If the component provides a gov.cca.ports.Go port (similar to a "main" subroutine), its go() method can be invoked to start the main portion of the computation.
- Connections can be made and broken throughout the life of the component.
- All component ports are disconnected, and the framework calls releaseServices prior to calling the component's destructor.

The handle to the framework services object, which all CCA components obtain shortly after instantiation, can be used to access various framework services throughout the component's execution. This represents the main difference between a class and a component: a component dynamically accesses another component's functionality through dynamically connecting ports (requiring the presence of a framework), whereas classes in objectoriented languages call methods directly on instances of other classes.

257 2.4. Component Development with Bocca

Bocca [2] is a tool in the CCA tool chain that was designed to help users create, edit, and manage a set of SIDL-based entities, including CCA components and ports, that are associated with a particular project. Once a set of CCA-compliant components and ports has been prepared, one can use a CCA-compliant framework such as Ccaffeine to link components from this set together to create applications or composite models.

Bocca was developed to address usability concerns and reduce the de-264 velopment effort required for implementing multilanguage component appli-265 cations. Bocca was designed specifically to free users from mundane, time-266 consuming, low-level tasks so they can focus on the scientific aspects of their 267 applications. It can be viewed as a development environment tool that al-268 lows application developers to perform rapid component prototyping while 269 maintaining robust software- engineering practices suitable to HPC envi-270 ronments. Bocca provides project management and a comprehensive build 271 environment for creating and managing applications composed of CCA com-272 ponents. Bocca operates in a language-agnostic way by automatically in-273 voking the Babel compiler. A set of Bocca commands required to create a 274

²⁷⁵ component project can be saved as a shell script, so that the project can
²⁷⁶ be rapidly rebuilt, if necessary. Various aspects of an existing component
²⁷⁷ project can also be modified by typing Bocca commands interactively at a
²⁷⁸ Unix command prompt.

While Bocca automatically generates dynamic libraries, a separate tool can be used to create *stand-alone executables* for projects by automatically bundling all required libraries on a given platform. Examples of using Bocca are available in the set of tutorials called "A Hands-On Guide to the Common Component Architecture," written by the CCA Forum members [11].

284 2.5. Other Component-Based Modeling Projects

We briefly discuss several other component-based projects in the area of Earth system-related modeling.

- The Object Modeling System (OMS) [35] is a pure Java, object-oriented
 framework for component-based agro-environmental modeling.
- The Open Modeling Interface (OpenMI) [44] is an open-source software-289 component interface standard for the computational core of numerical 290 models. Model components that comply with this standard can be con-291 figured without programming to exchange data during computation (at 292 runtime). Similar to the CCA component model, the OpenMI standard 293 supports two-way links between components so that the involved mod-294 els can mutually depend on calculation results from each other. Linked 295 models may run asynchronously with respect to time steps, and data 296 represented on different geometries (grids) can be exchanged by using 297 built-in tools for interpolating in space and time. OpenMI was designed 298

primarily for use on PCs, using either the .NET or Java framework.
 CSDMS has experimented with OpenMI version 1.4 (version 2.0 was
 recently released) but currently uses a simpler component interface.

- The Earth System Modeling Framework (ESMF) [18, 20] is software for building and coupling weather, climate, and related models written in Fortran. ESMF defines data structures, parallel data redistribution, and other utilities to enable the composition of multimodel high-performance simulations.
- The Framework for Risk Analysis of Multi-Media Environmental Systems (FRAMES) [19] is developed by the U.S. Environmental Protection Agency to provide models and modeling tools (e.g., data retrieval and analysis) for simulating different environmental processes.

311 3. Problem Definition – Component-based Plug-and-Play Model-312 ing

Next we discuss the challenges that we faced in tackling the problem of creating plug-and-play modeling capabilities that can be extended and actively used by the CSDMS community.

316 3.1. Attributes of Earth Surface Process Models

The Earth surface process modeling community has *numerous* models, but it is difficult to couple or reconfigure them to solve new problems. The reason is that they are a heterogeneous set.

• The models are written in *many different languages*, which may be object-oriented or procedural, compiled or interpreted, proprietary or

322	open-source, etc. Languages do not all offer the same data types and
323	features, so special tools are required to create "glue code" necessary
324	to make function calls across the <i>language barrier</i> .
325	The models typically are not designed to "talk" to each other and do
326	not follow any particular set of conventions.
327	The models generally have a $geographic$ context and are often used in
328	conjunction with GIS (Geographic Information System) tools.
329	The generally consist of one or more arrays (1D, 2D, or 3D) that are
330	being advanced in time according to differential equations or other rules $% \left({{{\bf{n}}_{\rm{s}}}} \right)$
331	(i.e., we are not modeling molecular dynamics).
332	The models use different input and output file formats.
333	The models are often open source. Even many models that were orig-
334	inally sold commercially are now available as open-source code, for

example parts of Delt3D from Deltares and many EDF (Energie du Francais) models.

337 3.2. Difficulties in Linking Models

Linking together models that were not specifically designed from the outset to be linkable is often surprisingly difficult, and a brute-force approach to the problem often requires a significant investment of time and effort. The main reason is that two models may differ in may ways. The following list of possible differences illustrates this point.

• The models are written in different languages, making conversion timeconsuming and error-prone.

345	• The person doing the linking may not be the author of either model,
346	and the code is often not well-documented or easy to understand.
347	• Models may have different dimensionality (1D, 2D, or 3D).
348	• Models may use different types of grids (e.g., rectangles, triangles, and
349	Voronoi cells).
350	• Each model has its own time loop or "clock."
351	• The numerical scheme may be either explicit or implicit.
352	3.3. Design Criteria
353	The technical goals of a component-based modeling system include the
354	following.
355	• Support for <i>multiple operating systems</i> (especially Linux, Mac OS X,
355 356	• Support for <i>multiple operating systems</i> (especially Linux, Mac OS X, and Windows).
355 356 357	 Support for <i>multiple operating systems</i> (especially Linux, Mac OS X, and Windows). Language interoperability to support code contributions written in pro-
355 356 357 358	 Support for <i>multiple operating systems</i> (especially Linux, Mac OS X, and Windows). Language interoperability to support code contributions written in procedural languages (e.g., C or Fortran) as well as object-oriented lan-
355 356 357 358 359	 Support for multiple operating systems (especially Linux, Mac OS X, and Windows). Language interoperability to support code contributions written in procedural languages (e.g., C or Fortran) as well as object-oriented languages (e.g., Java, C++, and Python).
355 356 357 358 359 360	 Support for multiple operating systems (especially Linux, Mac OS X, and Windows). Language interoperability to support code contributions written in procedural languages (e.g., C or Fortran) as well as object-oriented languages (e.g., Java, C++, and Python). Support for both structured and unstructured grids, requiring a spatial
 355 356 357 358 359 360 361 	 Support for multiple operating systems (especially Linux, Mac OS X, and Windows). Language interoperability to support code contributions written in procedural languages (e.g., C or Fortran) as well as object-oriented languages (e.g., Java, C++, and Python). Support for both structured and unstructured grids, requiring a spatial regridding tool.
 355 356 357 358 359 360 361 362 	 Support for multiple operating systems (especially Linux, Mac OS X, and Windows). Language interoperability to support code contributions written in procedural languages (e.g., C or Fortran) as well as object-oriented languages (e.g., Java, C++, and Python). Support for both structured and unstructured grids, requiring a spatial regridding tool. Platform-independent GUIs and graphics where useful.
 355 356 357 358 359 360 361 362 363 	 Support for multiple operating systems (especially Linux, Mac OS X, and Windows). Language interoperability to support code contributions written in procedural languages (e.g., C or Fortran) as well as object-oriented languages (e.g., Java, C++, and Python). Support for both structured and unstructured grids, requiring a spatial regridding tool. Platform-independent GUIs and graphics where useful. Use of well-established, open-source software standards whenever pos-

- Use of open-source tools that are mature and have well-established communities, avoiding dependencies on proprietary software whenever possible (e.g., Windows, C#, and Matlab).
- Support for *parallel computation* (multiprocessor, via MPI standard).
- Interoperability with other coupling frameworks. Since code reuse is a fundamental tenet of component-based modeling, the effort required to use a component in another framework should be kept to a minimum.
- *Robustness and ease of maintainenance*. It will clearly have many software dependencies, and this software infrastructure will need to be
 updated on a regular basis.
- Use of *HPC tools and libraries*. If the modeling system runs on HPC architectures, it should strive to use parallel tools and models (e.g., VisIt, PETSc, and the ESMF regridding tool).
- Familiarity. Model developers and contributors should not be required to make major changes to how they work.

Expanding the last bullet, developers should not be required to convert 380 to another programming language or use invasive changes to their code (e.g., 381 use specified data structures, libraries, or classes). They should be able to 382 retain "ownership" of the code and make continual improvements to it; some-383 one should be able to componentize future, improved versions with minimal 384 additional effort. The developer will likely want to continue to use the code 385 outside the framework. However, some degree of code refactoring (e.g., break-386 ing code into functions or adding a few new functions) and ensuring that the 387

code compiles with an open-source compiler are considered reasonable requirements. It is also expected that many developers will take advantage of various built-in tools if doing so is straightforward and beneficial.

391 3.4. Interface vs. Implementation

The word *interface* may be the most overloaded word in computer science. In each case, however, it adheres to the standard, English meaning of the word that has to do with a boundary between two items and what happens at the boundary.

Many people hear the word interface and immediately think of the in-396 terface between a human and a computer program, which is typically either 397 a command-line interface or a graphical user interface (GUI). While such in-398 terfaces are an interesting and complex subject, this is usually not what 390 computer scientists are talking about. Instead, they tend to be interested 400 in other types of interface, such as the one between a pair of software com-401 ponents, or between a component and a framework, or between a developer 402 and a set of utilities (i.e., an API or a software development kit). 403

Within the present context of component programming, we are interested 404 primarily in the interfaces between components. In this context, the word 405 interface has a specific meaning, essentially the same as in the Java pro-406 gramming language. An interface is a user-defined entity/type, similar to 407 an abstract class. It does not have any data fields, but instead is a named 408 set of methods or member functions, each defined completely with regard to 409 argument types and return types but without any actual implementation. A 410 CCA *port* is simply this type of interface. Interfaces are the name of the 411 game when it comes to the question of reusability or "plug and play." Once 412

an interface has been defined, one can ask the question: Does this compo-413 nent have interface A? To answer the question, we merely have to look at the 414 methods (or member functions) that the component has with regard to their 415 names, argument types, and return types. If a component does have a given 416 interface, then it is said to *expose* or *implement* that interface, meaning that 417 it contains an actual *implementation* for each of those methods. It is fine 418 if the component has additional methods beyond the ones that constitute a 419 particular interface. Thus, it is possible (and frequently useful) for a single 420 component to expose multiple, different interfaces or ports. For example, 421 multiple interfaces may allow a component to be used in a greater variety 422 of settings. An analogy exists in computer hardware, where a computer or 423 peripheral may actually have a number of different ports (e.g., USB, serial, 424 parallel, and ethernet) to enable it to communicate with a wider variety of 425 other components. 426

The distinction between *interface* and *implementation* is an important 427 theme in computer science. The word pair *declaration* and *definition* is used 428 in a similar way. A function (or class) declaration tells what the function 429 does (and how to interact with or use it) but not how it works. To see how 430 the function actually works, we need to look at how it has been defined or 431 implemented. C and C++ programmers are familiar with this idea, which 432 is similar to declaring variables, functions, classes, and other data types in a 433 header file with the file name extension .h or .hpp, and then defining their 434 implementations in a separate file with extension .c or .cpp. 435

⁴³⁶ Of course, most of the gadgets that we use every day (from iPods to cars) ⁴³⁷ are like this. We need to understand their interfaces in order to use them (and interfaces are often standardized across vendors), but often we have no
idea what is happening inside or how they actually work, which may be quite
complex.

While the tools in the CCA tool chain are powerful and general, they do 441 not provide a ready interface for linking geoscience models (or any domain-442 specific models). In CCA terminology, *port* is essentially a synonym for 443 interface and a distinction is made between ports that a given component uses 444 (uses ports), and those that it provides (provides ports) to other components. 445 Note that this model provides a means of bidirectional information exchange 446 between components, unlike dataflow-based approaches (e.g., OpenMI) that 447 support unidirectional links between components (i.e., the data produced by 448 one component is consumed by another component). 449

Each scientific modeling community that wishes to make use of the CCA 450 tools is responsible for designing or selecting component interfaces (or ports) 451 that are best suited to the kinds of models they wish to link together. This is 452 a big job that involves social as well as technical issues and typically requires 453 a significant time investment. In some disciplines, such as molecular biology 454 or fusion research, the models may look quite different from ours. Ours tend 455 to follow the pattern of a 1D, 2D or 3D array of values (often multiple, 456 coupled arrays) advancing in time. However, our models can still be quite 457 different from each other with regard to their dimensionality or the type 458 of computational grid they use (e.g., rectangles, triangles or polygons), or 459 whether they are implicit or explicit in time. 460

461 3.5. Granularity

While components may represent any level of granularity, from a simple 462 function to a complete hydrologic model, the optimum level appears to be 463 that of a particular physical process, such as infiltration, evaporation, or 464 snowmelt. At this level of granularity researchers are most often interested 465 in swapping out one method of modeling a process for another. A simpler 466 method of parameterizing a process may apply only to simplified special cases 467 or may be used simply because there is insufficient input data to drive a more 468 complex model. A different numerical method may solve the same governing 460 equations with greater accuracy, stability, or efficiency and may or may not 470 use multiple processors. Even the same method of modeling a given process 471 may exhibit improved performance when coded in a different programming 472 language. But the physical process level of granularity is also natural for 473 other reasons. Specific physical processes often act within a domain that 474 shares a physically important boundary with other domains (e.g., coastline 475 and ocean-atmosphere), and the fluxes between these domains are often of 476 key interest. In addition, experience shows that this level of granularity 477 corresponds to GUIs and HTML help pages that are more manageable for 478 users. 479

⁴⁸⁰ A judgment call is frequently needed to decide whether a new feature ⁴⁸¹ should be provided in a separate component or as a configuration setting ⁴⁸² in an existing component. For example, a kinematic wave channel-routing ⁴⁸³ component may provide both Manning's formula and the law of the wall as ⁴⁸⁴ different options to parameterize frictional momentum loss. Each of these ⁴⁸⁵ options requires its own set of input parameters (e.g., Manning's *n* or the roughness parameter, z_0). We could even think of frictional momentum loss as a separate physical process, under which we would have a separate Manning's formula and law of the wall components. Usually, the amount of code associated with the option and usability considerations can be used to make these decisions.

Some models are written in such a way that decomposing them into sep-491 arate process components is not really appropriate, because of some special 492 aspect of the model's design or because decomposition would result in an 493 unacceptable loss of performance (e.g., speed, accuracy, or stability). For 494 example, *multiphysics models*—such as Penn State Integrated Hydrologic 495 Model (PIHM)—represent many physical processes as one large, coupled set 496 of ODEs that are then solved as a matrix problem on a supercomputer. 497 Other models involve several physical processes that operate in the same do-498 main and are relatively tightly coupled within the governing equations. The 490 Regional Ocean Modeling System (ROMS) is an example of such a model, 500 in which it may not be practical to model processes such as tides, currents, 501 passive scalar transport (e.g., T and S), and sediment transport within sep-502 arate components. In such cases, however, it may still make sense to wrap 503 the entire model as a component so that it may interact with other models 504 (e.g., an atmospheric model, such as WRF, or a wave model, such as SWAN) 505 or be used to drive another model (e.g., a Lagrangian transport model, such 506 as LTRANS). 507

⁵⁰⁸ 4. Designing a Modeling Interface

A component interface is simply a named set of functions (called methods) that have been defined completely in terms of their names, arguments and return values. The purpose of this section is to explain the types of functions that are required and why. The functions that define an interface are somewhat analogous to the buttons on a handheld remote control—they provide a caller with fine-grained control of the model component.

515 4.1. The "IRF" Interface Functions

Most Earth-science models initialize a set of state variables (often as 1D, 516 2D, or 3D arrays) and then execute of series of timesteps that advance the 517 variables forward in time according to physical laws (e.g., mass conservation) 518 or some other set of rules. Hence, the underlying source code tends to follow 519 a standard pattern that consists of three main parts. The first part consists 520 of all source code prior to the start of the time loop and serves to set up 521 or *initialize* the model. The second part consists of all source code within 522 the time loop and is the guts of the model where state variables are updated 523 with each time step. The third part consists of all source code after the 524 end of the time loop and serves to tear down or *finalize* the model. Note 525 that root-finding and relaxation algorithms follow a similar pattern even if 526 the iterations do not represent timestepping. A time-independent model 527 can also be thought of as a time-stepping model with a single time step. 528 For maximum plug-and-play flexibility, each of these three parts must be 529 encapsulated in a separate function that is accessible to a caller. It turns out 530 that we get more flexibility if the function for the middle phase is written to 531

⁵³² accept the start time and end time as arguments.

For lack of a better term, we refer to this Initialize(), Run_Until(), Finalize() pattern as an *IRF interface*. All of the model coupling projects that we are aware of use this pattern as part of their component interface, including CSDMS, ESMF, OMF, and OpenMI. An IRF interface is also used as part of the Message Passing Interface (MPI) for communication between processes in high-performance computers.

To see how an IRF interface is used when coupling models, consider two models, Models A and B, that do not have this interface. To combine them into a single model, where one uses the output of the other during its time loop, we would need to cut the code from within Model A's time loop and paste it into Model B, or vice versa. The reason is that both models were designed to control the time loop and cannot reliquish this control.

545 4.1.1. Initialize (Model Setup)

The initialize step puts a model into a valid state that is ready to be executed. Mostly this involves initializing variables or grids that will be used within the execution step. Temporary files that the execution step will read from or write to should also be opened here.

550 4.1.2. Run_Until (Model Execution)

The run step advances the model from its current state to a future state. For time-independent models the run step simply executes the model calculation and updates the model state so that future calls will not require executing the calculations again. Encapsulating only the code *within* the time loop allows an application to run the model to intermediate states. This is necessary to allow an application to query the model's state for the purposes of (for instance) printing output or passing state data to another model.

559 4.1.3. Finalize (Model Termination)

The finalize step cleans up after the model is no longer needed. The main purpose of this step to make sure that all resources a model acquired through its life have been freed. Most often this will be freeing allocated memory, but it could also be freeing file or network handles. Following this step, the model should be left in an invalid state such that its run step can no longer be called until it has been initialized again.

566 4.2. Getter and Setter Interface Functions

A basic IRF interface, while important, really provides only the core functionality of a model coupling interface. A complete interface will also require functions that enable another component to request data from the component (a getter) or change data values (a setter) in the component. These are typically called within the Initialize() or Run_Until() methods.

572 4.2.1. Value Getters

Limiting access to the model's state to be through a set of functions allows control of what data the model shares with other programs and how it shares that data. The data may be transferred in two ways. The first is to give the calling program a copy of the data. The second is to give the actual data that is being used by the model (in C, this would mean passing a pointer to a value). The first has the advantage that it hides implementation details of the model from the calling program and limits what the calling program can do to the model. However, the downside of the first method is
that communication will be slower (and could be significantly so, depending
on the size of the data being transferred).

583 4.2.2. Value Setters

Variables in a model should be accessed and changed only through in-584 terface methods. This approach ensures that users of the interface are not 585 able to change values that the interface implementor does not want them 586 to change. This also detaches the programmer using the interface from the 587 model implementation, thus freeing the model developer to change details of 588 the model without an application programmer having to make any changes. 589 The setter can also perform tasks other than just setting data. For in-590 stance, it might be useful if the setter checked to make sure that the new 591 data is valid. After the setter method sets the data, it should ensure that 592 the model is still in a valid state. 593

The Get_Value() and Set_Value() methods can be general in terms of supporting different grid or mesh types, but it should be possible to bypass that generality and use simple, raster-based grids to keep things simple and efficient when the generality is not needed.

⁵⁹⁸ CSDMS has wrapped two open-source regridding tools that can act as ⁵⁹⁹ services (see Section 9) that other components can use when communicating ⁶⁰⁰ with one another (an example regridding scenario is shown in Figure 2). The ⁶⁰¹ first is from the ESMF project. It is implemented in Fortran and is designed ⁶⁰² to use multiple processors on a distributed memory system. It supports ⁶⁰³ sophisticated options such as mass-conservative interpolation. The second ⁶⁰⁴ tool is the multithreaded tool included in the Java SDK for OpenMI.



(c) Voronoi cells before regridding.(d) After regridding to raster cells.Figure 2: Regridding example.

The Get_Value() and Set_Value() methods should optionally allow specification (via indices) of which individual elements within an array that are to be obtained or modified. We often need to manipulate just a few values, and we don'twant to transfer copies of entire arrays (which may be large) unless necessary.

Each component should understand what variables will be requested from

it; and if those represent some function of its state variables (e.g., a sum
or product), then that computation should be done by the component and
offered as an output variable rather than passing several state variables that
must then be combined in some way by the caller.

In order to support dynamically typed languages like Python, additional interface functions may be required in order to query whether the variable is currently a scalar or a vector (1D array) or a grid.

618 4.3. Self-Descriptive Interface Functions

Two additional methods for a modeling interface would enable a caller to 619 query what type of data the component is able to use as input or compute 620 as output. These would typically not require arguments and would simply 621 return the names of all the possible input or output variables as an array of 622 strings, for example Get_Input_Item_List() and Get_Output_Item_List(). An-623 other type of self-descriptive function would be a function like Get_Status() 624 that returns the component's current status as a string from a standardized 625 list. 626

627 4.4. Framework Interface Functions

A component typically needs some additional methods that allow it to be instantiated by and communicate with a component-coupling framework. For example, a component must implement methods called __init__(), getServices(), and releaseServices() in order to be used within a CCA-compliant framework.

633 4.5. Autoconnection Problem

A key goal of component-based modeling is to create a collection of com-634 ponents that can be coupled together to create new and useful composite 635 models. This goal can be achieved by providing every component with the 636 same interface, and this is the approach used by OpenMI. A secondary goal, 637 however, is for the coupling process to be as automatic as possible, that is, 638 to require as little input as possible from users. To achieve this goal, we need 639 some way to group components into categories according to the functionality 640 they provide. This grouping must be readily apparent to both a user and the 641 framework (or system) so that it is clear whether a particular pair of compo-642 nents are *interchangeable*. But what should it mean for two components to 643 be interchangeable? Do they really need to use identical input variables and 644 provide identical output variables? Our experience shows that this definition 645 of interchangeable is unnecessarily strict. 646

To bring these issues into sharper focus, consider the physical process of 647 infiltration, which plays a key role in hydrologic models. As part of a larger 648 hydrologic model, the main purpose of an infiltration component is to com-649 pute the infiltration rate at the surface, because it represents a loss term in 650 the overall hydrologic budget. If the domain of the infiltration component 651 is restricted to the unsaturated zone, above the water table, then it may 652 also need to provide a vertical flow rate at the water table boundary. Thus, 653 the main job of the infiltration component is to provide fluxes at the (top 654 and bottom) boundaries of its domain. To do this job, it needs variables 655 such as flow depth and rainfall rate that are outside its domain and com-656 puted by another component. Hydrologists use a variety of different methods 657

and approximations to compute surface infiltration rate. The Richards 3D 658 method, for example, is a more rigorous approach that tracks four state vari-659 ables throughout the domain; on the other hand, the Green-Ampt method 660 makes a number of simplifying assumptions so that it computes a smaller 661 set of state variables and does not resolve the vertical flow dynamics to the 662 same level of detail (i.e., piston flow, sharp wetting front). As a result, the 663 Richards 3D and Green-Ampt infiltration components use a different set of 664 input variables and provide a different set of output variables. Nevertheless, 665 they both provide the surface infiltration rate as one of their outputs and can 666 therefore be used "interchangeably" in a hydrologic model as an "infiltration 667 component." 668

The infiltration example illustrates several key points that are transfer-669 able to other situations. Often a model, such as a hydrologic model, breaks 670 the larger problem domain into a set of subdomains where one or more pro-671 cesses are relevant. The boundaries of these subdomains are often physical 672 interfaces, such as surface/subsurface, unsaturated/saturated zone, atmo-673 sphere/ocean, ocean/seafloor, or land/water. Moreover, the variables that 674 are of interest in the larger model often depend on the fluxes across these 675 subdomain boundaries. 676

Within a group of interchangeable components (e.g., infiltration components), there are many other implementation differences that a modeler may wish to explore, beyond just how a physical process is parameterized. For example, performance and accuracy often depend on the numerical scheme (explicit vs. implicit, order of accuracy, stability), data types used (float vs. double), number of processors (parallel vs. serial), approximations used, the ⁶⁸³ programming language, or coding errors.

Autoconnection of components is important from a user's point of view. 684 Components typically require many input variables and produce many out-685 put variables. Users quickly become frustrated when they need to manually 686 create all these pairings/connections, especially when using more than just 687 two or three components at a time. The OpenMI project does not support 688 the concept of auto-connection or interchangeable components. When using 689 the graphical Configuration Editor provided in its SDK, users are presented 690 with droplists of input and output variables and must select the ones to be 691 paired. Doing so requires expertise and is made more difficult because there 692 is so far no ontological or semantic scheme to clarify whether two variable 693 names refer to the same item. 694

The CSDMS project currently employs an approach to autoconnection that involves providing interfaces (i.e. ,CCA ports) with different names to reflect their intended use (or interchangeability), even though the interfaces are the same internally.

⁶⁹⁹ 5. Current CSDMS Component Interface

This section contains a concise list of the current CSDMS IRF and getter/setter interfaces, which must be implemented by any compliant components.

703 5.1. The IRF Interface

The following methods comprise the IRF interface described in more detail in Section 4.1.

```
CMI_INITIALIZE (handle, filename)
706
     OUT
                   handle
                                      handle to the CMI object
707
                                      path to configuration file
     IN
                   filename
708
    CMI_RUN_UNTIL (handle, stop_time)
709
     IN
                   handle
                                      handle to the CMI object
710
     IN
                                      simulation time to run model until
                   stop_time
711
    CMI_FINALIZE (handle)
712
     INOUT
                   handle
                                      handle to the CMI object
713
714
    5.2. Value Getters and Setters
715
       The following methods comprise the CSDMS getter/setter interface dis-
716
    cussed in Section 4.2.
717
    CMI_GRID_DIMEN (handle, value_str, dimen)
718
     IN
                   handle
                                      handle to the CMI object
     IN
                  value_str
                                      name of the value to get
719
                                      length of each grid dimension
     OUT
                   dimen
    CMI_GRID_RES (handle, value_str, res)
720
     IN
                   handle
                                      handle to the CMI object
     IN
                  value_str
                                      name of the value to get
721
     OUT
                                      grid spacing for each dimension
                   res
```

722 CMI_GET_GRID_DOUBLE (handle, value_str, buffer)

	IN	handle	handle to the CMI object
723	IN	value_str	name of the value to get
	OUT	buffer	initial address of the destination values
724	CMI_SET_GRI	D_DOUBLE (handle	e, value_str, buffer, dimen)
	IN	handle	handle to the CMI object
725	IN	value_str	name of the value to get
125	IN	buffer	initial address of the source values
	IN	dimen	grid dimension
726	CMI_GET_TIN	/IE_SPAN (handle, s	pan)
	IN	handle	handle to the CMI object
727	OUT	span	start and end times for the simulation
728	CMI_GET_ELE	EMENT_SET (handl	e, value_str, element_set)
	IN	handle	handle to the CMI object
729	IN	value_str	name of the value to get
	OUT	buffer	model ElementSet
730	CMI_GET_VAI	LUE_SET (handle, v	alue_str, value_set)
	IN	handle	handle to the CMI object
731	IN	value_str	name of the value to get
	OUT	buffer	model ValueSet
732	CMI_SET_VAL	_UE_SET (handle, va	alue_str, value_set)
	IN	handle	handle to the CMI object
733	IN	value_str	name of the value to get
	IN	buffer	model ValueSet

734 6. Component Wrapping Issues

In this section we discuss several methods for creating components based
on existing codes by using an approach often referred to as *wrapping*.

737 6.1. Code Reuse and the Case for Wrapping

Using computer models to simulate, predict, and understand Earth sur-738 face processes is not a new idea. Many models exist, some of which are fairly 739 sophisticated, comprehensive, and well tested. The difficulty with reusing 740 these models in new contexts or linking them to other models typically has 741 less to do with how they are implemented and more to do with the interface 742 through which they are called (and to some extent, the implementation lan-743 guage.) For a small or simple model, little effort may be needed to rewrite 744 the model in a preferred language and with a particular interface. Rewriting 745 large models, however, is both time-consuming and error prone. In addition, 746 most large models are under continual development, and a rewritten version 747 will not see the benefits of future improvements. Thus, for code reuse to be 748 practical, we need a *language interoperability tool*, so that components dont 749 need to be converted to a different language, and a wrapping procedure that 750 allows us to provide existing code with a new calling interface. As suggested 751 by its name, and the fact that it applies to the "outside" (interface) of a com-752 ponent vs. its "inside" (implementation), wrapping tends to be noninvasive 753 and is a practical way to convert existing models into components. 754

755 6.2. Wrapping for Object-Oriented Languages

⁷⁵⁶ Component-based programming is essentially object-oriented program-⁷⁵⁷ ming with the addition of a framework. If a model has been written as a

class, then it is relatively straightforward to modify the definition of this 758 class so that it exposes a particular model-coupling interface. Specifically, 759 one could add new methods (member functions) that call existing methods, 760 or one could modify the existing methods. Each function in the interface 761 has access to all of the state variables (data members) without passing them 762 explicitly; it also has access to all the other interface functions. In object-763 oriented languages one commonly distinguishes between private methods that 764 are intended for internal use by the model object and *public methods* that are 765 to be used by callers and that may comprise one or more interfaces. (Some 766 languages, like Java, make this part of a method's declaration.) 767

In order for this model object to be used as a component in a CCA-768 compliant framework like Ccaffeine, it must also be "wrapped" by a CCA 769 implementation file (or IMPL file). The CCA tool chain has tools such as 770 Babel and Bocca that are used to autogenerate an IMPL-file template. For 771 a model that is written in an object-oriented and Babel-supported language 772 (e.g., C++, Python, or Java), the IMPL file needs to do little more than 773 add interface functions like setServices and releaseServices that allow the 774 component to communicate with and be instantiated by the framework. The 775 interface functions used for intercomponent communication (i.e., passing data 776 and IRF) can simply be inherited from the model class. Inheritance is a 777 standard mechanism in object-oriented languages that allows one interface 778 (set of methods) to be extended or overridden by another. Note that the 779 IMPL file may have its own Initialize() function that first gets the required 780 CCA ports and then calls the Initialize() function in the model's interface. 781 But the function that gets the CCA ports can simply be another function 782

in the model's interface that is used only in this context. Similarly, the 783 IMPL file may have a Finalize() function that calls the Finalize() function 784 of the model and then calls a function to release the CCA ports that are no 785 longer needed. It is desirable to keep the IMPL files as clean as possible, 786 which means adding some CCA-specific functions to the model's interface. 787 For example, a CSDMS component would have (1) functions to get and 788 release the required CCA ports, (2) a function to create a tabbed-dialog 789 (using CCA's so-called parameter ports), and (3) a function that prints a 790 language-specific traceback to stdout if an exception occurs during a model 791 run. 792

793 6.3. Wrapping for Procedural Languages

Languages such as C or Fortran (up to 2003) do not provide objectoriented primitives for encapsulating data and functionality. Because componentbased programming requires such encapsulation, the CCA provides a means to produce object-oriented software even in languages that do not support it directly. We briefly describe the mechanism for creating components based on functionality implemented in a procedural language (e.g., an existing library or model).

A class in object-oriented terminology encapsulates some set of related functions and associated data. To wrap a set of library functions, one can create a SIDL interface or class that contains a set of methods whose implementations call the legacy functions. The new interface does not have to mirror existing functions exactly, presenting a nonintrusive opportunity for redesigning the publicly accessible interfaces presented by legacy software. The creation of class or component wrappers also enables the careful definition of namespaces, thus reducing potential conflicts when integrating with other classes or components. The SIDL definitions are processed by Babel to generate IMPL files in the language of the code being wrapped. The calls to the legacy library can then be added either manually or by a tool, depending on how closely the SIDL interface follows the original library interface.

Function argument types that appear in the SIDL definition can be han-813 dled in two ways: by using a SIDL type or by specifying them as *opaque*. 814 SIDL already supports most basic types and different kinds of arrays found 815 in the target languages. Any user-defined types (e.g., structs in C or de-816 rived types in Fortran) must have SIDL definitions or be passed as opaques. 817 Because opaques are not accessible from components implemented in a dif-818 ferent language, they are rarely used. Model state variables that must be 819 shared among components can be handled in a couple of ways. They can 820 be encapsulated in a SIDL class and accessed through get/set methods (e.g., 821 as described in Section 4.2). Recently Babel has added support for defining 822 structs in SIDL, whose data members can be accessed directly from multiple 823 languages. 824

SIDL supports namespacing of symbols through the definition of packages 825 whose syntax and semantics are similar to Java's packages. In languages that 826 do not support object orientation natively, symbols (e.g., function names) 827 are prefixed with the names of all enclosing packages and parent class. This 828 approach greatly reduces the potential build-, link-, or runtime name conflicts 829 that can result when multiple components define the same interfaces (e.g., 830 the initialize, run, and finalize methods). These naming conventions can be 831 applied to any code, not only SIDL-based components. 832

Implementors working in non object-oriented languages should encapsu-833 late their model's state data in an object that is opaque to the application 834 programmer. Memory within the object is not directly accessible by the user 835 but can be accessed through an opaque handle, which exists in user space. 836 This handle is passed as the first argument to each of the interface functions 837 so that they can operate on a particular instance of a model. For example, 838 in C, this handle could simply be a pointer to the object and in Fortran, the 839 handle could be an index into a table of opaque objects in a system table. 840

Model handles are allocated and deallaocated in the initialize and finalize interface functions, respectively. For allocate calls, the initialize functions are passed an OUT argument that will contain a valid reference to the object. For deallocation, the finalize function accepts an INOUT variable that provides a reference to the object to be destroyed and sets the object to an invalid state.

847 6.4. Guidelines for Model Developers

⁸⁴⁸ Developers can follow several relatively simple follow so that it becomes ⁸⁴⁹ much easier to create a reusable, plug-and-play component from their model ⁸⁵⁰ source code. Given the large number of models that are contributed to the ⁸⁵¹ CSDMS project, it is much more efficient for model developers to follow ⁸⁵² these guidelines and thereby "meet us halfway" than for CSDMS staff to ⁸⁵³ make these changes after code has been contributed. This can be thought of ⁸⁵⁴ as a form of load balancing.

855 6.4.1. Programming Language and License

- Write code in a Babel-supported language (C, C++, Fortran, Java, Python).
- If code is in Matlab or IDL, use tools like I2PY to convert it to Python.
 Python (with the numpy, scipy, and matplotlib packages) provides a
 free work-alike to Matlab with similar performance.
- Make sure that code can be compiled with an open-source compiler (e.g., gcc and gfortran).
- Specify what type of open-source license applies to your code. Rosen
 [41] provides a good, online, and open-source book that explains open source licensing in detail. CSDMS requires that contributions have an
 open source license type that is compliant with the standard set forth
 by the Open Source Initiative.
- 868 6.4.2. Model Interface
- Refactor the code to have the basic IRF interface (5.1).
- If code is in C or Fortran, add a model name prefix to all interface
 functions to establish a namespace (e.g., ROMS_Initialize()). C code
 can alternatively be compiled as C++.
- Write Initialize() and Run_Until() functions that will work whether the component is used as a driver or *nondriver*.
- Provide getter and setter functions (4.2.1).
- Provide functions that describe input and output *exchange items* (4.2.1).

- Use descriptive function names (e.g., Update_This_Variable).
- Remove user interfaces, whether graphical, command line or otherwise,
 from your interface implementation. This avoids incompatible user
 interfaces competing with one another.
- 881 6.4.3. State Variables
- Decide on an appropriate set of state variables to be maintained by the component and made available to callers.
- Attempt to minimize data transfer between components (as discussed above).
- Use descriptive variable names.
- Carefully track each variable's units.
- 888 6.4.4. Input and Output Files
- Do not hardwire configuration settings in the code; read them from a configuration file (text).
- Do not use hardwired input filenames.
- Read configuration settings from text files (often in Initialize()). Do
 not prompt for command-line input. If a model has a GUI, write code
 so it can be bypassed; use the GUI to create a configuration file.
- Design code to allow separate input and output directories that are read from the configuration file. This approach allows many users to use the same input data without making copies (e.g., test cases). It is

- frequently helpful to include a *case prefix* (scenario) and a *site prefix* (geographic name) and use them to construct default output filenames.
- Establish a namespace for configuration files (e.g., ROMS_input.txt vs. input.txt).
- If large arrays are to be stored in files, save them as binary vs. text. (e.g., this is the case with NetCDF)
- Provide self-test functions or unit tests and test data. One self-test could simply be a "sanity check" that uses trivial (perhaps hard-coded) input data. When analytic solutions are available, these make excellent self-tests because they can also be used to check the accuracy and stability of the numerical methods.
- Do not create and write to output files within the interface implementa-⁹¹⁰ tion. If this is not possible, output files should be well documented and ⁹¹¹ allow for a naming convention that reduces the possibility of naming ⁹¹² conflicts.
- 913 6.4.5. Documentation
- Help CSDMS to provide a standardized, HTML help page.
- Help CSDMS to provide a standaridized, tabbed-dialog GUI.
- Make liberal use of comments in the code.

917 7. The CSDMS Modeling Tool (CMT)

As explained in Section 2.3, Ccaffeine is a CCA-compliant framework for connecting components to create applications. From a user's point of

view, Ccaffeine is a low-level tool that executes a sequence of commands in a 920 Ccaffeine script. The (natural language) commands in the Ccaffeine scripting 921 language are fairly straightforward, so it is not difficult for a programmer to 922 write one of these scripts. For many people, however, using a graphical 923 user interface (GUI) is preferable because they don'thave to learn the syntax 924 of the scripting language. A GUI also provides users with a natural, visual 925 representation of the connected components as boxes with buttons connected 926 by wires. It can also prevent common scripting errors and offer a variety of 927 other convenient features. The CCA Forum developed such a GUI, called 928 Ccafe-GUI, that presented components as boxes in a palette that can be 929 moved into an arena (workspace) and connected by wires. It also allows 930 component configurations and settings to be saved in BLD files and instantly 931 reloaded later. Another key feature of this GUI is that, as a lightweight and 932 platform-independent tool written in Java, it can be installed and used on 933 any computer with Java support to create a Ccaffeine script. This script can 934 then be sent to a remote, possibly high-performance computer for execution. 935 While the Ccafe-GUI was certainly useful, the CSDMS project realized

⁹³⁶ While the Ccafe-GUI was certainly useful, the CSDMS project realized ⁹³⁷ that it could be improved and extended in numerous ways to make it more ⁹³⁸ powerful and more user-friendly. In addition, these changes would serve not ⁹³⁹ only the CSDMS community but could be shared back with the CCA com-⁹⁴⁰ munity. That is, the new GUI works with any CCA-compliant components, ⁹⁴¹ not just CSDMS components. The new version is called CMT (CSDMS ⁹⁴² Modeling Tool). Significant new features of CMT 1.5 include the following.

943

• Integration with a powerful visualization tool called VisIt (see below).

944

• New, "wireless" paradigm for connecting components (see below).

945	• A login dialog that prompts users for remote server login information.
946 947	• Job management tools that are able to submit jobs to processors of a cluster.
948 949	• "Launch and go": launch a model run on a remote server and then shut down the GUI (the model continues running remotely).
950	• New File menu entry: "Import Example Configuration."
951	• A Help menu with numerous help documents and links to websites.
952	• Ability to submit bug reports to CSDMS.
953	• Ability to do file transfers to and from a remote server.
954 955	• Help button in tabbed dialogs to launch component-specific HTML help.
956	• Support for droplists and mouse-over help in tabled dialogs.
957 958	• Support for custom project lists (e.g., projects not yet ready for re- lease).
959	• A separate "driver palette" above the component palette.
960	• Support for numerous user preferences, many relating to appearance.
961	• Extensive cross-platform testing and "bulletproofing."
962	The CMT provides integrated visualization by using VisIt. VisIt $\left[47\right]$ is an
963	open-source, interactive, parallel visualization and graphical analysis tool for

viewing scientific data. It was developed by the U.S. Department of Energy 964 Advanced Simulation and Computing Initiative to visualize and analyze the 965 results of simulations ranging from kilobytes to terabytes. VisIt was designed 966 so that users can install a client version on their PC that works together with 967 a server version installed on a high-performance computer or cluster. The 968 server version uses multiple processors to speed rendering of large data sets 969 and then sends graphical output back to the client version. VisIt supports 970 about five dozen file formats and provides a rich set of visualization features, 971 including the ability to make movies from time-varying databases. The CMT 972 provides help on using VisIt in its Help menu. CSDMS uses a service com-973 ponent to provide other components with the ability to write their output 974 to NetCDF files that can be visualized with VisIt. Output can be 0D, 1D, 975 2D, or 3D data evolving in time, such as a time series (e.g., a hydrograph), 976 a profile series (e.g., a soil moisture profile), a 2D grid stack (e.g., water 977 depth), a 3D *cube stack*, or a scatter plot of XYZ triples. 978

Another innovative feature of CMT 1.5 is that it allows users to toggle 970 between the original, wired mode and a new wireless mode. CSDMS found 980 that displaying connections between components with the use of wires (i.e., 981 red lines) did not scale well to configurations that contained several compo-982 nents with multiple ports. In wireless mode, a component that is dragged 983 from the palette to the arena appears to broadcast what it can provide (i.e., 984 CCA provides ports) to the other components in the arena (using a con-985 centric circle animation). Any components in the arena that need to use 986 that kind of port get automatically linked to the new one; this is indicated 987 through the use of unique, matching colors. In cases where two components 988

⁹⁸⁹ in the arena have the same *uses port* but need to be connected to different
⁹⁹⁰ providers, wires can still be used.

CSDMS continues to make usability improvements to the CMT and used 991 the tool to teach a graduate-level course on surface process modeling at the 992 University of Colorado, Boulder, in 2010. Several features of the CMT make 993 it ideal for teaching, including (1) the ability to save prebuilt component 994 configurations and their settings in BLD files, (2) the File >> Import Ex-995 ample Configuration feature, (3) a standardized HTML help page for each 996 component, (4) a uniform, tabbed-dialog GUI for each component, (5) rapid 997 comparison of different approaches by swapping one component for another, 998 (6) the simple installation procedure, and (7) the ability to use remote re-990 sources. 1000



Figure 3: CMT screenshot.

¹⁰⁰¹ 8. Providing Components with a Uniform Help System and GUI

Beyond the usual software engineering definition of a component, a useful component will be one that also comes bundled with metadata that describes the component and the underlying model that it is built around. While creating a component as described in the preceding sections is important, it is of equal importance to have a well-documented component that an end user is able to easily use.

With a plug-and-play framework where users easily connect, interchange, 1008 and run coupled models, there is a tendency for a user to treat components 1009 as black boxes and ignore the details of the foundation that each component 1010 was built upon. For instance, if a user is unaware of the assumptions that 1011 underlie a model, that user may couple two components for which coupling 1012 does not make sense because of the physics of each model. The user may 1013 attempt to use a component in a situation where it was not intended to 1014 be used. To combat this problem, components are bundled with HTML 1015 help documents, which are easily accessible through the CMT, and describe 1016 the component and the model that it wraps. These documents include the 1017 following. 1018

- Extended model description (along with references)
- Listing and brief description of the component's uses and provides ports
- Main equations of the model
- Sample input and output
- Acknowledgment of the model developer(s)

A complete component also comes with metadata supplied in a more structured format. Components include XML description files that describe their user-editable input variables. These description files contain a series of XML elements that contain detailed information about each variable including a default value, range of acceptable values, short and long descriptions, units, and data type.

- 1030 <entry name=velocity>
- 1031 <label>River velocity</label>
- 1032 <help>Depth-averaged velocity at the river mouth</help>
- 1033 <default>2</default>
- 1034 <type>Float</type>
- 1035 <range>
- 1036 <min>0</min>

```
1037 <max>5</max>
```

1038 </range>

1039 <units>m/s</units>

1040 </entry>

Using this XML description, the CMT automatically generates a graphi-1041 cal user interface (in the form of tabbed dialogs) for each CSDMS component. 1042 Despite each model's input files being significantly different, this provides 1043 CMT users with a uniform interface across all components. Furthermore, the 1044 GUI checks user input for errors and provides easily accessible help within 1045 the same environment—none of which is available in the batch interface of 1046 most models. A special type of CCA provides port called a parameter port 1047 is also used in the creation of the tabbed dialogs. 1048

Nearly every model gathers initial settings from an input file and then 1049 runs without user intervention. Ultimately, any user interface that wraps a 1050 model must generate this input file for the component to read as part of its 1051 initialization step. The above XML description along with a template input 1052 file allows this to happen. Once input is gathered from the user, a model-1053 specific input file is created based on a template input file provided with each 1054 component. A valid input file is created based on \$-based substitutions in this 1055 template file. Instead of actual values, the template file contains substitution 1056 placeholders of the form **\$identifier**. Each identifier corresponds to an 1057 entry name in the XML description file and, upon substitution, is replaced 1058 by the value gathered from an external user interface (the CMT GUI, for 1059 instance). 1060

1061 9. Framework Services: "Built-in" Tools That Any Component 1062 Can Use

Developers (e.g., CSDMS staff) may wish to make certain low-level tools 1063 or utilities available so that any component (or component developer) can use 1064 them without requiring any action from a user. These tools can be encapsu-1065 lated in special components called *service components* that are automatically 1066 instantiated by a CCA framework on startup. The services or methods pro-1067 vided by these components are then called *framework services*. Unlike other 1068 components, which users may assemble graphically into larger applications, 1069 users do not interact with service components directly. However, a compo-1070 nent developer can make calls to the methods of service components through 1071 service ports. The use of service components allows developers to maintain 1072

¹⁰⁷³ code for a shared functionality in a single place and to make that function-¹⁰⁷⁴ ality available to all components regardless of the language they are written ¹⁰⁷⁵ in (or which address space they are in). CSDMS uses service components for ¹⁰⁷⁶ tasks such as (1) providing component output variables in a form needed by ¹⁰⁷⁷ another component (e.g., spatial regridding, interpolation in time, and unit ¹⁰⁷⁸ conversion) and (2) writing component output to a standard format such as ¹⁰⁷⁹ NetCDF.

Any CCA component can be "promoted" to a service component. A de-1080 veloper simply needs to add lines to its setServices() method that register it as 1081 a framework service. CCA provides a special port for this, gov.cca.ports.Ser-1082 viceRegistry, with three methods: addService(), addSingletonService(), and 1083 removeService(). If a developer then wants another component to be able to 1084 use this framework service, a call to the gov.cca.Services.getPort() method 1085 must be added within its setServices() method. (A similar call must be added 1086 in order to use CCA parameter ports and ports provided by other types of 1087 components.) Note that the setServices() method is defined as part of the 1088 gov.cca.Component interface. 1089

CCA components are designed for use within a CCA-compliant frame-1090 work (like Ccaffeine) and may make use of service components. But what if 1091 we want to use these components outside of a CCA framework? One option 1092 is to encapsulate a set of functionality (e.g., a service component) in a SIDL 1093 class and then "promote" this class to (SIDL) component status through in-1094 heritance and by adding only framework-specific methods like setServices(). 1095 (Note that a CCA framework is the entity that calls a component's setSer-1096 vices() method as described in Section 2.3.) This approach can be used to 1097

provide both component and noncomponent versions of the class. Compiling
the noncomponent version in a Bocca project generates a library file that we
can link against or, in the case of Python, a module that we can import.

1101 10. Current Contents of the CSDMS Component Repository

At the time of this publication the CSDMS model repository contains 1102 more than 160 models and tools. Of those, 50 have been converted into 1103 components as described in this paper and can be used in coupled modeling 1104 scenarios with the CMT or through the component composition interfaces 1105 supported by Ccaffeine. An up-to-date list is maintained at the CSDMS we-1106 biste. As with the model repository as a whole, CSDMS components cover 1107 the breadth of surface dynamics systems. Hydrologic components cover vari-1108 ous scales ranging from basin-scale (the entire TopoFlow [39] suite of models 1109 consists of 15 components that cover infiltration, meteorology, and channel 1110 dynamics; HydroTrend [4, 23]) to reach-scale (the one-dimensional sediment 1111 transport models of Parker [38]). Terrestrial components include models of 1112 landscape evolution (Erode, and CHILD [45]), geodynamics (Subside [21]) 1113 and cryospherics (GC2D [22]). Coastal and marine models include Ashton-1114 Murray Coastal Evolution Model [4, 5], Avulsion [4], and the stratigraphic 1115 model sedflux [21]. The component repository also contains modeling tools 1116 such as the ESMF and OpenMI SDK grid mappers, and file readers and 1117 writers for standard file formats (NetCDF, VTK, for example). 1118

1119 11. Conclusions

CSDMS uses a component-based approach to integrated modeling and 1120 draws on the combined power of many different open-source tools such as 1121 Babel, Bocca, Ccaffeine, the ESMF regridding tool, and the VisIt visualiza-1122 tion tool. CSDMS also draws on the combined knowledge and creative effort 1123 of a large community of Earth-surface dynamics modelers and computer sci-1124 entists. Using a variety of tools, standards, and protocols, CSDMS converts 1125 a heterogeneous set of open-source, user-contributed models into a suite of 1126 plug-and-play modeling components that can be reused in many different 1127 contexts. Components that encapsulate a physical process usually repre-1128 sent an optimal level of granularity. Standards that CSDMS has adopted 1129 and promotes include CCA, NetCDF [34], HTML, OGC (Open Geospatial 1130 Consortium) [37], MPI (Message Passing Interface) [32] and XML [48]. 1131

All the software that underlies CSDMS is installed and maintained on its 1132 high-performance cluster. CSDMS members have accounts on this cluster 1133 and access its resources using a lightweight, Java-based client application 1134 called the CSDMS Modeling Tool (CMT) that runs on virtually any desktop 1135 or laptop computer. This approach can be thought of as a type of *community* 1136 *cloud* since it provides remote access to numerous resources. This centralized 1137 cloud approach offers many advantages including (1) simplified maintenance, 1138 (2) more reliable performance, (3) automated backups, (4) remote storage 1139 and computation (user's PC remains free), (5) ability for many components 1140 (such as ROMS) and tools (such as VisIt and ESMF's regridder) to use 1141 parallel computation, (6) requiring to install only a lightweight client on their 1142 PC, (7) little technical support needed by users, and (8) ability to submit 1143

and run multiple jobs.

Babel's support of the Python language has proven very useful. Python 1145 is a modern, open-source, object-oriented language with source code that 1146 is easy to write, read and maintain. It runs on virtually any platform. It 1147 is useful for system administration, model integration, rapid prototyping, 1148 high-level tool development, visualization (via the matplotlib package) and 1149 numerical modeling (via the numpy package). Bocca is written in Python, the 1150 VisIt visualization package has a powerful Python API, and ESRI's ArcGIS 1151 software now uses Python as its scripting language ([10]). Many third-party 1152 geographic information system (GIS) tools implemented in Python are also 1153 available. With the numpy, scipy, and matplotlib packages, Python provides 1154 a work-alike to commercial languages like Matlab with similar performance. 1155 Other Python packages that CSDMS has found useful are suds (for SOAP-1156 based web services) and PyNIO (an API for working with NetCDF files). 1157

Several exciting opportunities exist for further streamlining and expand-1158 ing the capabilities of CSDMS. One area of particular interest is how CS-1150 DMS can provide its members with multiple paths to parallel computation. 1160 Software may be designed from the outset to use multiple processors, or be 1161 refactored to do so, often using MPI or OpenMP. But this is not easy and 1162 typically requires a multiyear investment. Another way to harness the power 1163 of parallelism is to modify code to take advantage of numerical toolkits such 1164 as PETSc (Portable Extensible Toolkit for Scientific Computation) [6, 7, 8] 1165 that contain parallel solvers for many of the differential equations that are 1166 used in physically based models. A third way is to for models written in 1167 array-based languages such as IDL, Matlab [31] and Python/NumPy [42] to 1168

¹¹⁶⁹ use array-based functions and operators that have been parallelized. This ¹¹⁷⁰ approach, although available only in commercial packages at present, is at-¹¹⁷¹ tractive for several reasons: (1) developers in these languages already know ¹¹⁷² to avoid spatial loops and use the array-based functions whenever possible ¹¹⁷³ for good performance, (2) most of these array-based functions are straightfor-¹¹⁷⁴ ward to parallelize, and (3) developers need only import a different package ¹¹⁷⁵ to take advantage of the parallelized functions.

Web services provide many additional opportunities. Peckham and Goodall [40] have demonstrated how CSDMS components can use CUAHSI-HIS [13] web services to retrieve hydrologic data, but CSDMS components could also offer their capabilities as web services.

CSDMS is also interested in *automated component wrapping*, which can be achieved by adding special annotation keywords within comments in the source code. If the code is sufficiently annotated, it is possible to write a flexible tool to wrap the component with any desired interface. Unfortunately, most existing code has not been annotated in this way, and it is typically necessary to involve the code's developer in the annotation process.

1186 Acknowledgments

CSDMS gratefully acknowledges major funding through a cooperative agreement with the National Science Foundation (EAR 0621695). Additional work was supported by the Office of Advanced Scientific Computing Research, Office of Science, U.S. Dept. of Energy, under Contracts DE-AC02-06CH11357 and DE-FC-0206-ER-25774.

1192 References

- [1] Allan, B., Armstrong, R., Lefantzi, S., Ray, J., Walsh, E., Wolfe, P.,
 2010. Ccaffeine a CCA component framework for parallel computing.
 http://www.cca-forum.org/ccafe/.
- [2] Allan, B. A., Norris, B., Elwasif, W. R., Armstrong, R. C., Dec. 2008.
 Managing scientific software complexity with Bocca and CCA. Scientific
 Programming 16 (4), 315–327.
- [3] Armstrong, R., Gannon, D., Geist, A., Keahey, K., Kohn, S., McInnes,
 L., Parker, S., Smolinski, B., 1999. Toward a Common Component Architecture for high-performance scientific computing. In: Proc. 8th IEEE
 Int. Symp. on High Performance Distributed Computing.
- [4] Ashton, A., Kettner, A. J., Hutton, E. W. H., 2011. Progress in coupling
 between coastline and fluvial dynamics. Computers & Geosciences (this
 issue).
- [5] Ashton, A., Murray, A. B., Arnoult, O., 2001. Formation of coastline
 features by large-scale instabilities induced by high-angle waves. Nature
 414, 296–300.
- [6] Balay, S., Brown, J., Buschelman, K., Eijkhout, V., Gropp, W. D.,
 Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., Zhang,
 H., 2010. PETSc users manual. Tech. Rep. ANL-95/11 Revision 3.1,
 Argonne National Laboratory.

- [7] Balay, S., Brown, J., Buschelman, K., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., Zhang, H., 2011. PETSc web
 page. Http://www.mcs.anl.gov/petsc.
- [8] Balay, S., Gropp, W. D., McInnes, L. C., Smith, B. F., 1997. Efficient management of parallelism in object oriented numerical software
 libraries. In: Arge, E., Bruaset, A. M., Langtangen, H. P. (Eds.), Modern
 Software Tools in Scientific Computing. Birkhäuser Press, pp. 163–202.
- 1220 [9] Bernholdt D. (PI), 2010. TASCS Center.
 1221 http://www.scidac.gov/compsci/TASCS.html.
- Н., 2005.[10] Buttler, AprilJune Α guide to the python uni-1222 for ArcUser Mag.Available online verse esriusers. at1223 http://www.esri.com/news/arcuser/. 1224
- [11] CCA Forum, 2010. A hands-on guide to the Common Component Ar chitecture. http://www.cca-forum.org/tutorials/.
- [12] CSDMS, 2011. Community Surface Dynamics Modeling System (CS DMS). http://csdms.colorado.edu.
- [13] CUAHSI, 2011. Consortium of Universities for the Advancement of the
 Hydrological Sciences Inc. . http://www.cuahsi.org.
- [14] Dahlgren, T., Epperly, T., Kumfert, G., Leek, J., 2007. Babel User's
 Guide. CASC, Lawrence Livermore National Laboratory, UCRL-SM230026, Livermore, CA.

- [15] de St. Germain, J. D., Morris, A., Parker, S. G., Malony, A. D., Shende,
 S., May 15-17 2002. Integrating performance analysis in the Uintah
 software development cycle. In: Proceedings of the 4th International
 Symposium on High Performance Computing (ISHPC-IV). pp. 190–206.
- URL http://www.sci.utah.edu/publications/dav00/ishpc2002.pdf
- [16] Diachin L. (PI), 2011. Center for Interoperable Tech-1239 nologies for Advanced Petascale Simulations (ITAPS). 1240 http://www.scidac.gov/math/ITAPS.html. 1241
- 1242 [17] EJB, 2011. Enterprise Java Beans Specification.
 1243 http://java.sun.com/products/ejb/docs.html.
- [18] ESMF Joint Specification Team, 2011. Earth System Modeling Frame work (ESMF) Website. http://www.earthsystemmodeling.org/.
- [19] FRAMES, 2011. Framework for Risk Analysis of Multi-Media Environ mental Systems (FRAMES). http://mepas.pnl.gov/FRAMESV1/.
- [20] Hill, C., DeLuca, C., Balaji, V., Suarez, M., da Silva, A., ESMF Joint
 Specification Team, 2004. The architecture of the Earth System Modeling Framework. Computing in Science and Engineering 6, 18–28.
- [21] Hutton, E. W. H., Syvitski, J. P. M., 2008. Sedflux-2.0: An advanced
 process-response model that generates three-dimensional stratigraphy.
 Computers & Geosciences 34 (10), 1319–1337.
- [22] Kessler, M. A., Anderson, R. S., Briner, J. P., 2008. Fjord insertion
 into continental margins driven by topographic steering of ice. Nature
 Geoscience 1, 365–369.

- [23] Kettner, A. J., Syvitski, J. P. M., 2008. Hydrotrend version 3.0: a
 climate-driven hydrological transport model that simulates discharge
 and sediment load leaving a river system. Computers & Geosciences
 34 (10), 1170–1183.
- [24] Keyes D. (PI), 2011. Towards Optimal Petascale Simulations (TOPS)
 Center. http://tops-scidac.org/.
- [25] Krishnan, S., Gannon, D., April 2004. XCAT3: A framework for CCA
 components as OGSA services. In: Proceedings of the 9th International
 Workshop on High-Level Parallel Programming Models and Supportive
 Environments (HIPS 2004). IEEE Computer Society, pp. 90–97.
- [26] Kumfert, G., April 2003. Understanding the CCA Specification Using
 Decaf. Lawrence Livermore National Laboratory.

1269 URL http://www.llnl.gov/CASC/components/docs/decaf.pdf

- [27] Larson, J. W., 2009. Ten organising principles for coupling in multiphysics and multiscale models. ANZIAM Journal 47, C1090–C1111.
- [28] Larson, J. W., Norris, B., 2007. Component specification for parallel
 coupling infrastructure. In: Gervasi, O., Gavrilova, M. L. (Eds.), Proceedings of the International Conference on Computational Science and
 its Applications (ICCSA 2007). Vol. 4707 of Lecture Notes in Computer
 Science. Springer-Verlag, pp. 56–68.
- 1277 [29] Lawrence Livermore National Laboratory, 2011. Babel.
 1278 http://www.llnl.gov/CASC/components/babel.html.

- [30] Lucas R. (PI), 2011. Performance Engineering Research Institute
 (PERI). http://www.peri-scidac.org.
- [31] MathWorks, 2011. MATLAB The Language of Technical Computing.
 http://www.mathworks.com/products/matlab/.
- [32] Message Passing Interface Forum, 1998. MPI2: A message passing in terface standard. High Performance Computing Applications 12, 1–299.
- 1285 [33] NET, 2011. Microsoft .NET Framework. 1286 http://www.microsoft.com/net/.
- ¹²⁸⁷ [34] NetCDF, 2011. NetCDF. http://www.unidata.ucar.edu/packages/netcdf.
- 1288 [35] OMP, 2011. Object Modeling System v3.0.
 1289 http://www.javaforge.com/project/oms.
- [36] Ong, E. T., Larson, J. W., Norris, B., Jacob, R. L., Tobis, M., Steder,
 M., 2008. A multilingual programming model for coupled systems. In ternational Journal for Multiscale Computational Engineering 6, 39–51.
- [37] Open Geospatial Consortium, 2011. OGC Standards and Specifications.
 http://www.opengeospatial.org/.
- [38] Parker, G., 2011. 1d sediment transport morphodynam-1295 ics with applications to rivers and turbidity currents. 1296 http://vtchl.uiuc.edu/people/parkerg/morphodynamic_e-book.htm. 1297
- [39] Peckham, S., 2008. Geomorphometry and spatial hydrologic modeling.
 Vol. 33 of Geomorphometry: Concepts, Software and Applications. Developments in Soil Science. Elsevier, Ch. 25, pp. 579–602.

- [40] Peckham, S. D., Goodall, J. L., 2011. Driving plug-and-play components with data from web services: A demonstration of interoperability
 between CSDMS and CUAHSI-HIS. Computers & Geosciences (this issue).
- [41] Rosen, L., 2004. Open Source Licensing: Software Freedom and Intellec tual Property Law. Prentice Hall, http://rosenlaw.com/oslbook.htm.
- Image: [42] T. Oliphant et al., 2011. Scientific Computing Tools for Python –
 NumPy. http://numpy.scipy.org/.
- [43] The MCT Development Team, 2006. Model Coupling Toolkit (MCT)
 Web Site. http://www.mcs.anl.gov/mct/.
- 1311 [44] The OpenMI Association, 2011. The Open Modeling Interface
 1312 (OpenMI). http://www.openmi.org.
- [45] Tucker, G. E., Lancaster, S. T., Gasparini, N. M., Bras, R. L., 2001. The
 Channel-Hillslope Integrated Landscape Development (CHILD) Model.
 Academic/Plenum Publishers, pp. 349–388.
- [46] United States Department of Energy, 2011. SciDAC Initiative homepage.
 http://scidac.gov/.
- ¹³¹⁸ [47] VisIt, 2011. VisIt. http://wci.llnl.gov/codes/visit.
- 1319 [48] XML, 2011. Extensible Markup Language (XML).
 1320 http://www.w3.org/XML/.

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.

1321