

A Generic Component Model for Building Systems Software

GEOFF COULSON, GORDON BLAIR, PAUL GRACE, FRANCOIS TAIANI,
ACKBAR JOOLIA, KEVIN LEE, JO UHEYAMA, and
THIRUNAVUKKARASU SIVAHARAN
Lancaster University

Component-based software structuring principles are now commonplace at the application level; but componentization is far less established when it comes to building low-level systems software. Although there have been pioneering efforts in applying componentization to systems-building, these efforts have tended to target specific application domains (e.g., embedded systems, operating systems, communications systems, programmable networking environments, or middleware platforms). They also tend to be targeted at specific deployment environments (e.g., standard personal computer (PC) environments, network processors, or microcontrollers). The disadvantage of this narrow targeting is that it fails to maximize the genericity and abstraction potential of the component approach. In this article, we argue for the benefits and feasibility of a generic yet tailorable approach to component-based systems-building that offers a uniform programming model that is applicable in a wide range of systems-oriented target domains and deployment environments. The component model, called *OpenCom*, is supported by a reflective runtime architecture that is itself built from components. After describing OpenCom and evaluating its performance and overhead characteristics, we present and evaluate two case studies of systems we have built using OpenCom technology, thus illustrating its benefits and its general applicability.

Categories and Subject Descriptors: D.4.7 [Operating Systems]: Organization and Design—*Distributed systems; real-time systems and embedded systems*; D.2.6 [Software Engineering]: Programming Environments—*Programmer workbench*

General Terms: Design, Standardization, Management

Additional Key Words and Phrases: Component-based software, computer systems implementation

ACM Reference Format:

Coulson, G., Blair, G. Grace, P., Taiani, F., Joolia, A., Lee, K., Ueyama, J., and Sivaharan, T. 2008. A generic component model for building systems software. *ACM Trans. Comput. Syst.* 26, 1, Article 1 (February 2008), 42 pages. DOI=10.1145/1328671.1328672 <http://doi.acm.org/10.1145/1328671.1328672>

This work is supported by the U.K. EPSRC.

Authors' address: Computing Department, Lancaster University, Lancaster LA1 4WA, England; email: G. Coulson and P. Grace: {g.coulson,p.grace}@lancaster.ac.uk; G. Blair, F. Taiani, A. Joolia, K. Lee, and T. Sivaharan: {gordon,f.taiani,joolia,leek,t.sivaharan}@comp.lancs.ac.uk; J. Ueyama: joueyama@ic.unicamp.br.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org. © 2008 ACM 0734-2071/2008/02-ART1 \$5.00 DOI 10.1145/1328671.1328672 <http://doi.acm.org/10.1145/1328671.1328672>

1. INTRODUCTION

Application-level component-based software is now well established. Prominent examples include browser plug-ins [Mozilla Organization 2005], JavaBeans and Enterprise JavaBeans [Sun Microsystems 2005], the CORBA Component Model [Object Management Group 1999], Microsoft's .NET [Microsoft 2005], and ICENI Grid components [Furmento et al. 2002]. Overall, some key characteristics of the component approach are as follows: (i) it promotes a high degree of genericity and abstraction in software design, implementation, and deployment, which leads to higher programmer productivity; (ii) it facilitates flexible configuration of software (and, potentially, runtime reconfiguration), and (iii) it fosters third-party software reuse [Emmerich 2002].

A broadly accepted definition of software components is that they are “units of composition with contractually-specified interfaces and explicit context dependencies only . . . that can be deployed independently and are subject to composition by third parties” [Szyperski 1998, p. 34]. Some component models restrict composition to build time or load time, but many (including all of the above) additionally support the composition of components at *runtime*. The “explicit context dependencies” aspect of the definition means that third parties can straightforwardly deploy independently developed components into established runtime software environments. Deployment is straightforward because it is clear and explicit what a newly deployed component expects from its environment.

Although component models are widely used at the application level, it is far less common to see the component approach being exploited in the construction of *low-level systems software* such as embedded systems, operating systems, communications systems, programmable networking environments, or middleware platforms. This is primarily because systems environments are typically far more demanding than application environments in terms of complexity, performance, and resource constraints. Nevertheless, the potential of the component approach appears in principle just as compelling in the systems area as it is in the applications area, and this view has been borne out by a number of pioneering efforts over the last few years. For example, proposals for component platforms for building *embedded systems* include Pebble [Magoutis et al. 2000], PECOS [Winter et al. 2002], and Koala [Van Ommering et al. 2000]; proposals for componentized *operating systems* (OSs) include THINK [Fassino et al. 2002], OSKit [Ford et al. 1997], and MMLITE [Helander and Forin 1998]; proposals for componentized *programmable networking environments* include VERA [Karlín and Peterson 2001], MicroACE [Johnson and Kunze 2003], and Netbind [Campbell et al. 2002]; and proposals for componentized *middleware platforms* include LegORB [Roman et al. 2000], k-Components [Dowling and Cahill 2001], and various JavaBeans-based approaches (e.g., Bruneton and Riveill [2000] and Joergensen et al. [2000]).

However, all of these efforts suffer from the key limitation that they are *narrowly targeted*. This applies in two senses: (i) in terms of the *target domain* at which they are aimed (i.e., embedded systems, OSs, etc.), and (ii) in terms of the intended *deployment environment* in which they will operate (e.g., most

of the above-mentioned technologies have been deployed only on conventional desktop machines as opposed to more “exotic” deployment environments like personal digital assistants (PDAs), embedded hardware, or network processors). The disadvantage of this narrow targeting is that it fails to maximize the genericity and abstraction potential of the component approach—for example it locks systems component programmers into narrow, nontransferable skill sets and areas of expertise. Also, it fails to support the construction of component-based systems that span target domains and/or deployment environments (e.g., embedded middleware, or OSs for network processors).

In this article, we discuss a *general-purpose* component-based systems-building technology called *OpenCom*. OpenCom tries to maximize the genericity and abstraction potential of the component based programming model *while at the same time* supporting a principled approach to supporting the unique requirements of a wide range of target domains and deployment environments. This is achieved by splitting the programming model into a simple, efficient, minimal kernel, and then providing on top of this a principled set of extension mechanisms that allow the necessary tailoring. We also recognize and support a *separation of roles* between programmers who use the extension mechanisms to realize an OpenCom-based platform in a given deployment environment, and programmers who then use this environment to develop a target system (e.g., an embedded system, OS, etc.).

In more detail, the design of OpenCom tries to address the following requirements:

- (1) *Target domain independence*. A general-purpose systems-building technology should provide only generic and fundamental functionality that is independent of the specialist needs of any particular target domain. For example, a generic technology should *not* inherently support characteristics such as real-time execution, sand-boxing, or 24×7 availability. This is because such characteristics carry an inevitable cost which should not be incurred where they are not required. Nevertheless, a general-purpose technology should be *inherently tailorable* and *extensible* so that it can be specialized in a natural and explicitly supported way to meet such needs where required. The same consideration applies to runtime reconfigurability: the technology should provide a basis for this but should not *dictate the policies* that control and manage it.
- (2) *Deployment environment independence*. The technology should be straightforwardly deployable in a wide range of deployment environments from PCs, to supercomputers, to set-top boxes, to resource-poor PDAs, to bare-iron embedded systems with no OS support, to networks-on-a-chip. This implies *simplicity* (for ease of porting), *small memory footprint*, and *programming language independence*. More fundamentally, it again implies inherent support for *tailorability* and *extensibility*. A key aspect of extensibility at the deployment-environment level is that it should be straightforward to expose idiosyncratic hardware and software features of the underlying deployment environment (e.g., multiple processors, hardware hashing units, optimized




Fig. 1. Overall OpenCom architecture.

interprocess communication (IPC) channels, memory hierarchies, etc.) in terms of the native abstractions of the generic component-based programming model.

- (3) *Negligible overhead.* As well as incurring only a small memory overhead, the technology should impose as small a demand as possible on other resources—especially processing resources. This particularly implies that the “in-band” execution path [Coulson et al. 2004] of target systems should be as independent as possible of any runtime support provided by the technology (e.g., intercomponent communication should not be reliant on a kernel-mediated message passing service). In addition, exposing deployment-environment-specific features in terms of generic programming model abstractions (as discussed above) should incur as small a performance penalty as possible—ideally a penalty of zero.

The remainder of this article is structured as follows: first, Section 2 provides an overview and motivation of the general OpenCom approach. Then Section 3 presents in detail our minimal component-based programming model. Next, Sections 4 and 5 discuss the key “extension mechanisms” that are (optionally) layered on top of the OpenCom kernel and are instrumental in providing the necessary tailorability and extensibility. Section 6 then analyzes the inherent performance characteristics and overheads of the OpenCom approach, and Section 7 presents case studies of our recent use of OpenCom, including its extension layers, in building nontrivial systems. Finally, related work is discussed in Section 8, and our conclusions are offered in Section 9.

2. OVERALL APPROACH

We approach the satisfaction of the requirements identified above—that is, target domain independence, deployment-environment independence, and negligible overhead—through the architecture illustrated in Figure 1. At the heart of the architecture is a minimal *component runtime kernel* that supports the basic services of *loading* and *binding* components. This is discussed in detail in Section 3. A runtime kernel is required to be able to support dynamic systems which have an inherent need for runtime reconfigurability (e.g., extensible OSs, active networking nodes, adaptive middleware, etc. [Blair et al. 2004]). The kernel lies immediately above the (hardware and/or software) *deployment environment*. The kernel is policy free, and its application programmer’s interface

(API) is target-system and deployment-environment independent. For static systems, it is used only to initially configure the system—when configuration is complete, it can be unloaded so that it does not consume any resources. In dynamic systems, the kernel continues to exist at runtime. However, even here its resource demands are minimal, as shown in Section 6.

Above the kernel is a layer of so-called *extensions* which enhance the basic loading and binding-based programming model in accordance with the needs of various target domains and deployment environments. This layer thus plays a central role in providing the tailorability and extensibility that OpenCom aims to deliver. The extensions are independently and optionally deployable and configurable (via the kernel). Importantly, the extensions are themselves implemented as components, so there is no essential boundary between the extensions and target system “layers” and thus no inherent layering overhead. The extensions that we currently employ fall into the two main classes: first, *platform extensions*, discussed in Section 4, provide structured support for tailorability and extensibility at the deployment environment level—essentially, this layer addresses the above-mentioned requirement to efficiently expose unique features of deployment environments in terms of generic component-based abstractions. Second, *reflective extensions*, discussed in Section 5, provide generic support for target system reconfiguration—that is, inspecting, adapting, and extending the structure and behavior of dynamic systems at runtime [Maes 1987]. These reflective extensions build on and extend inherently reflective features of the kernel such as explicitly represented cross-component bindings and support for extensible metadata (see Section 3). We also provide a set of *security extensions*; these, however, are not as mature as the other extensions and are not discussed further in this article.

A key architectural feature of the OpenCom approach is its extensive use of the notion of *component frameworks* [Szyperski 1998]. Component frameworks work at a coarser granularity than components, and contribute a generic approach to the structuring and extensibility of software through component composition. In OpenCom, a component framework (CF) is a tightly coupled set of components that (i) cooperates to address some focused area of concern; (ii) provides a well-defined extension protocol that accepts additional “plug-in” components that modify or extend the CF’s behavior; and (iii) constrains [Clarke et al. 2001] how these plug-ins may be organized. As an example, we have a protocol stacking CF that accepts protocol components as its plug-ins, and constrains its plug-ins to be composed into linear stacks [Coulson et al. 2002b].

OpenCom CFs typically employ *runtime pluggability* as well as merely the design-time or build-time pluggability that is found in many component models. For example, a CF may be represented at runtime as a “root” component that exports an operation that accepts plug-in components as its arguments. Internally, this root component discovers the interfaces supported by plug-ins (using reflection as explained later), and configures and binds them in a manner appropriate to the CF (e.g., in the above-mentioned protocol stacking CF, the root component would discover and bind the interfaces of its plug-ins to generate a linear stack topology). We have also explored a more sophisticated approach in which CF constraints are expressed at design time in terms of an

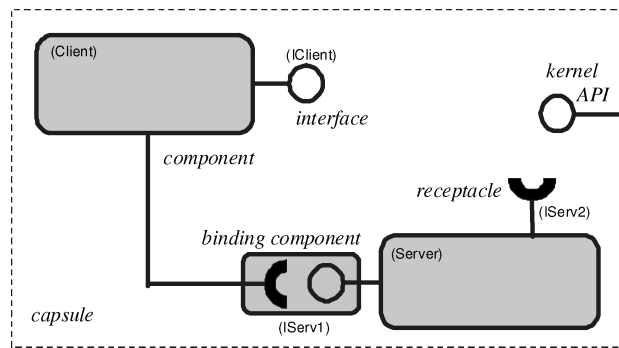


Fig. 2. Elements of the kernel-level programming model.

architecture description language (specifically, ACME [Garlan et al. 2000]) and are compiled to generate CF-specific constraint policing code [Joolia et al. 2005]. The key point, however, is that different CFs can adopt different approaches to pluggability and constraint. In fact, CFs do not inherently require anything beyond the facilities provided by the foundational component model; that is, minimal CFs can be viewed simply as architectural patterns.

As will become clear in the remainder of this article, CFs provide structure and extensibility at all levels of the architecture. At the level of the platform extensions, CFs are provided that, for example, support plug-in “loader” and “binder” components. Similarly, at the level of the reflective extensions, plug-ins take the form of, for example, operation interceptors. At the target-system level, plug-ins are applied in such areas as protocol stacking (as discussed above), thread scheduling, packet forwarding, memory management, or user interaction; and they define plug-ins and constraints that make sense in those domains.

3. THE PROGRAMMING MODEL AND THE KERNEL API

3.1 Programming Model

A high-level view of the elements of the OpenCom component-based programming model is given in Figure 2 (the legends in brackets refer to the example given in Section 3.3 below). *Capsules* are containers into which components are loaded, instantiated, and composed. Each capsule defines a namespace for its contained component instances (hereafter, we refer to component instances simply as *components*), and offers the OpenCom kernel API, which is discussed in Section 3.2 below. Capsules do not recognize any nesting or hierarchical organization of their contained components, although such organizations can be conceptually superimposed on this basic “flat” organization through the use of ADL-based CFs or other such *formalisms*.

Components are encapsulated units of functionality which interact with other components in their containing capsule exclusively through so-called *interaction points*, of which there are two types: *interfaces* and *receptacles*. Component types are templates from which components (instances) can be instantiated at runtime. Each component type is defined by a name, the set of the interaction points it supports, and a set of statically defined

`<name, type, value>` properties (this so-called *static property* facility is used to associate arbitrary metadata with a component type which is available at run-time using the `getprop()` kernel call; see below). *Interfaces* are units of service provision offered by components. Components types may support any number of interfaces including zero (in Figure 2, each of the large components has one interface and one receptacle; and the small component has none). The use of multiple interfaces is useful in embodying separations of concern (e.g., between base functionality and component management). Interfaces are defined in terms of sets of operation signatures and associated datatypes. *Receptacles* are “required interfaces” that make explicit the dependencies of a component on other components. Components may support any number of receptacles. Receptacles are key to supporting the *third-party* mode of deployment and composition inherent in a component-oriented environment: when third-party-deploying a component into a capsule, one knows by looking at its receptacles precisely which other components must be present to satisfy the component’s dependencies.

For language independence, we use the OMG’s IDL interface definition language [Object Management Group 1995] to define interfaces, receptacles, and component types. Interfaces and receptacles are expressed using the standard interface definition syntax.¹ Component types, on the other hand, employ the following extended syntax (which is similar to a subset of the OMG “component” syntax):

```

<comp_type_defn> ::= ‘componentType’ <comp_name> ‘{’
                    <static_props> <provides_and_uses> ‘}’
<static_props> ::= (<prop_name> ‘:’ <prop_type>
                    ‘=’ <prop_value> ‘;’)*
<provides_and_uses> ::= (<provides> | <uses>)*
<provides> ::= ‘provides’ <interface_name> ‘;’
<uses> ::= ‘uses’ <interface_name> ‘;’

```

The “provides” clause refers to interfaces supported by the component type and the “uses” clause refers to its receptacles. The elements `<comp_name>` and `<interface_name>` are strings; the latter refers to the names of IDL interfaces defined elsewhere. Similarly, `<prop_name>` is a string, and `<prop_type>` refers to an IDL data type defined elsewhere; `<prop_value>` is a value of the appropriate `<prop_type>`.

Note that we do not provide any facilities at the component-type definition level for “nesting” component definitions, or to specify static bindings between receptacles and interfaces. The aim is to maximize the simplicity of the programming model. Nevertheless, such higher-level facilities can be straightforwardly built on top where required. For example, we have employed the ACME ADL to specify such concerns, and have also experimented with an XML-based formalism [Joolia et al. 2005]. However, when such formalisms are employed, they always compile down to the basic OpenCom programming

¹Actually, receptacles do not need to be defined explicitly as each receptacle is implicitly defined in terms of its associated interface.

model—that is, a “flat” component structure in which all receptacle-interface bindings are created at runtime.

A *binding* is a runtime association between a single interface and a single receptacle. Interfaces may participate in multiple bindings, whereas each receptacle may only participate in a single binding at a time. Like component deployment, the creation of bindings is inherently *third party* in nature. That is, bindings can be created by any component within the capsule, not only by the “first-party” components whose interface or receptacle is actually participating in the binding. Each binding is represented by a component (the small component in Figure 2) that is implicitly created by the kernel. The semantics of these components (often called *binding components*) are identical to those of any other component—with the following exception: when such a component is destroyed, the interface/receptacle association that it represents is also destroyed. The OpenCom specification allows kernel implementations to employ binding components that themselves support any number of interfaces and receptacles. Binding components with zero interfaces/receptacles (such as the one shown in the figure) are sufficient for lightweight implementations in a resource-poor deployment environment; but less constrained implementations are free to support binding components that have interfaces that, for example, support operations to obtain the identifiers of the bound interface/receptacle, or to insert interceptors. Such facilities, however, are more typically supported by extension binders, as discussed in Section 4.2.

3.2 The Runtime Kernel API

Each capsule embodies a single kernel instance which offers the following runtime API²:

```
interface Kernel {
    typedef struct template {long id};
    typedef struct component {long id};
    typedef struct interface {long id};
    typedef struct receptacle {long id};

    status load(in string component_type_name, out template t);
    status instantiate(in template t, out component c);
    status unload(in template t);
    status destroy(in component c);
    status bind(in interface i, in receptacle r, out component binding);
    status putprop(in long entity_UID, in string key, in any value);
    status getprop(in long entity_UID, in string key, out any value);
    long register(in long proposed_UID);
    status notify(in ICallback callback);
}
```

This API is deliberately minimal in nature and has been specified on the basis of considerable experience and experimentation with runtime component model

²The API, specified here in OMG IDL, has been slightly simplified for presentational purposes.

APIs over the past few years [Clarke et al. 2001; Coulson et al. 2003; Grace et al. 2005]. The basic philosophy is to build the API in terms of two very primitive system-level facilities: *dynamic loading* (in the shape of *load()*, *instantiate()*, *unload()*, and *destroy()*); and *dynamic linking* (in the shape of *bind()*).³ This API then supports the implementation of all other areas of systems-related functionality (e.g., concurrency, protection, distribution, etc.) as well-defined components that build on the basic kernel-level loading and linking services. The fact that the “loading and binding layer” is offered as a well-defined component model which is consistently reused in the higher layers lends great coherence, uniformity, and flexibility to the approach, while at the same time allowing the kernel to be implemented in an extremely small, efficient, and policy-free manner.

We now discuss the API in detail. The struct definitions are used to name the four types of entities comprehended by the kernel (i.e., templates, components, interfaces, and receptacles). Wrapping the UIDs of these entities in structs facilitates language-independent type safety at a cost of only *sizeof(long)* bytes of memory overhead per UID. *Load()* loads a named component type into the capsule and returns a “template”—that is, an in-memory representation of a component type (including its executable code) which can subsequently be instantiated using *instantiate()*. Instantiation is separated from loading to assist the user in controlling the tradeoff between memory economy and instantiation latency. For example, a programmer may choose to load and instantiate on demand (high instantiation latency, but with the benefit of only incurring memory overhead when a given component is actually required); or alternatively she may choose to preload templates so that instances can later be created quickly (low instantiation latency at the expense of having the templates occupy memory between the load and instantiate steps). *Unload()* unloads the specified template (to free up memory), and *destroy()* destroys a component instance. *Unload()* fails if there are extant instances of the target template; and *destroy()* fails if any of the target component’s receptacles or interfaces are currently bound. As the kernel itself is modeled as a component which exports its API as an OpenCom interface, *destroy()* can be used to *remove the kernel itself*—this frees up the memory used by the kernel but leaves all components and bindings untouched; the effect is to forgo the possibility of making subsequent runtime changes in the capsule. *Bind()* is used to create a binding between a specified receptacle and interface. As bindings are represented by components, *destroy()* is used to remove a binding. The arguments to *bind()* (i.e., an interface and a receptacle of the to-be-bound components) can be obtained from an internal kernel “registry” which is accessed via the *putprop()* and *getprop()* calls. All components are required on instantiation to call *putprop()* to store in the registry their interfaces and receptacles, using built-in *key* arguments of “I” and “R,” respectively. These interfaces and receptacles can then be retrieved, given the

³Having said that, we do obviously provide some additional calls that have been found to be generally useful—that is, basic facilities to (i) manage metadata associated with component model elements (i.e., *putprop()* and *getprop()*), (ii) manage names (i.e., *register()*), and (iii) reflect on calls being made on the API (i.e., *notify()*).

target component’s identifier, using *getprop()*. Apart from these built-in keys, higher-level CFs and extensions can use the registry facility to attach arbitrary metadata (using keys that they themselves define) to any component model entity (i.e., templates, components, interfaces, or receptacles).

Finally, the purpose of *register()* and *notify()* is to provide specific support to the extensions, as discussed below in Sections 4 and 5. *Register()* allocates a UID for any newly created entity and stores it in the registry. It is used when entities are created by platform extensions rather than by the kernel itself. The role of *notify()* is to assist reflective extensions in obtaining and maintaining information about relevant activity in the capsule, and to serve as the basis of a *policy enforcement point* for security and consistency management purposes. When a callback is registered with *notify()*, every subsequent call on the kernel (i.e., of *load()*, *bind()*, etc.) is reported to the callback. More specifically, the callback is invoked *twice* for each kernel call: the first invocation is made *before* the associated kernel call has been made; it reports the “in” arguments of the kernel call (e.g., the *name* argument to *load()*). The second callback is made *after* the associated kernel call has been made; it reports the “out” argument values that have come back from the kernel call. It is possible to use this callback facility to “veto” a kernel call by returning a specific value from the first callback. This prevents the kernel call from being executed and as a consequence prevents the second callback from occurring; an error code is returned to the caller of the kernel call.

Despite its minimality and relative ease of realisation (see Section 6), the above-described API already provides a powerful and self-contained component-based programming model. It is nevertheless still quite limited in terms of its tailorability and extensibility. For example, it supports only a single (implicit) mechanism for loading and binding components, and it does not provide any specific support for principled reconfiguration beyond the basic capability to dynamically create and destroy components and bindings. The extension layers described in Sections 4 and 5 build on the basic kernel functionality to specifically address such concerns.

3.3 Programming Language Bindings

We have realized the above-described programming model and kernel in Java, C, and C++, and in a range of deployment environments (see the case studies in Section 7). All the language bindings employ an IDL compiler to generate glue code, and to define, according to standard OMG-defined programming language mappings, the language-specific representations of the four OpenCom entities (i.e., templates, components, interfaces, and receptacles), and common data types like ints, strings, etc.

Due to its simplicity and accessibility, it is most useful to use the Java programming language to exemplify language binding issues. In Java binding, components and receptacles are represented as Java classes, and component interfaces are represented as Java interfaces that the component class “implements.” The IDL compiler generates the receptacle classes and also a per-component class called `._component_name` from which the programmer’s

component implementation class should inherit. This generated class makes available the kernel API operations to its user-defined child class (this is done by transparently prebinding a per-component receptacle called a *kernel* to the kernel interface), encapsulates the declaration and initialization of the receptacle classes, and uses *putprop()* to register the component type's static properties with the kernel. The child class itself may also, of course, register any further "dynamic" properties with the kernel as it sees fit.

Given this preamble, it should be straightforward to understand the following simple example, which refers to the component topology shown in Figure 2. First, we define the interfaces used in the example:

```
interface IServ1 {int op1(int i, int j);}
interface IServ2 {int op2(int i);}
interface IClient {int setup(); int call(char op, int arg1,
    int arg2);}
```

Next, we define the two component types: the *Server* type provides an *IServ1* interface and has a receptacle for *IServ2*; and the *Client* type provides *IClient* and has a receptacle for *IServ1*:

```
componentType Server {version:int=1; provides IServ1; uses IServ2;}
componentType Client {provides IClient; uses IServ1;}
```

Next, we write Java application code corresponding to these specifications. This might appear as follows:

```
public class Server extends _Server implements IServ1 {
    public Server() { super (); }
    public int op1(int a, int b) {return ...;} /* defined in IServ1 */
}
public class Client extends _Client implements IClient {
    public Binding b;

    public Client() { super(); }
    public int setup() { /* defined in IClient */
        Template t_serv = kernel.load("Server");
        Component serv = kernel.instantiate(t_serv.id);
        OCM_IRefList ilit = (OCM_IRefList)kernel.getprop(serv.id, 'I');
        IServ1 i_IServ1 = (IServ1)ilit.getIRef("IServ1");
        kernel.bind(r_IServ1.id, i_IServ1.id, b);
    }
    public int call(char op, int arg1, int arg2) { /* defined in IClient */
        ...
        int result = r_IServ1.op1(arg1, arg2); /* call the bound receptacle */
        return result;
    }
}
```

Client loads, instantiates and (first-party⁴) binds to *Server* when its *IClient.setup()* operation is called. It subsequently calls *Server.IServ1.op1()* when its *IClient.call()* operation is called.

Much of the machinery is defined behind the scenes in the automatically generated *_Server* and *_Client* classes. This includes the definition of the receptacle classes (which are conventionally named as *r_<interface_name>*—that is, *Client.r_IServ1* in our example) and the initialization of these in the constructor. The constructor also includes generated code that stores the component type’s static properties into the kernel, together with the component’s interfaces and receptacles (using the “standard” “I” and “R” keys mentioned above). The mirror-image of this latter mechanism, together with a helper class called *OCM_IRefList*, is employed to obtain the *IServ1* interface from the newly instantiated *Server* component. Then *bind()* is called to bind the *Client*’s *r_IServ1* receptacle to the just-obtained *i_IServ1* interface. The kernel type-checks the arguments passed to *bind()* using Java reflection.

The C and C++ language bindings follow a similar pattern, but suitably adapted to the target language’s capabilities. For example, in the C binding, a component is represented by a *.c* file that contains implementations of all the component’s interfaces’ operations. The user must provide *startup()* and *shutdown()* lifecycle-management functions, and take responsibility for allocating and deallocating within them the memory for receptacles and interfaces. The IDL compiler generates a per-component header file that contains the necessary definitions of receptacles and interfaces represented as structs, with pointer-to-function members representing operations, and a per-type unique identifier (UID) member to help with type checking (the latter is used by the kernel to type-check the receptacle and interface arguments passed to *bind()*). The action of *bind()* is simply to assign the function pointers in the interface struct to those in the receptacle struct, resulting in an extremely minimal and efficient implementation of component binding. Interface operations are represented as user-provided C functions which conventionally take as their first argument the UID of the component instance being invoked. Operation invocations on bound receptacles are realized using per-operation IDL-compiler-generated *CALL_<opname>(<receptacle>, <arglist>)* macros which transparently dereference the given receptacle to determine and call the appropriate target C function with the appropriate first argument.

4. THE PLATFORM EXTENSIONS

4.1 Overview

The platform extensions augment the basic programming model elements discussed above with new, optional, abstractions and services that play a major role in delivering the tailorability and extensibility promised by OpenCom. The platform extensions are of three kinds, caplets, loaders, and binders, and they are collectively supported by a CF called the *Platform Extensions CF*.

⁴We provide a complementary example that illustrates third-party binding in Section 4.4.2.

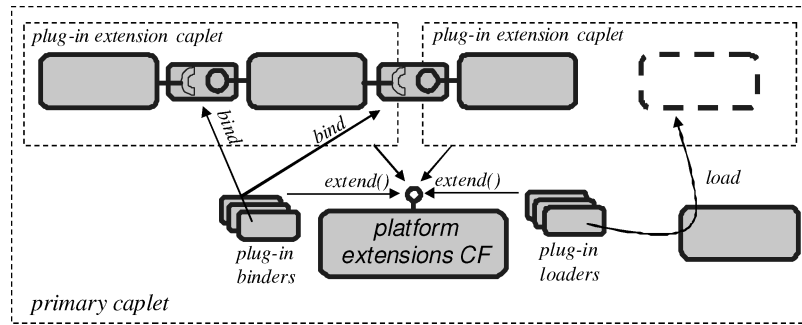


Fig. 3. The programming model extended with caplets, loaders, and binders.

Figure 3 visualizes caplets, loaders, and binders as being “plugged into” the platform extensions CF using an *extend()* operation (see later).

In brief, *caplets* are specialized plug-in component-support environments that can be dynamically instantiated within a capsule; *loaders* are plug-ins that are responsible for loading components into caplets in various different ways; and *binders* are plug-ins that are responsible for creating bindings both within and across a capsule’s caplets in various different ways. In the context of the Platform Extensions CF, the “original” capsule environment (i.e., before the Platform Extensions CF was loaded) is referred to as the *primary caplet*, and all subsequently loaded caplets are referred to as *extension caplets*. Similarly, the “original” loader and binder (which are implicit behind the kernel’s *load()* and *bind()* calls) are referred to, respectively, as the *primary loader* and the *primary binder*, and all subsequently loaded ones are known as *extension loaders* and *extension binders*. We motivate plug-in caplets, loaders, and binders, and discuss them in detail, in Section 4.3 below.

As well as hosting these plug-ins, the Platform Extensions CF recognizes and embodies an implicit system development methodology that draws a clean distinction between two distinct “programmer roles” as follows:

- (1) The deployment environment programmer (hereafter called the *environment programmer*) creates suitable caplets, loaders, and binders for a particular deployment environment using facilities native to that environment (this relates to the pink areas in Figures 2, 3, and 4).
- (2) The target system programmer (hereafter called the *system programmer*) then develops target systems using the APIs described in Sections 3 and 4.2, together with a specific palette of caplets, loaders, and binders that has been provided by the environment programmer (this relates to the blue areas in Figures 2, 3, and 4).

Distinguishing these two roles is key to OpenCom’s approach of offering a simple and generic programming model in a diverse and extensible environment of target domains and deployment environments. The use of the two roles is discussed in Section 4.4; they are further discussed in the context of case studies in Section 7.

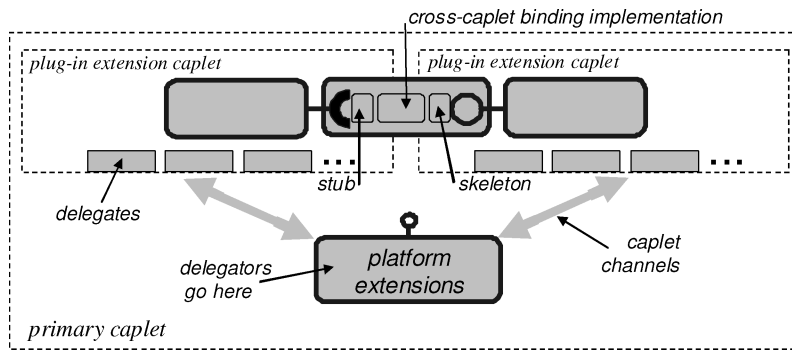


Fig. 4. Environment programmer concepts.

4.2 The Platform Extensions CF's API

The API offered by the Platform Extensions CF is as follows:

```
interface Platform_Extensions {
    status extend(in extension_type ext_type,
                 in string ext_name, out component ext);
    status load(in string name, in component loader,
               in component caplet, out template t);
    status instantiate(in template t, out component c);
    status unload(in template t);
    status destroy(in component c);
    status bind(in interface i, in receptacle r,
                in component binder, out component binding);
    status setdefaultextension(in extension_type ext_type,
                               in component ext);
    status notify(in ICallback callback);
}
```

It will first be observed that the API is similar to, but builds on, the basic kernel API given in Section 3. *Extend()* calls on the primary loader to load and instantiate a component that will play the role of an extension. The *ext_type* argument is an enumerated type {CAPLET, LOADER, BINDER} which specifies which one of the three possible extension types is intended; the *ext_name* argument then refers to the name of a component type that will play the corresponding role (i.e., the role of a caplet, a loader, or a binder). *Load()* is like the similarly named operation in the kernel API; the difference is that this version allows a particular loader and caplet to be specified. Likewise, *bind()* allows the specification of a particular binder. Every time a component model element (i.e., a template, component, interface, or receptacle) is created, the creating component must record appropriate registry entries pertaining to the element using *kernel.register()* and *kernel.putprop()* (e.g., a loader that creates a new component instance should register the latter and the latter's interaction points). This gives newly created entities a UID and makes them visible to the rest of the system as if they were created as a result of calls on the kernel API.

Setdefaultextension() is used to designate an extension of a given *extension_type* (i.e., a caplet, loader, or binder) as a default to be used in subsequent calls to *load()* and *bind()* if an argument of DEFAULT is passed as the (respectively) caplet, loader, and binder arguments of these calls. This “default” facility is useful in helping to manage the complexity of simultaneously supporting many caplets, loaders, and binders. The basic idea is to employ the “strategy” pattern [Gamma et al. 2004]: (i) the environment programmer writes a “metaloader” or “metabinder” that knows the set of loaders (binders) available in a given deployment environment, and the conditions under which each is used (e.g., to deal with specific component types; to load into specific caplets, etc.); (ii) the environment programmer uses *setdefaultextension()* to designate this metaloader (metabinder) as the default; and (iii) when the system programmer calls *load()* (or *bind()*), the metaloader (metabinder) dispatches to a specific caplet, loader, or binder according to its inbuilt knowledge.

Finally, *notifiy()* is identical in function to the similarly-named operation in the kernel API. Its use is discussed in Section 5.

4.3 Caplets, Loaders, and Binders

4.3.1 Caplets. In terms of *motivation*, caplets, as mentioned, are specialized component-support environments that can be dynamically instantiated within a capsule. There are *three* main motivations for caplets. The *first* is for different caplets to represent different technology domains in the underlying deployment environment. For example, if it was desired to build a system that comprised both C++ and Java components, this could be achieved by employing a separate caplet for each of the two language environments. In such cases, the caplets might typically be realized as OS processes, with one executing a Java Virtual Machine (JVM). Alternatively, if a deployment environment consisted of multiple bus-connected microcontrollers, each with its own private memory, a caplet could be used to represent each microcontroller/memory pair. The essential difference between supporting multiple caplets in the same capsule and simply employing multiple separate capsules is that in the former case all the contained components, regardless of which caplet they are in, see a common name space and a single kernel API instance. This enables the “third-party” loading and binding semantic of the kernel to operate transparently across caplets.

The *second* motivation for caplets is to provide privacy and isolation between components that are mutually distrustful, or which have different privileges. For example, when building an OS environment, one might implement the OS kernel as one caplet and the user space as another (or the user space could be represented using multiple caplets, one per process; or caplets could be used to impose protection domains in a single address space). A similar strategy could be adopted in an active networking environment where it was necessary to isolate user-provided functionality from system functionality so that the latter could not crash the former (see, e.g., Karlin and Peterson [2001]). To manage privacy and isolation, extension caplets can choose whether or not to allow their components access to the kernel: where required, extension caplets arrange

kernel access for their hosted components by providing a kernel proxy, and using a cross-caplet binder to bind this to the kernel interface in the primary caplet. But this arrangement can be selectively disallowed either by the primary caplet or by the extension caplet itself.

The *third* and final motivation for caplets is to support heterogeneous *component styles*—that is, different implementations of the abstract component concept. The component style supported by the primary caplet is known as the *primary component style*; styles supported by extension caplets are known as *extension component styles*. While the semantics of all component styles must conform to the general characteristics given in Section 3 (i.e., support for interfaces and receptacles, etc.), each extension component style is free to take its own position on a range of issues such as the following:

- (1) the layout of components on disc and in memory—this may be language/compiler/OS specific;
- (2) whether components can be instantiated multiple times or are singletons (i.e., can be instantiated only once);
- (3) whether components may support an arbitrary set of interaction points or if these are somehow constrained (e.g., in terms of numbers or types of interfaces, or numbers or types of operations in those interfaces);
- (4) whether components support fixed sets of interaction points or if these can be dynamically created;
- (5) whether components are represented as native executables or as interpreted code (e.g., Java).

There are two major reasons to support extension component styles. The first is to be able to accommodate components written using existing component models (e.g., for purposes of reuse, integration, or backward compatibility). For example, we could integrate Microsoft COM components and JavaBeans in a single system by providing a caplet for each of these component styles, together with a suitable cross-caplet binder. The second reason for supporting extension component styles is to support “specialized” styles. For example, in a primitive resource-poor deployment environment such as a microcontroller or sensor network element, we could define a minimal component style that imposed severe restrictions on the numbers of interfaces components can support, or the types of arguments that can be passed to operations (e.g., integers only). Nevertheless, such specialized styles still look exactly the same to external third-party code that deploys and binds components in the standard manner supported by the enclosing capsule’s kernel API. See Section 7.2 for a detailed discussion of such a case.

In terms of the *realization* of caplets, from the point of view of the system programmer, a caplet is simply a (primary-style) component that is loaded into the primary caplet (where the Platform Extensions CF itself resides). This primary-style component encapsulates all the deployment-environment-specific machinery needed to realize its particular instantiation of the caplet concept. The necessary machinery is created and/or initialized when the “caplet

component” is first instantiated (more detail is given below in Section 4.4.1). At the environment programmer level, however, there is a basic requirement that caplets must provide a basic communicational facility to enable them to interact with the rest of the deployment environment (e.g., so that loader and binders can work with them). To meet this requirement the primary style component that represents a caplet must implement the following interface:

```
interface Caplet {
    int createchannel(void);
    status destroychannel(in long channel);
    status sendmessage(in long channel, in any message);
    any receivemessage(in long channel);
}
```

In other words, this interface is required for a component to be recognized as a valid caplet by the Platform Extensions CF. The fact that caplets provide an execution environment for components but are otherwise passive is reflected in the form of this interface which, as can be seen, provides only generic message passing services and does not support any caplet-specific functionality. The way in which these message passing services are used by loaders and binders that will be associated with the caplet is discussed below in Section 4.4.1.

4.3.2 Loaders. In terms of *motivation*, extension loaders are used to load components into a capsule (or, more specifically, a caplet) in some particular manner. In many cases, extension loaders are closely associated with particular caplet types. For example, a caplet type that supports a particular component style would typically have an associated loader that knows how to load and instantiate components of this style. However, the concept of pluggable loaders has a much wider applicability than this. In particular, separating the loader and caplet concepts allows one to associate several loaders with a particular caplet type, or to share common loaders across multiple caplet types. It also allows us to provide different loaders with specialized semantics and behaviors. For example, different loaders might get component templates from different places (e.g., from different repositories or over the network). Similarly, loaders might perform security checks on the templates they load and/or instantiate, or validate particular properties, or perform special behaviors on loading/instantiation. As an example of the latter, a loader might use reflection (see Section 5) to transparently analyze a component’s receptacles when loading it, and then recursively preload the full set of components on which it depends; or another loader could load balance across a set of processor/memory units managed within a single caplet.

In terms of *realization*, as with caplets, a plug-in loader is, to the system programmer, simply a primary-style component that offers a facade that hides arbitrary deployment-environment-specific functionality. To be recognized as a loader by the Platform Extensions CF, loader components are required to implement the following interface, the operations of which are called by the

Platform Extensions CF as a result of prior calls of the latter’s corresponding calls:

```
interface Loader {
    status load(in string name, in component caplet, out template t);
    status instantiate(in template t, out component c);
    status unload(in template t);
    status destroy(in component c);
}
```

4.3.3 *Binders*. The motivation for plug-in binders is to represent different “binding mechanisms” in the underlying deployment environment. For example, different binders can abstract over binding mechanisms such as interrupts, traps, special buses, shared RAM, optimized register transfers, nearest-neighbor registers in pipeline architectures, or OS-level IPC calls. The plug-in binder abstraction makes all such features uniformly available to the component programmer both within and across caplets, and between components of a common style or different styles. In addition, binders can support special behaviors such as operation interception, performing security checks on invocations, or supporting “actor”-like concurrency models in which thread context switches take place while a thread is executing inside a binding.

Binders can vary widely in their complexity. At the more complex end of the scale, bindings that operate across caplets or component styles may need to incorporate stubs and skeletons to mediate between components that assume different calling conventions (perhaps because they were generated by different compilers), or which employ different representations of types or different language semantics. In some cases, these stubs and skeletons may be automatically generated by an IDL compiler in the classic “middleware” style; in other cases they might be hand coded for performance reasons. In all cases, however, the necessary complexity is completely encapsulated within the specific binder and is thus hidden from the user of the generic component-based programming model.

Note that the possibility of cross-caplet binding raises questions concerning the degree of coupling and the distribution of the caplets that comprise a single capsule. The answers are deployment environment specific, but all bindings within an OpenCom capsule are at minimum assumed to be *reliable* in the sense that the semantic of making a call over a cross-caplet binding should be indistinguishable, apart from a slightly greater latency, from that of making a call over an intracaplet binding. That is, the binding must not drop or reorder any messages. Given this, one would not usually expect the degree of coupling and distribution within a capsule to be so loose and widely distributed that cross-caplet binders would need to implement complex middleware-like functionality in order to maintain the required degree of reliability. If a target system must operate in such a loose and widely distributed environment, the preferred approach would be to design the system not as a set of caplets within a single capsule, but as a number of capsules containing middleware-like CFs (as described in the case study of Section 7.3).

In terms of realization, like caplets and loaders, plug-in binders are simply primary-style components that offer a facade that hides arbitrary deployment-environment-specific functionality. To be recognized by the Platform Extensions CF, they are required to implement the following interface:

```
interface Binder {
    status bind(in interface i, in receptacle r, out component binding);
    status destroy(in component binding);
}
```

4.4 Programmer Roles

4.4.1 The Environment Programmer Role: Creating Platform Extensions. As mentioned, the Platform Extensions CF sees plug-in caplets, loaders, and binders simply as primary style components that support specified interfaces. Behind these interfaces, however, can lurk a great deal of deployment-environment-specific complexity that the environment programmer (but not the system programmer) must deal with. For example, consider a “Java caplet” that encapsulates a JVM and supports a Java-based OpenCom component style. Having been instantiated with a call of `extend()`, such a caplet might proceed by forking a new process to run a JVM, and then establish contact with this new process using some OS-specific IPC mechanism—for example, a UNIX pipe.

In such cases, that is, whenever there is a physical separation between the primary-style extension component and the software that implements the extension, we typically employ a “delegator-delegate” pattern in which the extension component plays the role of a *delegator* and the “separate software” plays the role of a *delegate*. The two players communicate, using some appropriate protocol, over the channels supported by caplets. The use of the delegator-delegate pattern is illustrated in Figure 4. (Figure 4 also shows the binder-related environment programmer concepts of stubs, skeletons, and cross-caplet binding implementations—these are dynamically created by plug-in binders as required.)

As a more concrete example of the use of the delegator-delegate pattern, consider the following sequence of steps that might be taken in instantiating a new loader for the above-mentioned Java caplet example:

- (1) In the primary caplet, the loader component (delegator) is loaded using `extend()`.
- (2) This delegator opens a channel to the Java caplet using `createchannel()` in the Java caplet’s interface.
- (3) The delegator uses a Java-caplet-specific protocol to send (using `sendMessage()`) a command to the Java caplet delegate to ask it to load the loader’s delegate (loading the loader’s delegate can be done using any appropriate means; as we are in the environment programmer domain here, this is outside the scope of the OpenCom programming model).
- (4) The delegator opens a new channel and uses this to communicate, using a loader-specific protocol, with its newly established delegate.

One would use similar steps to instantiate a binder plug-in for the Java caplet. Cross-caplet binders employ a single delegator and multiple delegates—one in each of the caplets it knows how to deal with. We consider more concrete examples of delegator-delegate-based extensions in Section 7.2.

4.4.2 The System Programmer Role: Using Platform Extensions. To further appreciate the distinction between the environment and system programmer roles, consider the simple program below in which an system programmer builds on a set of platform extensions that have been provided by an environment programmer. In this program, the system programmer instantiates and a third party binds a primary-style component in the primary caplet to a component in the “Java caplet” discussed above. The system programmer also employs an extension loader that loads components into the Java caplet (as discussed above), and an environment programmer provided an extension binder that binds components across the two caplet types.

```
template jtemp, ptemp;
component jcaplet, jloader, cbinder, jcomp, pcomp, binding;
interface ifaces[N];
receptacle recpts[N];

/* set up the Java caplet with a Java specific loader and binder */
extend(CAPLET, ‘MyJavaCaplet’, jcaplet);
extend(LOADER, ‘MyJavaLoader’, jloader);
extend(BINDER, ‘MyJavaToPrimaryBinder’, cbinder);

/* load and instantiate the components */
load(‘PrimaryComp1’, PRIMARY, PRIMARY, ptemp);
load(‘JavaComp1’, jloader, jcaplet, jtemp);
instantiate(ptemp, pcomp);
instantiate(jtemp, jcomp);

/* obtain interaction points on the components and bind them */
getprop(pcomp.id, ‘R’, recpts);
getprop(jcomp.id, ‘I’, ifaces);
bind(ifaces[0], recpts[0], cbinder, binding);
```

The key point to notice is the transparency and generality of this system programmer code. First, the structure of the code is independent of the underlying caplet structure and caplet types involved: it would be essentially similar if we were dealing with only one caplet, or if the primary caplet ran on a PC and the extension caplet represented a primitive microcontroller attached to the PC by a PCI bus and supported a component style in which components were realized as small segments of machine code (see Section 7.2). Second, the code could be executed unchanged with the same effect from within any component in the capsule (including one running in the microcontroller caplet cited above). This means that system programmers do not need to know or care in which caplet their components will execute, and that the choice of an optimal target caplet

for a newly loaded component can be left to the platform (i.e., to an extension loader or metaloader).

5. THE REFLECTIVE EXTENSIONS

5.1 Overview

The purpose of the OpenCom reflective extensions is to support the construction of *dynamic* target systems that need to change or evolve during their execution in a controlled and principled manner. To achieve this, the reflective extensions provide generic support for inspecting, adapting, and extending the structure and behavior of systems at runtime. They also help to maintain an architectural separation of concerns between *system building* (or “base-level” programming) and *system configuration and adaptation* (or “metaprogramming”). This distinction is orthogonal to the distinction between the environment and system programmer roles discussed above.

Following Maes [1987], we view the essence of reflection as enabling the inspection and manipulation of “causally connected metamodels” of a software system. Causally connected metamodels are representations of some aspect of the system under consideration, and they expose a so-called metainterface through which the representation can be inspected and manipulated. *Causal connection* means that (i) runtime changes made to a metamodel (effected via its metainterface) cause corresponding changes to immediately be reflected in the represented system; and (ii) changes in the represented system that occur due to some external cause are similarly reflected in the metamodel. A key principle of our approach to reflection support is to provide an *extensible set* of orthogonal metamodels, each of which is optional and can be dynamically loaded when required, and unloaded when no longer required (assuming no dependencies). We have found in our experimentation to date that the metamodels described in the following subsection are particularly useful.

5.2 Example Reflective Metamodels

5.2.1 *The Interface Metamodel.* This provides two related capabilities: (i) to dynamically discover (at runtime) details of the interaction points of a component in terms of their operation signatures; and (ii) to perform “dynamic invocations” on dynamically discovered interfaces. Together, these capabilities enable components to invoke interfaces which were not known to them at compile time (i.e., they need not support the requisite receptacles). Essentially, these capabilities are similar to Java core reflection except that they work at the OpenCom level and are therefore programming language independent. One possible drawback of the interface metamodel is that its dynamic invocation capability allows one to bypass the architectural structure of an OpenCom system in terms of its explicit bindings. But in some circumstances it is important to support such dynamic behavior. For example, it is particularly useful in supporting generic functionality such as debugging, component database browsing, or generic bridging (cf. the CORBA dynamic invocation interface or DII [Object Management Group 1995]). It is also required to support the architecture

and interception metamodels discussed below with the necessary typing information.

In implementation, the interface metamodel is realized as a singleton primary-style component that uses the OpenCom kernel's *getprop()* and *putprop()* APIs to store and obtain the interface types associated with a component, and pointers to IDL definitions of these types in an encapsulated IDL repository. It implements dynamic invocation by offering a generic *invoke()* API which is modeled on the CORBA DII interface. The arguments to *invoke()* comprise the target interface and operation, together with a stack of argument values that is built manually by the invoking component on the basis of the runtime IDL type information provided by the metamodel. Internally, the metamodel establishes a binding to the target interface in the normal manner (i.e., using the *bind()* call) so that the called component is unaware of the fact that it is being invoked in an unusual manner.

5.2.2 The Architecture Metamodel. This represents the topology of the current set of components within a capsule. It is used primarily to achieve coarse-grained topological inspection, adaptation, and extension of the structure of a dynamic target system. For example, in a media-streaming scenario, we have used it to dynamically manage the set of media codecs in use when a mobile PDA migrates between fixed and wireless networks [Blair et al. 2004]. This involves first inspecting the underlying component topology to locate the codec component, and then adapting/extending the topology to effect corresponding change (e.g., to replace the codec). The metamodel provides a “graph-oriented” API in which components are represented as nodes and bindings as arcs. Inspection is carried out by traversing the graph, and adaptation/extension is achieved by adding or removing nodes or arcs (e.g., adding a node results in the deployment of a new component). An example of the use of the architecture meta-model is given in Section 7.2.

In implementation, the architecture metamodel is realized as a singleton primary-style component that provides a topological view of the capsule contents based on (i) the raw component/interface/receptacle data in the kernel's registry; (ii) the typing information maintained by the interface metamodel; and (iii) current binding information as provided by a *notify()* callback (see Section 3). When the architecture metamodel is used to adapt/extend the system topology (e.g., by adding a node as above), it effects the necessary changes by using appropriate kernel calls to load components, create bindings, etc., as required.

5.2.3 The Interception Metamodel. This metamodel, which is a version of probably the most widely explored reflective mechanism in general use [Kon et al. 2002], exposes the process of invoking an operation in a component's interface. More specifically, a metainterface is provided that allows the metaprogrammer to insert arbitrary code elements called *interceptors* within bindings, such that an interceptor is executed whenever an operation is invoked across the binding (more specifically, either before, or after, or both before and after, the invocation). Such an interceptor might, for example, audit the pattern of

invocations and their arguments for debugging purposes, or dispatch invocations to an alternative object instance (“hooking”), or insert a security or concurrency control check on an invocation. Interception is especially useful in adaptation scenarios; for example, in the above-mentioned media streaming/mobile computing scenario, an interceptor on a low-level protocol component could be used to monitor and detect the conditions under which a codec should be replaced by means of the architecture metamodel. In addition, interception can be used as a basis for dynamic aspect-oriented programming [Bencomo et al. 2005].

A commonly cited disadvantage of interception is that it incurs an inherent performance overhead whether or not an interceptor is actually installed [Coulson et al. 2004]. This is because interception typically requires bindings to support a level of indirection. Our realization of the interception metamodel uses the plug-in binder concept to sidestep this disadvantage. That is, we provide both interception-capable and non-interception-capable binders and select from these according to requirement—that is, we choose an interception-capable binder only where we are likely to need interception. If we choose wrongly we can straightforwardly recover by destroying the current binding and rebinding using a different binder. An additional advantage of perbinder interception is that we can provide alternative models of interception that use different underlying implementations offering different tradeoffs.

5.3 Controlling Access to Reflective Metamodels

Reflection is a powerful and general technique, and its use should always be constrained to minimize programmer errors. Our approach to providing such constraint is to limit the set of components that can access the reflective metamodels: in particular, access is typically given to CFs but not to their plug-ins. As well as preventing spurious access, this helps ensure that metamodels are accessed only when conditions are “safe”; for example, a CF might restrict component replacement via the architecture metamodel to situations in which no invocations are currently being made on interaction points owned by the “old” component. Further, a CF could define a suitable state-transfer protocol to carry over essential state from the old component to the new one.

6. PERFORMANCE AND OVERHEADS

We now discuss the inherent performance properties and overheads of Open-Com. This section offers a generic treatment; more specific performance evaluations involving measurements of particular systems constructed using the technology are given in Section 7.

In assessing inherent overheads, we recognize a key distinction between *in-band* and *out-of-band* execution [Coulson et al. 2004]. In-band execution refers to segments of code that are repeatedly executed in the normal course of events and are therefore particularly performance sensitive; out-of-band execution, on the other hand, refers to code segments that are executed only “occasionally” to the extent that their impact on system performance is negligible. On the basis of this distinction, it can be seen that the kernel inherently incurs zero in-band

execution overhead—this is because it is only involved in out-of-band operations, namely, loading, instantiating, and destroying components, and creating and destroying bindings. These operations are typically only invoked when a system is being (re)configured; that is, an established component topology need not make *any* calls on the kernel. Note that this is quite unlike the situation in operating system microkernels, which are unavoidably involved in critical in-band operations (such as interrupt handling or thread scheduling). Note also that the “out-of-band” characterization, with the exception of the overhead introduced by interception-capable binders (see below), also applies to the reflective metamodels.

Of course, these observations do not imply that the performance of target systems constructed using OpenCom is entirely unaffected by the use of the technology! In particular, the following overhead-contributing factors are important:

- (1) overheads inherent in the primary and extension component styles used (e.g., per-component memory overhead);
- (2) performance overheads inherent in bindings (whether created by the primary binder or extension binders);
- (3) the granularity of the componentization of the target system (which, in turn, affects the relative impact of the above two factors).

To gain insight into these overheads we now provide a brief overview of our “reference” kernel implementation⁵ and present some basic measures of its performance and overheads. The reference implementation realises its primary caplet in terms of a standard Linux process (it also runs under Windows). The primary component style is based on the standard ELF executable format emitted by the GNU C++ compiler. The primary loader is based on Linux Shared Objects (or Windows DLLs), and the primary binder uses standard C++ vtables to realize bindings. As well as the primary binder, we have two extension binders that work with primary-style components in the primary caplet: the first of these offers an implementation of the interception metamodel; the second optimises away the vtable apparatus by replacing the usual vtable-based indirected call in the caller’s code segment with a simple CALL instruction. While this CALL-based extension binder saves some overhead, it can only be used where the component on the receptacle side is a singleton. Note that the CALL-based binding works by modifying the code associated with the receptacle in the caller’s code segment. Such code rewriting approaches are sometime dangerous; our modification, however, is small, well defined, and constrained in scope.

Using our reference implementation, we carried out the following measurements of inherent overhead, the results of which are summarized in Table I:

⁵In addition to our C++-based reference implementation, we have Java and C implementations, as discussed in Section 3.3. In all experiments, we used a Dell Precision 340 series workstation equipped with 4×1.6 -GHz Pentium processors and 512 MB of RAM, and running Linux Redhat 8.0.

Table I. Summary of Measures of Inherent Overhead

| | OpenCom | Native C++ |
|--|--------------|-----------------------------------|
| Memory footprint of kernel | 32 kbytes | N/A |
| Memory footprint of null component | 36 bytes | 20 bytes (null C++ object) |
| Mem. footprint of receptacle/interface | 28 bytes | N/A |
| Component loading time | 9.8 μ s | 7 μ s (null C++ object) |
| Component instantiation time | 0.47 μ s | 0.28 μ s (null C++ object) |
| Time to creating a primary binding | 2.4 μ s | 2.4 μ s (C++ vtable creation) |

Table II. Overhead of OpenCom Operation Invocation

| | Calls/s | C/C++ calls/s |
|----------------------------------|--------------------|--------------------------------------|
| Primary binding | 101×10^6 | 101×10^6 (C++ method calls) |
| Intercepted binding | 5.8×10^6 | N/A |
| Primary binding (nonnull op) | 0.54×10^6 | N/A |
| Intercepted binding (nonnull op) | 0.5×10^6 | N/A |
| CALL-based binding | 198×10^6 | 198×10^6 (C function calls) |

- (1) *Memory footprint of kernel.* We measured this as 32 kbytes. This is a modest overhead which should make the kernel deployable in many resource-scarce environments.
- (2) *Memory footprint of a null component.* The memory requirement of a null primary-style Component (i.e., one with no interfaces and receptacles and null initialization/finalization routines) is 36 bytes. This compares to an overhead of 20 bytes for a single null C++ object. In addition, the per-interface and per-receptacle overhead is 28 bytes (with an additional 5 bytes/operation).
- (3) *Component loading and instantiation time.* The time taken to load a single null primary-style component (averaged over a few million loads) was measured as 9.8 μ s (compared to 7 μ s for a Linux Shared Object containing a null C++ object); and the time to instantiate an already-loaded null component was measured as 0.47 μ s (compared to 0.28 μ s for a null C++ object). The small overheads here are attributable to the larger file size of the OpenCom component template (due to metadata, etc.), and the slightly more complex instantiation process, including interaction with the kernel registry.
- (4) *Time to create a primary binding.* This was measured as 2.4 μ s—which is identical to the time required to create a vtable in C++.

Next we measured the time taken to perform invocations across different types of bindings. The details, which are summarized in Table II, are as follows:

- (1) *Overhead of calls made across a primary binding.* For measuring this, we performed a series of invocations across a primary binding involving an interface with a single operation with no arguments and a void return value. We measured 101×10^6 calls/s. We then confirmed that, as expected, this is identical to the invocation rate achieved when calling a method in a simple C++ object (which employs vttables in an identical manner). Thus there is zero in-band overhead arising from the use of the primary binder.

- (2) *Overhead of calls made across an intercepted binding.* This test used the same interface as above but employed a binding created by our interception-capable extension binder. We used a null “before” interceptor and measured 5.8×10^6 calls/s. Comparing this overhead with that of the primary binder (see above), this represents a slowdown factor of 17.4—thus it is clear that there is a substantial cost incurred for interception. However, to put this into perspective, we repeated the comparison using a target operation with a more typical and representative nonnull body—namely, an empty loop of 1000 iterations. This time we measured 0.54×10^6 calls/s for the primary binding and 0.5×10^6 for the intercepted binding—a slowdown factor of only 1.08. This illustrates that the absolute overhead of calls is still very small. Bear in mind also that, as explained in Section 5, a system programmer only needs to incur this cost when interception is actually required.
- (3) *Overhead of calls made across a CALL-based binding.* This test again used the above-described interface but this time employed our CALL-based singleton binder. We measured 198×10^6 calls/s, thus revealing that singleton bindings yield a twofold speedup over primary (vtable-mediated) bindings. We also confirmed that, as expected, this is identical to the performance achieved when calling a null C function in a simple C program. Again, this indicates zero in-band overhead arising from the use of the binder abstraction.

Overall it can be fairly concluded from the above figures that the basic approach to building component-based systems adds little overhead to that of a traditional programming-language-based systems development environment.

7. CASE STUDIES

7.1 Overview

Having described and evaluated the basic OpenCom technology, we now report on our experiences with the technology and present two representative and contrasting case studies of the use of the technology to build nontrivial target systems.

In terms of experiences, OpenCom has been used by over 25 programmers at several universities, with different skill levels ranging from third-year undergraduates to experienced systems researchers. These programmers have applied OpenCom in a diverse range of systems domains including middleware (e.g., OpenORB [Parlavantzas 2005] or ReMMoC [Grace et al. 2003]), sensor networks (e.g., Gridkit [Hughes et al. 2006]), embedded systems (e.g., RUNES [Costa et al. 2007]), programmable networking (e.g., NETKIT [Coulson et al. 2003]), and overlay networks (Open Overlays [Grace et al. 2005]). Some of these systems have been quite large—for example, OpenORB employs over 40 components and comprises 63,000 lines of code, while ReMMoC employs over 25 components and comprises some 30,000 lines of code. Furthermore, over 10 OpenCom developers have worked on multiple projects in different domains; for example, in both mobile computing and sensor networks. On the basis of this experience, we can report anecdotally that it is indeed the case that, thanks

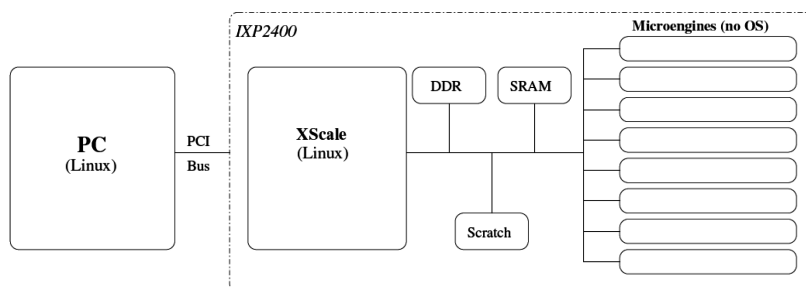


Fig. 5. Outline of an Intel IXP-based router.

to OpenCom’s uniformity of approach, a developer’s prior experience in one domain considerably reduces her or his time to develop in a second domain.

In terms of case studies, we focus in the following subsections on the above-mentioned OpenORB and NETKIT projects. For each case study, we discuss in context the application of the two programmer roles (system and environment programmers), and also provide an overall evaluation, including a performance evaluation.

7.2 Case Study 1: Programmable Networking Environments on Network Processors

7.2.1 Introduction. Network processors [Network Processing Forum 2005] are specialized multiprocessor devices that process and forward network packets at gigabit speeds. Unlike traditional high-speed routers, network processors (hereafter called *NPs*) process packets *in software*—primarily for reasons of flexibility and modifiability. They offer an ideal case study for a systems-level software component model because (i) they are widely acknowledged to be very difficult to program (e.g., they feature specialized hardware and often have no standard OS environment); and (ii) the software must be extremely efficient to meet the demands for gigabit forwarding speeds. In this section, we briefly describe work carried out in our recently completed NETKIT project [Coulson et al. 2003] which explored the use of NPs (specifically, the Intel IXP range of NPs [Intel 2004]) in building flexible programmable networking environments.

An Intel IXP-based router, as illustrated in Figure 5, consists of an IXP card connected to a host PC (running Linux) via a PCI bus. IXP cards comprise the following: an XScale processor running Linux and serving as a general-purpose control processor; an array of so-called microengine *reduced instruction set* (RISC) processors that are attached to each other and to the control processor through a series of buses and a hierarchical DDR/SRAM shared memory architecture; and a set of dedicated hardware elements (not shown) such as network interfaces and a hardware hash unit. The microengines are responsible for “fast-path” packet forwarding. They execute small code modules, written in C or assembler, that run in a “bare-iron” environment (i.e., no OS). They have general-purpose, although primitive, instruction sets, and also support specialized packet-forwarding functionality such as checksumming, hardware timing, and pseudorandom number generation. As well as having access to

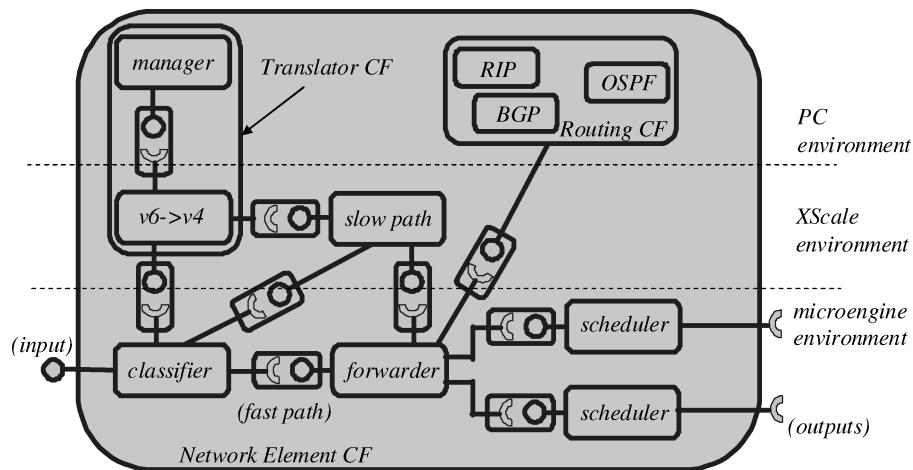


Fig. 6. Outline of an Intel IXP-based router.

shared memory, each microengine has a private instruction/data memory called a *microstore*.

7.2.2 Programmer Role Considerations. The *system programmer* role works entirely in terms of the abstract OpenCom concepts described in Sections 3 and 4. Using these concepts, we have developed a *Network Element CF*, and also a nested *Routing CF* that takes various routing protocols such as RIP, BGP, or OSPF as its plug-ins, and a nested *Translator CF* that is capable of translating between different versions of IP. These are illustrated in Figure 6, which also indicates which of the three environments (i.e., the PC, the XScale, or the microengines) that each CF and component resides in. Essentially the *Network Element CF* [Coulson et al. 2003] defines “hot spots” in a software router architecture that accept different component-based implementations of network elements such as classifiers, forwarders, and schedulers. It is thus closely related to programmable networking systems such as VERA [Karlín and Peterson 2001]. However, the key point is that by building on the set of plug-in caplets, loaders, and binders provided by the environment programmer (see below), the CF developer is able to largely disregard the underlying complexity and heterogeneity of the IXP architecture. Thus, when the *Network Element CF* accepts a new forwarder provided by its user, and must bind this to an existing classifier and a scheduler, it does not need to act differently depending on which microengines these components reside on, or what communication mechanism is being used to bind their interfaces and receptacles.

The power and generality of the OpenCom approach is also evident in dynamic reconfiguration scenarios involving the reflective metamodels. Consider, for example, a scenario in which a *Network Element CF* user dynamically installs an IPv6-to-IPv4 protocol translator. The initial CF configuration, illustrated in Figure 6, comprises several components on the router’s fast path, namely, a classifier and a forwarder, scheduling components, and a component for processing IP options on the slow path. These various components are

assigned by the Network Element CF to caplets in a way that best exploits the hardware capabilities of the IXP-based router; for example, the CF would typically deploy the fast-path components in the microengine caplets, the IP options component in a XScale caplet, and the Routing CF in an XScale caplet or even a PC caplet.

Initially, the Translator CF accepts from the user a $v6 \rightarrow v46$ protocol translator. This is offered to the CF via a CF-specific interface which uses the interface metamodel to check the to-be-loaded component type for conformance to the CF's criteria for plug-ins. The Translator CF then employs a kernel-level metaloader, as described in Section 4.1, to transparently select an appropriate loader for the component type. The selected loader again uses the interface metamodel to check that the to-be-loaded component conforms to its own requirements (e.g., that microengine-hosted components support only single-integer arguments/return values, as discussed above). The Translator CF then uses the OpenCom kernel's *getprop()* operation to locate the classifier's receptacle, and, by manipulating a CF-specific API that encapsulates and constrains the architecture metamodel, arranges for this to be bound to the translator. Here, the CF checks that the proposed binding does not violate its rules for structural integrity; having verified this, the CF asks the metabinder to actually do the work. It could additionally add an interceptor to the binding to, say, monitor and log the number of IPv6 packets actually forwarded.

The first decision to be taken by the OpenCom *environment programmer* role when deploying OpenCom on an Intel IXP-based router is the scoping of capsules—for example should there be separate capsules for the PC, XScale, and microengines, or should a single capsule span the entire router? We chose a single all-encompassing capsule on the basis that this promotes the highest degree of integration between the different areas of the router hardware. The next decision is where to locate the capsule's primary caplet. Given that our reference kernel implementation runs under Linux, and that the microengines themselves are too primitive to support the kernel, the choice was between the host PC and the XScale processor. We chose the XScale because it is more “central”—that is, it is in more direct contact with the microengines—but the PC would also have been possible. Next we must define a suitable set of extension caplets and corresponding component styles, loaders, and binders. In the XScale and PC environments, the caplets were implemented as standard OS processes, but in the microengine environment we designed a specialized caplet type (see below) that runs on each microengine. Note that other caplet-to-microengine mappings are possible. For example, it would have been possible to encapsulate all the microengines within a single caplet whose loader “intelligently” assigns components to specific microengines on the basis of their current loading, and perhaps even migrates already-running components between microengines according to some resource management policy.

In the remainder of this section, we focus on the above-mentioned microengine caplet, which is of most interest in illustrating the generality of our systems-building approach. The implementation of the microengine caplet itself is relatively straightforward: its only function is to establish a caplet channel from the primary caplet on the XScale to the target microengine. We realized

the caplet channel pattern in terms of libraries provided by Intel that allow one to directly access a microengine's microstore from the XScale processor. Thus no explicit caplet delegate was required. The *component style* employed by the microengine caplet is noteworthy in being strictly and severely constrained: components may have a maximum of one interface and one receptacle, and these furthermore may only support operations that accept and return a single integer. The main reason for these restrictions is that they enabled us to base the component style on an already-existing module convention adopted by Intel and the Netbind project [Campbell et al. 2002]. This considerably eased the task of developing a broad palette of functionality in terms of classifiers, filters, forwarders, schedulers, etc. Our microengine component style extends the Intel/Netbind convention by incorporating an on-disc metadata-enhanced format that supplements the bare executable with pointers to the corresponding IDL interfaces, and other metadata. We actually designed two variants of the microengine component style: the first was a "basic" version that supports only "singleton" components. With this version, if we require, say, a pipeline of n filter components, the code of the filter needs to be loaded n times, which wastes scarce microstore memory. The second, more "advanced," variant supports multiple instantiation of components. While this is more flexible, it has the corresponding disadvantage of being slightly less efficient. In particular, it assumes a level of indirection in bindings (cf. vtables); also, it involves separating out instance variables and modifying the instructions that access these, and it involves modifying intercomponent calls to pass an instance identifier on the stack.

We now discuss the loader and binder plug-ins that we designed for the microengine caplet. The *loader* is of interest because it provides the illusion of dynamic loading/instantiation despite the fact that the microengine hardware only allows modification of its microstore memory when the microengine has been stopped [Intel 2004]. The basic capability provided by the hardware is to stop the processor, access microstore locations, and then restart execution at a hard-wired microstore address. To achieve transparent loading, it is therefore necessary to not only load the new component, but also to patch the hard-wired restart address so that subsequent execution jumps to the point at which it left off. In addition, to avoid microstore fragmentation when components are being repeatedly loaded and unloaded, the loader needs to occasionally relocate components within the microstore in a transparent manner. The loader is also responsible for enforcing, by inspecting the metadata attached to an on-disc component template, the above-mentioned component-style specific restrictions on interfaces and receptacles.

We have implemented a range of *binders* for the microengine caplet. The simplest of these is an intramicroengine binder for the above-mentioned "basic" component style. This uses the caplet channel to access the microengine memory, and implements a binding by "morphing" a jump instruction in the component supporting the receptacle, to refer to the designated entry point of the component supporting the interface (this technique was pioneered in the Netbind project [Campbell et al. 2002]). The necessary entry and exit point information is obtained (via the kernel's registry) from metadata attached to the packaged

component, which is transformed by the loader from relative to absolute offsets. The intramicroengine binder for the “advanced” multi-instantiation component style instead uses the caplet channel to initialize a per-binding indirection table and instance variable vector. We have also implemented a range of *cross-caplet* binders that bind microengine-hosted components to (i) components in other microengine caplets, (ii) components in XScale caplets (running under Linux), and (iii) components in host PC caplets (also running under Linux). These are considerably more complex than the intramicroengine binders discussed above. In particular, the latter two require stubs and skeletons to map the parameter and return values to PCI bus packets. The microengine-side stubs/skeletons are hand coded rather than being generated automatically from IDL (this is another reason for severely constraining the interfaces of microengine components).

Finally, we have implemented a range of loaders for the XScale and PC caplet types, together with binders that perform intracaplet binding within these caplet types, and a binder that operates between the two. These are relatively straightforward and are implemented using delegators and delegates along the lines discussed in Section 4.4.1.

7.2.3 Overall Evaluation. This case study offers a glimpse of the abstraction power of the OpenCom programming model, and also exemplifies the use of domain-specific CFs and reflective metamodels to facilitate reconfiguration and ad hoc intervention at a high level of abstraction. Note particularly that the steps involved in the reconfiguration need not have been foreseen when the initial configuration was defined, and that they are entirely separate from the basic functionality of the components involved.

The case study is also useful in clarifying the benefits of the component-based programming model in “taming” hostile and heterogeneous deployment environments such as IXP NPs. In particular, the following can be seen:

- (1) The system programmer is completely shielded from the complex, low-level, heterogeneous and distributed nature of the NP-based router. Instead, she sees only components, loaders, binders, and caplets—all of which are operated on using standard and generic APIs. Even the presence of multiple caplets can be made transparent if the environment programmer has provided suitable metaloaders/binders.
- (2) Nevertheless, the implementation incurs no inherent in-band performance costs, as the component styles used are optimally tailored to their particular environments (e.g., the “basic” microengine component style is based directly on that provided by Intel⁶), and the binders provided by the environment programmer employ optimal, IXP-specific mechanisms such as address-patching and the use of nearest-neighbor registers.

⁶Obviously there is an additional overhead involved in the use of the “advanced” multiply instantiable component style; but given the extra functionality, this is to be expected. The key point is that if one compares like with like—that is, if one uses singleton components—the overhead is no greater than in the Intel/Netbind environments.

Table III. Throughput of an OpenCom-Based Network Bridge

| Number of Components | “Basic” binder/ Component Style | “Advanced” Binder/ Component Style |
|----------------------|---------------------------------|------------------------------------|
| 2 (sender/receiver) | 646.34 Mb/s | 634.52 Mb/s |
| 5 (i.e., 3 null) | 636.51 Mb/s | 614.51 Mb/s |
| 10 (i.e., 8 null) | 631.68 Mb/s | 575.23 Mb/s |
| 20 (i.e., 18 null) | 613.90 Mb/s | 442.10 Mb/s |

This case study also illustrates how the third-party nature of the programming model supports the control and management of the software in the primitive microengine environment. For example, from the primary caplet on the XScale one can incrementally load/unload microengine-based and PC-based components in a completely uniform manner. Similarly, bindings can be uniformly created between a microengine component and another component located anywhere else on the router. The uniformity is most clearly seen when we consider that a microengine component can equally straightforwardly initiate the loading and binding of components on the XScale or the host PC. Furthermore, the reflective metamodels apply uniformly to components in all three caplet types.

To validate the performance of OpenCom in the NP context, we compared the throughput of our Network Element CF with that of standard monolithic software. In both cases an IXP2400 was configured with a simple “bridging” program that ran on a single microengine and transferred packets between two 2.5-Gb/s fiber ports. In the “standard monolithic” case, the bridge was based on IXP2400 code released by Intel. In our case, the same code was refactored as two pipelined components: a receiver and a sender. We used the “basic” component style. Subsequently, we experimented with adding additional “null” components between the receiver and sender, and with using both of the component styles and binding types referred to above (i.e., the “basic” and “advanced” types). The throughput in the Intel case was 646.34 Mb/s. The results for the OpenCom case are shown in Table III.

Note that the OpenCom case with the two components and the “basic” binder adds *zero* overhead. Inevitably, overhead is incurred where additional components and/or the “advanced” component style and binding type are deployed, but this appears to be well within tolerable bounds given the additional functionality. In terms of memory use, the per-component overhead of the null components was as follows: basic component style: 40 b (i.e., one microengine instruction word); advanced component style: 120 b (i.e., three instruction words). Overall, the results show that it is perfectly feasible in performance terms to use the OpenCom approach to build low-level embedded software such as NP software.

7.3 Case Study 2: Reflective Middleware

7.3.1 Introduction. Middleware is distributed systems software that sits between (distributed) applications and an underlying network and end-system infrastructure. Its role is to shield applications from the complexity and heterogeneity of the underlying infrastructure by providing them with a distributed virtual machine. Due to diversification of both applications and infrastructures,

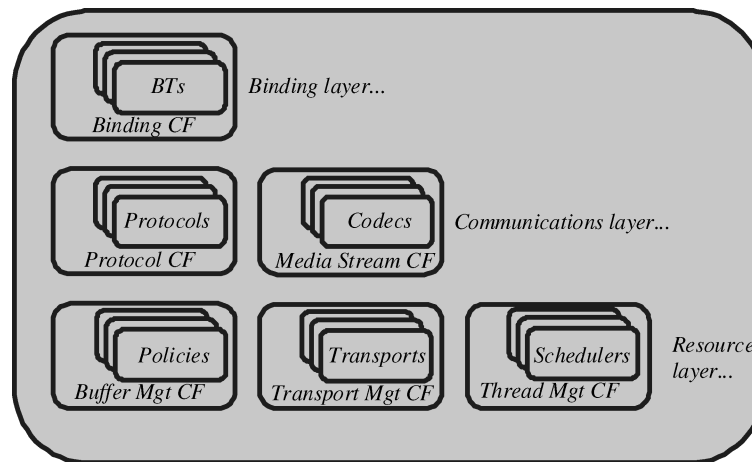


Fig. 7. The generic OpenORB architecture.

recent years have seen an explosion in the types of middleware being employed. Examples include traditional object-based middleware such as CORBA; component-based middleware such as Enterprise JavaBeans or the CORBA Component Model; Web Services and Grid middleware; asynchronous middleware based on events, tuples, or publish/subscribe; adaptive middleware for distributed real-time systems or mobile computing; embedded systems middleware; and sensor network middleware. This diversity makes it increasingly attractive to approach middleware construction using a framework approach in which tailored platforms can be constructed and customised in terms of a reusable and extensible set of components and CFs. We have built such a framework, called *OpenORB* [Coulson et al. 2002b; Parlavantzas 2005], and have used it to instantiate a number of different middleware platforms as discussed in this section.

7.3.2 Programmer Role Considerations. The system programmer is responsible for creating the OpenORB framework and populating it with a suitable set of CFs that can be combined by the user to build a variety of configurable and runtime reconfigurable middleware platforms. OpenORB (see Figure 7) is structured as a top-level CF (or *Top*) that is itself composed of three layers called the *Resource Layer*, the *Communications Layer*, and the *Binding Layer*. Each of these layers can itself contain a potentially extensible set of second-level CFs. The Top CF enforces the three-layer structure by ensuring that each plugged-in CF (and its components) only has access to interfaces offered by components in the same or lower layers. Furthermore, it imposes policies concerning layer composition and dynamic changes in layer composition. The plugged-in CFs address more focused subdomains of middleware functionality (e.g., binding establishment and thread management), and enforce appropriate subdomain-specific policies.

The Resource Layer CFs, respectively, handle buffers (the Buffer Management CF takes pluggable buffer allocation policies), transport connections (the

Transport Management CF takes pluggable transport protocols), and thread management (the Thread Management CF takes pluggable user-level thread scheduling policies). Next the Communications Layer contains Protocol and Media Stream CFs. The former accepts plug-in protocol components (e.g., GIOP or SOAP), which can be organized into stacks, and the latter accepts software codecs. Finally, the Binding Layer contains a Binding CF [Parlavantzas et al. 2003] that accepts *binding-type* (BT) implementations (e.g., remote object invocation, streaming connections, publish/subscribe, etc.). The Binding CF is the most complex of the CFs and is a crucial part of the OpenORB architecture because it determines the programming model offered to middleware users (e.g., standard CORBA or Web Services).

The OpenORB architecture fosters a considerable degree of component sharing across multiple middleware platforms. In particular, the Resource Layer provides buffer management, transport, and thread scheduling components that can potentially be used by many platforms including diverse ones such as CORBA, Web Services, and event-based middleware. The layer composition policies in the Top CF determine exactly which plug-ins are to be applied. The typical level of reuse, however, tends to diminish in the higher layers.

Overall, the OpenORB implementation consists of about 50,000 lines of C++ divided into five CFs and around 30 components. We have used this to build a standard CORBA environment [Coulson et al. 2002b], a platform for mobile computing [Grace et al. 2003] (which employs a Web Services API and a Dynamic Service Discovery CF in the binding layer) and a platform for grid computing [Grace et al. 2005]. Full details of the overall OpenORB approach are available in Parlavantzas [2005].

Turning now to the *environment programmer* role, unlike the programmable networking case study, OpenORB does not need to employ a large and complex set of caplets, loaders, and binders. This is because (with the exception of embedded and sensor network middleware) the underlying infrastructure tends to be accessed via relatively high-level OS APIs. However, this case study does require the OpenCom kernel to be ported to a range of end systems running a range of OSs. Because of the size and simplicity of the kernel, this has proved straightforward: we have easily ported the kernel to a range of PCs and PDAs running both Windows (including CE) and a variety of Unix flavors. The main requirement in terms of caplets has been to support both native and interpreted component styles. Therefore, we have Java-based caplets to go with native kernels and vice versa. We employ straightforward OS mechanisms (e.g., Java class loaders and IPC mechanisms such as pipes) to implement suitable loaders and cross-caplet binders that work within a single machine environment. We then build intermachine communication in terms of middleware-level communications CFs that include RPC protocols etc.

7.3.3 Overall Evaluation. Unlike the programmable networking case study, the OpenORB case study emphasises the system programmer role more than the environment role; it demonstrates the practicability of using OpenCom to develop a large and complex system which has evolved over a number of years and been worked on by a significant number of people. It




Fig. 8. Performance of OpenORB versus GOPI and Orbacus.

also validates the usefulness and practicality of the CF approach to building systems.

In terms of overhead, we have compared the performance of our OpenORB-based CORBA implementation with standard CORBA object request broker (ORB) implementations. The objective here is to evaluate the degree to which componentization impacts performance. The ORBs chosen for comparison were Orbacus 3.3.4 and GOPI v1.2 (GOPI [Coulson et al. 2002a] is a noncomponent-based CORBA implementation that we previously developed). All three ORBs employed CORBA GIOP v1.2. Orbacus is well known as one of the fastest and most mature CORBA-compliant commercial ORBs available and therefore represents the state of the art in CORBA performance. A comparison with GOPI, on the other hand, yields insight into the overhead of componentization as the GOPI code was extensively reused in our later component-based OpenORB implementation.

All three ORBs were tested on a Dell Precision 410MT workstation equipped with 256 MB RAM and an Intel Pentium III processor rated at 550 MHz. The operating system used was Microsoft XP and the compiler was Microsoft's cl.exe version 12.00.8804 with flags /MD/W3/GX/FD/O2. Our tests measured method invocations per second over the PC's loopback interface. A minimal IDL interface was employed that supported a single operation that took as its argument an array of octets and returned a different array of the same size. The implementation of this method at the server side was null.

The results of timing a large number of round-trip invocations using the above setup are shown in Figure 8. It can be seen that, for packets smaller than

1024 octets, OpenORB performs about the same as Orbacus, with GOPI running around 10% faster. As might be expected, there is a diminishing difference between all three systems as packet size increases; this is presumably due to the fact that the overhead of data copying begins to outweigh the cost of call processing. The relative overhead of OpenORB compared to GOPI can be attributed to the former's increased use of indirection (i.e., through bindings). In our test configuration, the data path for each GIOP invocation involved 67 bindings, 32 on the client side and 35 on the server side. Despite this additional work, it can be seen that the performance of OpenORB is entirely comparable to that of the noncomponentized ORBs. This demonstrates that structuring middleware platforms in terms of dynamically composable components do not necessarily incur a significant overhead.

Finally, based on our experience with both OpenORB and the noncomponent-based GOPI, we briefly comment on the relative development benefits of using a generic component model to build middleware. Although GOPI was already structured in a “modular” fashion, the informality of the module interfaces was, as usual, a limiting factor in fostering reuse and independent development. In contrast, we found that the greater rigor imposed by OpenCom—and especially by the use of receptacles and language-independent interfaces—considerably enhanced these factors as well as leading to, if anything, faster development time. In addition, we automatically benefited from the ability to mix programming languages, as well as to reconfigure component configurations dynamically.

8. RELATED WORK

OpenCom can usefully be positioned against three areas of related work: (i) application-level component models, (ii) systems-level component models that are narrowly targeted at specific target domains and/or deployment environments, and (iii) generic systems-level component models.

In terms of a comparison with application-level component models, OpenCom substantially differs from designs such as EJB [Sun Microsystem 2005], the CORBA Component Model [Object Management Group 1999], Microsoft's COM or .NET [Microsoft 2005], and ICENI [Furmento et al. 2002] in being far more lightweight. OpenCom's capsule concept is superficially related to the “container” concept espoused by many of these models, but the latter is far more complex and provides correspondingly richer functionality—for example: in terms of policy specification for security, event handling, transactions, and persistence. In contrast, OpenCom capsules—due to our goals of genericity, resource parsimony, and performance—are policy free and embody only minimal, low-level functionality (i.e., they are loading, binding, and registry related). OpenCom shares with these models an emphasis on third-party deployability of components. However, whereas software reuse is the primary reason for most of these systems to espouse third-party deployability, the main reason for third-party deployability in OpenCom is to facilitate system configurability and reconfigurability, and to enable primitive components in extension caplets to function as first-class players in the computational model (cf. the programmable

networking case study). OpenCom also differs from EJB and ICENI in being language independent.

We now turn to component models which aspire to support the development of low-level systems, but which cannot reasonably be designated as *generic*. Here, we restrict our remarks primarily to arguing why these systems cannot be so designated. In the *embedded systems* area, we have a wide range of component models such as Pebble [Magoutis et al. 2000], PECOS [Winter et al. 2002], PBO [Stewart et al. 1993], SaveCCM [Hansson et al. 2004], and Koala [Van Ommering et al. 2000]. Most of these are build-time-only technologies—components are not visible at runtime and therefore these systems do not support dynamic reconfiguration. One area that some of these systems (i.e., PECOS and PECT) *do* support, however, that OpenCom does not natively support, is the specification (at build-time) of real-time constraints such as cycle time or worst-case execution time. Such facilities are clearly important in certain real-time-critical areas. Our approach to providing such facilities, where needed, would be to provide a suitable “real-time systems” CF rather than building-in real-time properties into the programming model itself. A further observation is that many of these embedded systems technologies (e.g., PBO, SaveCCM, and Koala) are tightly coupled to a specific underlying OS environment and/or are programming language specific.

In the *OS building* area, we have systems such as THINK [Fassino et al. 2002], OSKit [Ford et al. 1997], and MMLITE [Helander and Forin 1998]. THINK has a component-based programming model similar to OpenCom and also employs the notion of plug-in bindings. However, it does not share OpenCom’s “kernel-style” architecture and it lacks any equivalent of OpenCom’s other platform extension facilities (i.e., caplets and loaders) and so is unable to abstract over as broad a range of deployment environments. It also lacks OpenCom’s reflective capabilities and has so far only been used to develop OSs in standard PC environments. OSKit is a Knit-based toolkit that consists specifically of OS-related components. Thus it is Knit (see below) rather than OSKit itself that is the “generic” model. MMLITE was an early attempt to adapt and apply Microsoft’s COM as a vehicle for building operating systems. Although it did support a degree of reconfiguration, MMLITE had no specific support (e.g., in terms of reflection and CFs) to control and manage this; in addition, due to its reliance on COM, it was not suitable for primitive deployment environments.

In the *programmable networking environments* area, we have a wide range of component systems that might be exemplified by VERA [Karlin and Peterson 2001], MicroACE [Johnson and Kunze 2003], and Netbind [Campbell et al. 2002]. Many of these systems are targeted at primitive deployment environments; but they largely achieve this by defining very limited and specific component models. VERA, for example, is more a component framework than a generic component model in that it supports plug-in components only at predesignated “hot spots.” Similarly, MicroACE is specifically designed for Intel IXP network processors—its notion of “component” is inherently “split” so that a part of a component exists on a microengine and part on the control processor. Apart from the fact that MicroACE does not support runtime reconfiguration, this architecture is clearly inapplicable to other types of deployment environment.

Again, Netbind, on which we have built our own programmable networking research (see Section 7.2), is barely a component model—it deals with informal “modules” rather than components and its binding mechanism is intimately tied to the IXP deployment environment.

Finally, in the *middleware platforms* area, we have systems such as LegORB [Roman et al. 2000], k-Components [Dowling and Cahill 2001], and various JavaBeans-based approaches (e.g., Bruneton and Riveill [2000]; Joergensen et al. [2000]). These systems tend to be flexible in terms of their support for dynamic reconfiguration, but none of them are suitable for deployment environments other than standard PC environments (or at best PDAs running a standard OS such as Windows CE).

Turning now to purportedly generic systems-level component models, we observe that as mentioned in the introduction, only a limited amount of work has so far been carried out on exploiting the component paradigm as a generic systems-building approach. The other main players in the field are Knit [Reid et al. 2000] and Fractal [Bruneton et al. 2004]. Knit was initially targeted primarily at operating systems (e.g., the OSKit system mentioned above), but it has also been successfully used to build software routers, although only in conventional PC environments. The main limitation of Knit is that it addresses purely *build-time* concerns: the component model is not visible at runtime, so there is no systematic support for dynamic component loading, still less managed reconfiguration. Fractal is much closer to OpenCom in its goals and approach. Like OpenCom, it supports bindings between interfaces as first-class objects, and it takes seriously the need for dynamic reconfiguration. However, although Fractal is designed to be abstract and generic, it seems to lack some of the flexibility of OpenCom, which potentially limits its applicability. In particular, it mandates a specific component style that provides quite rich (and therefore costly) behavior in several areas—that is, it supports composite components and has an architecture that describes how components can be built in terms of multiple Java classes (or C files), and its components need to include a “controller” part for lifecycle management. It is difficult to see how this relatively complex component style could, for example, be deployed in very primitive environments like network processor microengines or sensor nodes. Like OpenCom, Fractal supports reflection. This is primarily architectural reflection which employs an XML-based architectural description language that is available at runtime. OpenCom, on the other hand, supports an extensible range of reflective facilities, and leaves the choice of specific facilities to the environment programmer.

Finally, it is instructive to briefly compare OpenCom’s kernel with the common notion of OS microkernels (see, e.g., Rashid et al. [1989]). While there is an apparent similarity between the architecture diagram in Figure 2 and that of a typical microkernel OS, it can easily be shown that the resemblance is only superficial: microkernels are responsible for such in-band tasks as interrupt handling, thread scheduling, message passing, and paging. The OpenCom kernel, on the other hand, only supports simple compositional functionality—that is, loading and binding. The distinction is perhaps most clearly evident when one considers that it is straightforward to unload the kernel itself to save space (as long as no subsequent system reconfiguration is required).

9. CONCLUSIONS

We have presented a generic component-based approach to the construction of systems software, and have argued that this generic approach has significant advantages over the current crop of narrowly targeted systems-building components models. In particular, we argue that our approach maximizes the genericity and abstraction level of the component-based programming model while at the same time supporting a principled approach to the support of a wide range of target domains and deployment environments. OpenCom offers a flexible, extensible, and language-independent architecture that is based on a minimal component runtime kernel. Due to its simplicity, the OpenCom kernel can easily be deployed in a wide range of deployment environments and used to underpin the construction of both dynamic and static systems (in the latter case the kernel can, as mentioned above, be viewed simply as a highly flexible and configurable loader that is discarded once the system has booted).

On top of the kernel, the platform extensions layer adds the generic abstractions of caplets, loaders, and binders. These are instrumental in adapting OpenCom to heterogeneous deployment environments and in “taming” the idiosyncrasies of the deployment environment without overabstraction (which would lead to poor performance and lack of accessibility to useful deployment-environment-specific features). As shown in Section 7, a wide range of deployment environment idiosyncrasies can be exposed in a consistent and efficient manner through these abstractions. The result is a standard and general approach to system construction and reconfiguration based on a uniform model of third-party loading and binding (this is particularly well illustrated by the programmable networking case study).

Another important element of our platform extensions design is the notion of differentiated programmer roles: environment programmers populate a given deployment environment with suitable caplets, loaders, and binders; and system programmers construct the target system in terms of these. This differentiation structures the development of systems and provides a principled approach to the task of bridging the “implementation gap” between an abstract programming model and a concrete deployment environment. Also in the extensions layer, the reflective metamodels provide principled means of reconfiguring systems in terms of inspection, adaptation, and extension. Again, it is for the environment programmer to choose which of an (extensible) set of reflective metamodels should be made available in any given installation.

Component styles and component frameworks are two further features of our approach that strongly foster genericity. The ability (through caplets) to support an extensible set of component styles means that the component notion can be instantiated in almost any conceivable programming language and deployment environment (including microengines and sensor motes). Nevertheless, thanks to third-party loading and binding, the idiosyncrasies of different component styles are only visible to the developers who *write* components using some particular style. Other developers who *compose* (and reconfigure) systems in terms of these components can choose to remain oblivious to their internal heterogeneity. Our architecture is similarly agnostic in the area of component frameworks. Unlike, say, Fractal, which supports a particular instantiation of

the CF concept (which employs an associated XML-based architecture description language), OpenCom fosters generality by leaving the CF notion loosely specified as an architectural “pattern.” Nevertheless, this does not preclude the use of specific “metaframeworks” that can be used as optional extensions to generate CFs. For example, we have developed a metaframework called “*Plastik*” that employs an ADL to formally specify constraints on OpenCom CFs and automatically generates checking code to police these constraints at runtime [Joolia et al. 2005].

In our present and future work, we are continuing both to develop OpenCom and to evaluate it by building target systems in a wide variety of target domains and deployment environments. For example, a current EU-funded collaborative project is focusing on the use of an OpenCom-like component model in the development of real-time and embedded systems [Costa et al. 2007]. This project is also investigating the provision of fundamental support for security in OpenCom-like systems. The approach we are taking here is to provide a “security mediator” in the extensions layer that offers a minimal Trusted Computing Base, and then to provide higher-level mechanisms and policies in terms of CFs. We are also working on an extension of the above-mentioned Plastik metaframework CF for the specification and policing of real-time properties. Finally, we are investigating the potential of aspect-oriented systems development techniques in OpenCom-based real-time embedded environments as a more design-oriented approach to complex system configuration and reconfiguration.

ACKNOWLEDGMENTS

The authors would like to acknowledge the UK EPSRC for funding the vast majority of this research. We would also like to gratefully acknowledge our students over the past few years who have contributed in numerous ways: in particular, Nelly Bencomo, Michael Clarke, Anita Cleetus, Nikos Parlavantzas, Rajiv Ramdhany, and Nirmal Weerasinghe. Finally, we would like to acknowledge our colleagues in the EU-funded RUNES project (IST-004536) who contributed ideas in the later stages of the development of OpenCom. RUNES is currently developing a component-based middleware infrastructure for networked embedded systems that picks up and develops many of the concepts pioneered in OpenCom.

The OpenCom software is available online for download at <http://www.comp.lancs.ac.uk/computing/research/mpg/reflection/download.php>.

REFERENCES

- BENCOMO, N., BLAIR, G., COULSON, G., GRACE, P., AND RASHID, A. 2005. Reflection and aspects meet again: Runtime reflective mechanisms for dynamic aspects. In *Proceedings of the First Middleware '05 Workshop on Aspect Oriented Middleware Development (AOMD 05, Grenoble, France)*.
- BLAIR, G., COULSON, G., AND GRACE, P. 2004. Research directions in reflective middleware: The Lancaster Experience. In *Proceedings of the 3rd Workshop on Reflective and Adaptive Middleware (RM2004; co-located with Middleware 2004, Toronto, Ontario, Canada)*.
- BRUNETON, E., COUPAYE, T., LECLERC, M., QUEMA, V., AND STEFANI, J. B. 2004. An open component model and its support in Java. In *Proceedings of the 7th International Symposium on Component-Based Software Engineering (ICSE-CBSE7, Edinburgh, Scotland)*.
- BRUNETON, E. AND RIVEILL, M. 2000. JavaPod: An adaptable and extensible component platform. *Proceedings of the Reflective Middleware 2000*, (New York, NY).

- CAMPBELL, A. T., CHOU, S., KOUNAVIS, M.E., STACHTOS, V. D., AND VICENTE, J. B. 2002. NetBind: A binding tool for constructing data paths in network processor-based routers. In *Proceedings of the 5th IEEE International Conference on Open Architectures and Network Programming (OPENARCH'02)*.
- CLARKE, M., BLAIR, G. S., AND COULSON, G. 2001. An efficient component model for the construction of adaptive middleware. In *Proceedings of the IFIP/ACM Middleware 2001* (Heidelberg, Germany).
- COSTA, P., COULSON, G., GOLD, R., LAD, M., MASCOLO, C., MOTTOLA, L., PICCO, G. P., SIVAHARAN, T., WEERASINGHE, N., AND ZACHARIADIS, S. 2007. The RUNES middleware for networked embedded systems and its application in a disaster management scenario. In *Proceedings of the 5th Annual IEEE International Conference on Pervasive Computing and Communications (PERCOM'07, White Plains, NY)*.
- COULSON, G., BAICHO, S., AND MOONIAN, O. 2002a. A retrospective on the design of the GOPI middleware platform. *ACM Multimed J.* 8, 3, 340–352.
- COULSON, G., BLAIR, G. S., CLARKE, M., AND PARLAVANTZAS, N. 2002. The design of a highly configurable and reconfigurable middleware platform. *ACM Distrib. Comput. J.* 15 2, 109–126.
- COULSON, G., BLAIR, G. S., AND GRACE, P. 2004. On the performance of reflective systems software. In *Proceedings of the International Workshop on Middleware Performance (MP 2004, Phoenix, AZ)*; satellite workshop of the IEEE International Performance, Computing and Communications Conference (IPCCC 2004).
- COULSON, G., BLAIR, G. S., HUTCHISON, D., JOOLIA, A., LEE, K., UHEYAMA, J., GOMES, A.T., AND YE, Y. 2003. NETKIT: A software component-based approach to programmable networking. *ACM SIGCOMM Comput. Communi. Rev.* 33, 5, 55–66.
- DOWLING, J. AND CAHILL, V. 2001. The k-component architecture meta-model for self-adaptive software. In *Proceedings of Reflection 2001*. Lecture Notes in Computer Science, vol. 2192. Springer, Berlin, Germany.
- EMMERICH, W. 2002. Distributed component technologies and their software engineering implications. In *Proceedings of the 24th International Conference on Software Engineering* (Orlando, FL). 537–546.
- FASSINO, J.-P., STEFANI, J.-B., LAWALL, J., AND MULLER, G. 2002. THINK: A software framework for component-based operating system kernels. In *Proceedings of the Usenix Annual Technical Conference* (Monterey, CA).
- FORD, B., BACK, G., BENSON, G., LEPREAU, J., LIN, A., AND SHIVERS, O. 1997. The Flux OSKit: A substrate for kernel and language research. In *Proceedings of the 16th ACM Symposium on Operating Systems Principles* (Saint Malo, France). ACM Press, New York, NY, 38–51.
- FURMENTO, N., MAYER, A., MCGOUGH, S., NEWHOUSE, S., FIELD, T., AND DARLINGTON, J. 2002. ICENI: Optimisation of component applications within a grid environment. *J. Parallel Comput.* 28, 12, 1753–1772.
- GAMMA, E., HELM, R., JOHNSON, R., AND VISSIDES, J. 2004. Design patterns: Elements of reusable object-oriented software. Addison-Wesley, Reading, MA.
- GARLAN, D., MONROE, R. T., AND WILE, D. 2000. Acme: Architectural description of component-based systems. In *Foundations of Component-Based Systems*, G. T. Leavens and M. Sitaraman (Eds), Cambridge University Press, Cambridge, U.K. 47–68.
- GRACE, P., BLAIR, G. S., AND SAMUEL, S. 2003. ReMMoC: A reflective middleware to support mobile client interoperability. In *Proceedings of the Symposium on Distributed Objects and Applications (DOA 2003, Catania, Sicily, Italy)*.
- GRACE, P., COULSON, G., BLAIR, G. S., AND PORTER, B. 2005. Deep middleware for the divergent grid. In *Proceedings of IFIP/ACM/USENIX Middleware 2005*.
- HANSSON, H., AKERHOLM, M., CRNKOVIC, I., AND TORNGREN, M. 2004. SaveCCM—a component model for safety-critical real-time systems. In *Proceedings of the IEEE Euromicro Conference* (Special Session on Component Models for Dependable Systems).
- HELANDER, J. AND FORIN, A. 1998. MMLite: A highly componentized system architecture. In *Proceedings of the 8th ACM SIGOPS European Workshop* (Sintra, Portugal). 96–103.
- HUGHES, D., GREENWOOD, P., BLAIR, G., COULSON, G., PAPPENBERGER, F., SMITH, P., AND BEVEN, K. 2006. An intelligent and adaptable grid-based flood monitoring and warning system. In *Proceedings of the UK eScience All Hands Meeting*.

- INTEL. 2004. Intel IXP1200/2400 Network Processors. Go online to <http://www.intel.com/IXA>.
- JOERGENSEN, B. N., TRUYEN, E., MATTHIJS, F., AND JOOSEN, W. 2000. Customization of object request brokers by application specific policies. In *Proceedings of IFIP Middleware 2000* (New York, NY).
- JOHNSON, E. J. AND KUNZE, A. R. 2003. IXP2400/2800 programming: The complete microengine coding guide. Intel Press, Santa Ceara, CA.
- JOOLIA, A., BATISTA, T., COULSON, G., AND TADEU A. 2005. Mapping ADL specifications to an efficient and reconfigurable runtime component platform. In *Proceedings of the 5th Working IEEE/IFIP Conference of Software Architecture* (WICSA 5, Pittsburgh, PA).
- KARLIN, S. AND PETERSON, L. 2001. VERA: An extensible router architecture. In *Proceedings of the IEEE Conference on Open Architectures and Network Programming* (OPENARCH 2001, Anchorage, AK). 3–14.
- KON, F., COSTA, F., CAMPBELL, R., AND BLAIR, G. 2002. The case for reflective middleware. *Communi. ACM* 45, 6, 33–38.
- MAES, P. 1987. Concepts and experiments in computational reflection. In *Proceedings of OOP-SLA'87 ACM SIGPLAN Not.* 22, 147–155.
- MAGOUTIS, K., BRUSTOLONI, J. C., GABBER, E., NG, W. T., AND SILBERSCHATZ, A. 2000. Building appliances out of reusable components using pebble. In *Proceedings of the SIGOPS European Workshop 2000* (Kolding, Denmark), 211–216.
- MICROSOFT. 2005. Net Home Page. Go online to <http://www.microsoft.com/net>.
- MOZILLA ORGANIZATION. 2005. XPCOM project. Go online to <http://www.mozilla.org/projects/xpcom>.
- NETWORK PROCESSING FORUM. 2005. Go online to <http://www.npforum.org/>.
- OBJECT MANAGEMENT GROUP. 1995. *The Common Object Request Broker: Architecture and Specification*, 2.0 ed. MG, Needham, MA.
- OBJECT MANAGEMENT GROUP. 1999. CORBA components final submission. OMG Document orbos/99-02-05. OMG, Needham, MA.
- PARLAVANTZAS, N. 2005. Constructing modifiable middleware with component frameworks. Ph.D. dissertation. Computing Department, Lancaster University, Lancaster, U.K.
- PARLAVANTZAS, N., COULSON, G., AND BLAIR, G.S. 2003. An extensible binding framework for component-based middleware. In *Proceedings of the 7th IEEE International Enterprise Distributed Object Computing Conference* (EDOC 2003, Brisbane, Australia).
- RASHID, R., BARON, B., FORIN, A., GOLUB, D., JONES, M., JULIN, D., ORR, D., AND SANZI, R. 1989. Mach: A foundation for open systems. In *Proceedings of the 2nd Workshop on Workstation Operating Systems* (WWOS2).
- REID, A., FLATT, M., STOLLER, L., LEPREAU, J., AND EIDE, E. 2000. Knit: Component composition for systems software. In *Proceedings of the 4th Conference on Operating Systems Design and Implementation* (OSDI). 47–360.
- ROMAN, M., MICKUNAS, D., KON, F., AND CAMPBELL, R. H. 2000. LegORB. In *Proceedings of the IFIP/ACM Middleware 2000 Workshop on Reflective Middleware*. IBM Palisades Executive Conference Center, Yorktown Heights, NY.
- STEWART, D., VOLPE, R., AND KHOSLA, P. 1993. Design of dynamically reconfigurable real-time software using port-based objects. Robotics Institute, Rep. No. CMU-RI-TR-93-11. Carnegie Mellon University, Pittsburgh, PA.
- SUN MICROSYSTEMS. 2005. Go online to <http://java.sun.com/products/ejf/index.html>.
- SZYPERSKI, C. 1998. *Component Software: Beyond Object-Oriented Programming*. Addison-Wesley, Reading, MA.
- VAN OMMERING, R., VAN DER LINDEN, F., KRAMER, J., AND MAGEE, J. 2000. The Koala component model for consumer electronics software. *IEEE Comput.* 33, 3, 78–85.
- WINTER, M., GENBLER, T., CHRISTOPH, A., NIERSTRASZ, O., DUCASSE, S., WUYTS, R., AREVALO, G., MULLER, P., STICH, C., AND SCHONHAGE, B. 2002. Components for embedded software: the PECOS approach. In *Proceedings of the 2002 International Conference on Compilers, Architecture, and Synthesis for Embedded Systems* (CASES'02, Grenoble, France). ACM Press, New York, NY, 19–26.

Received January 2006; revised March 2007, November 2007; accepted November 2007