

Discovering Architectures from Running Systems

Bradley Schmerl, Jonathan Aldrich, David Garlan, Rick Kazman, and Hong Yan

Abstract—One of the challenging problems for software developers is guaranteeing that a system as built is consistent with its architectural design. In this paper, we describe a technique that uses runtime observations about an executing system to construct an architectural view of the system. In this technique, we develop mappings that exploit regularities in system implementation and architectural style. These mappings describe how low-level system events can be interpreted as more abstract architectural operations and are formally defined using Colored Petri Nets. In this paper, we describe a system, called DiscoTect, that uses these mappings and we introduce the DiscoSTEP mapping language and its formal definition. Two case studies showing the application of DiscoTect suggest that the tool is practical to apply to legacy systems and can dynamically verify conformance to a preexisting architectural specification.

Index Terms—Software architecture discovery, reverse engineering, architecture design tools and analyses.

1 INTRODUCTION

A well-defined software architecture is critical for the success of complex software systems. An architectural model consists of many views (e.g., [24]) of the system. One particularly useful view is the component and connector (C&C) view, which provides a high-level view of a system in terms of its principle runtime components (e.g., clients, servers, databases), their interactions (e.g., remote procedure call, event multicast, piped streams), and their properties (e.g., throughputs, latencies, reliabilities) [5], [31], [35]. As an abstract representation of a system, an architecture permits many forms of high-level inspection and analysis, allowing the architect to determine if a system's design will satisfy its critical quality attributes. Consequently, over the past decade, considerable research and development has gone into the development of notations, tools, and methods to support architectural design [7], [8], [27].

However, a persistent problem is determining whether a system as implemented has the architecture with which it was designed. Without some form of conformance guarantee, the validity of any architectural analysis will be suspect at best and completely erroneous at worst.

Two tool-based techniques have been used to determine or enforce relationships between a system's C&C software architecture and its implementation. The first is to ensure consistency by *construction*. This can be done by embedding

architectural constructs in an implementation language (e.g., as described by Aldrich et al. [2]), where analysis tools can check for conformance. Or, it can be done through *code generation*, using tools to create an implementation from a more abstract architectural definition [34], [37], [38].

Ensuring consistency by construction is effective when it can be applied since tools can guarantee conformance; unfortunately, it has limited applicability. In particular, it can usually be applied only in situations where engineers are required to use specific architecture-based development tools, languages, and implementation strategies. For systems that are composed of existing parts or that require a style of architecture or implementation outside those supported by generation tools, this approach does not apply.

The second technique is to ensure conformance by extracting an architecture from a system's code using static code analysis [17], [21], [28]. When an implementation is sufficiently constrained so that modularization and coding patterns can be identified with architectural elements, this technique can work well. Unfortunately, however, the technique is limited by an inherent mismatch between static, code-based structures (such as classes and packages) and the runtime structures that are the essence of most architectural descriptions [8], [14]. In particular, the actual runtime structures may not even be known until the program executes: Clients and servers may come and go dynamically, components (e.g., Dynamic Linked Libraries) not under direct control of the implementers may be dynamically loaded, and so forth. Indeed, determining the actual runtime architectural configuration of a system using static analysis is, in general, undecidable.

A third, relatively unexplored, technique is to determine the architecture of a system by *examining its runtime behavior*. The key idea here is that a system's execution can be monitored. Observations about its runtime behavior can then, in principal, be used to infer its dynamic architecture. This approach has the advantages that it applies to *any*

• B. Schmerl, J. Aldrich, D. Garlan, and H. Yan are with the School of Computer Science, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15221.

E-mail: {schmerl, aldrich, garlan, yh}@cs.cmu.edu.
 • R. Kazman is with the University of Hawaii, 2404 Maile Way, Honolulu, HI, and the Software Engineering Institute, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15221. E-mail: kazman@sei.cmu.edu.

Manuscript received 6 May 2005; revised 1 Mar. 2006; accepted 11 July 2006; published online 9 Aug. 2006.

Recommended for acceptance by S. Uchitel, I.H. Krüger, M. Broy, and J. Whittle.

For information on obtaining reprints of this article, please send e-mail to: tse@computer.org, and reference IEEECS Log Number TSE-0143-0505.

system that can be monitored, it gives an accurate image of what is actually going on in the real system, it can accommodate systems whose architecture changes dynamically, and it imposes no a priori restrictions on system implementation or architectural style.

However, there are a number of hard technical challenges in making this technique work. The most serious problem is finding mechanisms to bridge the abstraction gap: In general, low-level system observations do not map directly to architectural constructs. For example, the creation of an architectural connector might involve many low-level steps and those actions might be temporally interleaved with many other architecturally relevant actions. Moreover, there is likely no single architectural interpretation that will apply to all systems: Different systems will use different runtime patterns to achieve the same architectural effect and, conversely, there are many possible architectural elements to which one might map the same low-level events.

In this paper, we describe a technique to solve the problem of dynamic architectural discovery for a large class of systems. Our key idea and contribution is to provide a framework that allows one to map implementation styles to architecture styles. This mapping is defined conceptually as a Colored Petri Net [19] that is used at runtime to track the progress of the system and output architectural events when predefined runtime patterns are recognized. Thus, the mapping provides a way to identify when a program performs “architecturally significant” actions that produce architectural structures. An important additional feature of the approach is the ability to reuse such mappings across systems. In particular, we exploit regularity in implementation and architectural styles so that a single mapping can serve as an architectural extractor for a large collection of similar systems, thereby reducing the cost of writing each abstraction mapping while still providing flexibility.

In the remainder of this paper, we describe this approach in detail and describe the tool called *DiscoTect* that we have implemented. Section 2 presents the *DiscoTect* approach, including an overview of the *DiscoTect* framework and the *DiscoSTEP* language used for specifying mappings. We then outline a formal semantics for *DiscoSTEP* that specifies its meaning in terms of Colored Petri Nets. We illustrate this with a simple example. In Section 3, we describe the implementation of *DiscoTect*. Sections 4 and 5 present case studies that we use to evaluate the efficacy of *DiscoTect*. In Section 6, we position our work in the context of related research. Finally, in Section 7, we present our conclusions and future work.

2 DiscoTect

2.1 Technical Challenges

Any approach that supports dynamic discovery of architectures must address three challenges:

1. **monitoring**: observing a system’s runtime behavior,
2. **mapping**: interpreting that runtime behavior in terms of architecturally meaningful events, and
3. **architecture building**: representing the resulting architecture.

In this paper, we are primarily concerned with the second problem of bridging the abstraction gap between system observations and architectural effects. There are a number of issues that make this a hard problem. First, mappings between low-level system observations and architectural events are not usually one-to-one. Many low-level events may be completely irrelevant. More importantly, a given abstract event, such as creating a new architectural connector, might involve many runtime events, such as object creation and lookup, library calls to runtime infrastructure, initialization of data structures, and so forth. Conversely, a single implementation event might represent a series of architectural events. For example, executing a procedure call between two objects might signal the creation of a new connector and its attachment to the runtime ports of the respective architectural components. This implies the need for a technique that can keep track of intermediate information about mappings to an architectural model.

Second, architecturally relevant actions are typically interleaved in an implementation. At any given moment, a system might be midway through creating several components and their connectors. This implies that any attempt to recognize architectural events must be able to cope with concurrent intermediate states.

Third, there is no single gold standard for indicating what implementation patterns represent specific architectural events. Different implementations may choose different techniques for creating the same abstract architectural element. Consider the number of ways that one might implement pipes, for example. Indeed, one might even find multiple implementation approaches in the same system. Moreover, there is no single architectural style or pattern that can be used for all systems. For example, the use of sockets might be used to represent many different types of connectors.

To handle this variability of implementation strategies and possible architectural styles of interest, a language is required to define new mappings. Given a set of implementation conventions (which we will refer to as an *implementation style*) and a vocabulary of architectural element types and operations (which we will refer to as an *architectural style* [12]), we require a description that captures the way in which runtime events should be interpreted as operations on elements of the architectural style. Thus, each pair of implementation style and architectural style has its own mapping. A significant consequence is that these mappings can be reused across programs that are implemented in the same style.

In summary, taking these challenges into account produces the following requirements for an approach to dynamically determine the architecture of a running system:

1. It must handle M-N mappings between low-level system events and constructs at the architectural level.
2. It must be able to keep track of concurrent and interleaved sets of event traces that produce architecturally significant events.
3. It needs to have flexibility in mapping implementation style to architectural style.

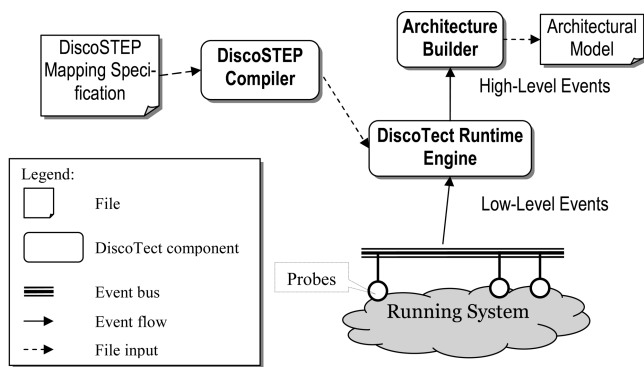


Fig. 1. The DiscoTect architecture.

2.2 The DiscoTect Approach

To address these concerns, we have adopted the approach illustrated in Fig. 1. Events captured from a running system are first filtered to select the subset of system observations that must be considered. The resulting stream of events is then fed to the DiscoTect Runtime Engine. The DiscoTect Engine takes in a specification of the mapping, written in a language called *DiscoSTEP* (Discovering Structure Through Event Processing). The DiscoTect engine constructs a Colored Petri Net from the mapping to recognize interleaved patterns of runtime events and, when appropriate, to produce a set of architectural operations as outputs. Those operations are fed to an Architecture Builder that incrementally creates an architectural model, which can then be displayed to a user or processed by architecture analysis tools. We now elaborate on each of the three main components in turn:

1. **DiscoSTEP Mapping Specification.** We have developed a language, called DiscoSTEP, for specifying mappings between low-level and architecture events. The execution semantics of these mappings is defined using Colored Petri Nets. We outline a translation from DiscoSTEP mappings to Colored Petri Nets in Section 2.4.
2. **DiscoTect Runtime Engine.** The runtime engine takes, as inputs, events from the running program and a DiscoSTEP specification and then runs the DiscoSTEP specification to produce architecture events. System runtime events are first intercepted and converted into XML (Extensible Markup Language) [45] streams by *Probes*. The resulting stream of events is then fed to the DiscoTect Runtime Engine, which uses the DiscoSTEP specification to recognize interleaved patterns of runtime events and, when appropriate, outputs a set of architectural events.
3. **Architecture Builder.** The architecture builder takes architectural events and incrementally constructs an architectural description, which can then be displayed to a user or processed by other architecture analysis tools.

2.3 Informal Introduction to DiscoSTEP

To handle the variability of implementation strategies and architectural styles of interest, we provide a language to

define mappings. DiscoSTEP specifies how to map system-level events to architectural-level events. To discover the architecture of a system, a program written in the DiscoSTEP language is compiled into byte code that can be interpreted by the DiscoTect engine. The DiscoTect engine processes the runtime events produced by the system and generates architectural events on the fly. The architectural events can be further consumed by an architectural builder to construct the architecture. A DiscoSTEP specification has three main ingredients:

1. **Events types.** The kinds of events, and their structure, that are consumed and produced by the Disco-STEP program. These are defined by an XML Schema.
2. **Rules.** Rules specify how to map a set of system-level events into architectural events and represent a single “step” in processing an event stream.
3. **Compositions.** Compositions of rules specify how output events from rules are passed as input events to other rules. Complex patterns of event sequences to be processed can be defined in this way.

In this section, we informally describe these three components. We then use a simple example to illustrate how they form a language to instruct event handling.

2.3.1 Events

In DiscoTect, we can capture *runtime events* such as method calls, CPU utilization, network bandwidth consumption, memory usage, etc. To generate these events at runtime, we probe, or instrument, the running system. To do this, we can use resource monitoring tools or code instrumentation tools such as AspectJ [23] and AspectC++ [36] that allow us to inject code into the target system. These events are input into a DiscoSTEP specification. The specification produces *architectural events*, generated as a result of processing the runtime events, which in turn are used to produce the software architecture.

Representation. We do not want the kinds of events consumed or produced by DiscoTect to be hardwired. We want the flexibility of customizing events for particular implementation and architectural styles. To do this, we use XML Schemas to specify the valid XML representations that can be used as events by DiscoSTEP and the events are therefore represented in XML. This allows us to perform structural type checking of the use of events in the DiscoSTEP specification. Fig. 2 illustrates the schema definition of a call type event and Fig. 3 provides an example of a valid XML event that conforms to this schema. In this paper, we use Acme representations for architectures. However, Acme is by no means the only architectural representation that could be supported by DiscoSTEP because we use an XML schema to define the output.

Input and Output. Our event processing engine takes runtime events as input and produces architectural events as output. For clarity, and for the purposes of checking the correctness of a DiscoSTEP specification, we partition DiscoSTEP event types into *inputs* and *outputs*. For example, the call type event described above would be specified as an input, whereas an event like “create_component” would be specified as an output event.

```

<element name="call">
  <complexType>
    <attribute name="method_name" type="string" />
    <attribute name="callee_id" type="string" />
    <attribute name="return_id" type="string" />
  </complexType>
</element>

```

Fig. 2. Defining even schema for use in DiscoSTEP.

2.3.2 Rules and Compositions

Rules determine what events should be processed and generate new events that are either used to construct the architecture or are fed into other rules to allow recognition of complex implementation patterns. Rules are composed of *inputs*, *outputs*, *triggers*, and *actions*. Input and output blocks declare input events that a rule cares about and the output events that a rule can generate. Triggers are predicates over the input events. Actions are assignments to the output events. When a trigger returns true upon the arrival of input events, the actions in the corresponding rule are activated to instantiate the output events. Since events are represented in XML, we use a well-defined XML Query language called XQuery [44] to describe triggers and actions. This allows us to specify a large range of predicates and manipulations over XML, rather than inventing a new syntax.

A rule only represents a fragment of what is needed to capture an architecture. Multiple rules can be assembled into a *composition* to handle event sequences. In a composition, rules are connected to each other by a set of input/output bindings. In the simplest case, a rule consumes an event. However, in many cases, we want the rule to be continually active; for example, a rule for recognizing a port on a component needs to fire multiple times so that multiple ports on the same component can be recognized. For this case, DiscoSTEP provides a bidirectional binding (denoted by <->). This means that a rule fires on an input without consuming that input (the event is implicitly reproduced as an output and fed back to the rule as input). Bidirectional bindings therefore are a convenient shorthand for two directional bindings.

Section 3.4.4 provides examples of rules and compositions.

2.3.3 Informal Runtime Semantics

Informally, DiscoTect runs DiscoSTEP specifications in the following sequence:

1. When an event is received by DiscoTect, associate the event with any rule that accepts that type of event.
2. For any rule that has a value for each of its inputs, evaluate the trigger.
3. If a trigger matches for a set of input events, execute the action with that set of events.
4. For output events that are composed with other rules, send the event as inputs to those rules. For output events not composed with other rules, emit them from DiscoTect as architectural events.

For each rule, an input event can be considered to be a set of events that match that type. Triggers match events to be

```

<call method_name="java.net.ServerSocket.accept"
      callee_id="19efb05"
      return_id="1d1acd3" />

```

Fig. 3. An example "call" event.

used in the rules. Formally, we model an event as colored tokens (where the color is given by the type) and a rule as a transition. For each event, a token is generated and put in a place before a transition. Rules consume tokens and put them in places after transitions. We discuss this formal definition in Section 2.4.

2.3.4 An Example DiscoSTEP Specification

To illustrate the concept of events and the use of rules and compositions, we now profile a simple program written in Java. In doing so, we illustrate how to specify events using DiscoSTEP and how DiscoTect uses this specification to generate an architectural description. The example is a simple system that implements a chat server.¹ The chat server creates a server socket and announces its intention to accept connections. When a client connects to the waiting server, a new thread (of type ClientThread) is started which forwards all messages from that client to all connected clients.

While this system is simple, it allows us to illustrate the concepts of DiscoSTEP. In later sections, we provide more complex case studies. In the remainder of this section, we show how we instrumented this system, describe its DiscoSTEP specification, and discuss how DiscoTect processes this specification to produce an architectural description.

The architectural description for this system is based on a client-server style. Informally, it is constructed as follows: When the server is created in the program, this maps to a server type component in the architecture; ClientThreads map to client components and the socket connection maps to an architectural connector between each client and server.

Instrumentation. The act of instrumenting a system to produce runtime events is not a novel aspect of DiscoTect. In fact, wherever possible, we use off-the-shelf technologies for our instrumentation. For Java-based systems, we have used AspectJ to define instrumentation *aspects* that are woven into the compiled bytecode of the programs. These aspects emit events when methods of interest enter or exit and when objects are constructed.

The aspects can reflectively retrieve information about the runtime environment of, for example, a call, to ascertain the calling object, the instance of the object that was called, the argument values and types that were passed to the method, the method signature, etc. The aspects are written to emit XML elements that conform to a schema expected by DiscoTect. For example, to instrument the ChatServer, we wove in aspects to emit events when methods were called and when objects were constructed.

Runtime Events. Two types of runtime events were collected from this running system: call events and init events. A call event is reported when a method is invoked.

1. The example system, including Java code and DiscoSTEP specification, is presented in Appendix A, which can be found on the Computer Society Digital Library at <http://computer.org/tse/archives.htm>.

Similarly, an object instantiation produces an init event. Take the following two events for example:

```
<init constructor_name="ServerSocket" instance_id="10">
```

```
<call method_name="ServerSocket.accept"
  callee_id="10" return_id="11" />
```

An init event is generated when

```
ServerSocket ss = new ServerSocket(1111)
```

is executed; a call event is triggered by an execution of a method call. For example, the call event above is emitted by the following statement execution:

```
Socket socket = ss.accept()
```

Because multiple ClientThreads can run concurrently, some of the runtime events, such as `InputStream.read` and `OutputStream.write`, show up in random order and, hence, may be interleaved with each other.

A fuller trace of the events that we retrieved when running the program is available in Fig. 10 of Appendix A, which can be found on the Computer Society Digital Library at <http://computer.org/tse/archives.htm>. These events can be fed into DiscoText either in real-time or offline, after the program has completed running.

DiscoSTEP Program. A DiscoSTEP program that specifies how to handle the interleaved events between the client and server was specified to capture how to map system events into architectural events. The full specification is given in Fig. 11 of Appendix A, which can be found on the Computer Society Digital Library at <http://computer.org/tse/archives.htm>; in this section, we discuss some of the rules and how these are combined with the event trace to produce the architecture. DiscoText takes the runtime events from the ChatServer to produce architectural events that construct a Client Server style representation of the system.

Fig. 4 shows a fragment of the DiscoSTEP program that includes two rules, and how they are composed. The `CreateServer` rule constructs an architecture Server component. It takes the input event under inspection to be an init event named `$e`. The output events include the string event `$server_id` and the `create_component` event `$create_server`. The condition for triggering this rule is that the `constructor_name` attribute of `$e` contains the string "ServerSocket." If the rule is triggered, the following action is taken: `$server_id` is assigned the id of the object constructed in the init event and an architecture event that constructs a server component named with the id of the newly created instances is assigned to `$create_server`.

The `$server_id` output from the `CreateServer` rule is fed to the `ConnectClient` rule, which has two inputs: `$e` and `$server_id`. Once these inputs are received by `ConnectClient`, the trigger will check to see if any events are calls to `ServerSocket.accept`. If so, output events `$client_id`, `$create_client`, and `$create_cs_connection` are assigned appropriate values to construct both the client component and the connector connecting it with the previously created server component.

Instead of being specific to this particular ChatServer program, our client server event processing program is

```
rule CreateServer {
  input { init $e; }
  output { string $server_id; create_component $create_server; }
  trigger { ? contains($e/@constructor_name, "ServerSocket") ? }
  action = { ?
    let $server_id := $e/@instance_id;
    let $create_server :=
      <create_component name="{ $server_id }"
        type="ServerT" />;
  }
}

rule ConnectClient {
  input { call $e; string $server_id; }
  output {
    create_component $create_client;
    create_connector $create_cs_connection;
    string $client_id;
  }
  trigger { ?
    contains($e/@method_name, "ServerSocket.accept")
    and $e/@callee_id = $server_id
  ? }
  action = { ?
    let $client_id := $e/@return_id;
    let $create_client :=
      <create_client name="{ $client_id }" type="ClientT" />;
    let $create_cs_connection :=
      <create_connector name=concat($client_id,"-", $server_id)
        type="CSConnectorT"
        end1="{ $server_id }" end2="{ $client_id }" />;
  ? }
}

composition {
  CreateServer.$server_id -> ConnectClient.$server_id;
  ...
}
```

Fig. 4. The DiscoSTEP rule to create a Server component.

generic enough to be applicable to *any* client server applications implemented with the same style (with, at the most, some minor changes in the triggers).

Both the compositions and the rules are well encapsulated. Rules are self-contained specifications, communicating with each other via inputs and outputs; compositions function as glue that assembles the rules. We can reuse compositions by applying them to a different system and reuse rules by assembling them with a different composition (and adding new rules if necessary).

2.3.5 Satisfying the Requirements

In this section, we revisit the requirements for DiscoSTEP that we introduced in Section 2.1 and discuss how DiscoSTEP meets these requirements.

- *Allow M-N mappings between events.* Since a DiscoSTEP rule can have an arbitrary number of inputs and outputs, this requirement is simply met by DiscoSTEP.
- *Keep track of information for use in subsequent stages.* Input events and output events are data structures that can be passed from one rule to the next. These data structures are used to store information that can be passed between rules. Compositions define how this data is passed between rules.
- *Cope with concurrent states.* The informal execution semantics defined in Section 2.3.3 describe how input events are propagated to each rule that can accept an event of that type. In this way, these events can start multiple execution threads to cope with concurrent states. A rule will wait until it gets a set of input events that match a trigger before firing. In

this way, interleaved threads of information can be managed.

- *Be flexible with respect to mapping implementation style to architectural style.* The fact that the types of events consumed and produced by are specified in a DiscoSTEP mapping means that multiple implementation styles and architectural styles can be manipulated by DiscoTect. Furthermore, though not described in detail in this paper, the abstract syntax of DiscoSTEP specifies that a composition itself may have input and output events, as well as subcompositions. In this way, compositions can be combined hierarchically to form more complex mappings. So, for example, it is possible to take a composition that identifies a mapping between system file usage and a data repository architectural style and combine that with a mapping that recognizes the construction of a pipe-filter architecture to define the mapping for a pipe-filter system that retrieves and stores data in files.

In sum, meeting these requirements means that a DiscoSTEP mapping *captures the way in which runtime events following an implementation style should be interpreted as operations on elements of an architectural style.*

2.4 Formal Definition of DiscoSTEP

To define the execution semantics of a DiscoSTEP program, and to formally explain how the mappings are interpreted, we use Colored Petri Nets (CP-nets) [19]. In [46], we informally described the semantics of DiscoTect mappings in terms of state machines. However, this semantics was awkward because of the need to retain multiple active states in the state machine to model the concurrency in the model. We believe that Colored Petri Nets are a more appropriate formalism for describing the semantics of DiscoSTEP mappings because their tokens provide a rich way of representing concurrent system states.

The DiscoSTEP language is formally described and modeled so that type checking and consistency checking can be done at the language level. The more interesting definition comes from how DiscoSTEP is formally modeled using Colored Petri Nets. In this section, we summarize how we translate a DiscoSTEP program into its equivalent Colored Petri Net; Appendix B, which can be found on the Computer Society Digital Library at <http://computer.org/tse/archives.htm>, contains the full semantic definition of DiscoSTEP. Informally, we translate the DiscoSTEP constructs into CP-net constructs in the following way:

- The types for the CP-net are formed out of the union of event types used as input and output from all the rules used in a DiscoSTEP composition.
- Each DiscoSTEP rule becomes a CP-net transition.
- Each group of connected input and output events becomes a CP-net place.
- Each DiscoSTEP trigger becomes a guard in the CP-net.
- The color for a place is derived from the connected input and output event type used to construct that place, which are guaranteed to be equivalent because of the DiscoSTEP type semantics.

- Each DiscoSTEP action becomes an action in the CP-net.
- CP-net arcs and nodes are constructed from the composition of the places and rules.

The complication in the translation comes from forming the places and dealing with the DiscoSTEP bidirectional binding. These complications are a direct result of keeping rules and composition separate.

2.4.1 Formally Modeling the Example

The ChatServer DiscoSTEP program uses the following types: `string`, `init`, `call`, `create_component`, `create_connector`, and `update_component`. By applying the translational semantics, we obtain the color sets for the CP-net as:

$$\Sigma = \{\text{string, init, call, create_component, create_connector, update_component}\}.$$

The transitions in the CP-net are created as:

$$T = \{\text{CreateServer, ConnectClient, ClientIO, ClientRead, ClientWrite, UpdateServer}\}.$$

By applying other translations, the inputs and outputs are translated into CP-net places, arcs, and node functions, the triggers and actions are translated into CP-net guards defined from transitions into predicates, and arc expressions are defined from arcs to *XQuery* expressions. Fig. 5 shows the resulting net. Note that the backward arcs from, for example, the `ClientRead` transition to the `$client_id` place are formed through the unification process in the translational semantics because the `$client_id` output of `ConnectClient` is bound to the inputs of more than one rule (`ClientIO`, `ClientRead`, and `ClientWrite`).

This CP-net representation is then “executed” in our toolset. When an event comes into the net, a token is created for each place that is relevant for that event so that it can be compared with the associated guards for interested transitions. There is no priority on competing transitions—each matching set of tokens is processed by the transition when they match.

3 IMPLEMENTATION OF DISCOTECT

Recall from Section 2 that, to provide a framework for discovering architectures, we need to solve three challenges. In this section, we discuss our response to each challenge.

Monitoring. We use existing probing technologies to extract monitoring events. In the case studies in Sections 4 and 5, we used AspectJ to monitor object creation, method invocation, etc. We provide a library that allows aspects to produce system events formatted as XML strings which are placed on a JMS event bus to be consumed by DiscoTect.

Mapping. We have implemented the DiscoTect engine in Java, which follows the implementation outlined in this paper. The DiscoSTEP language is implemented using JavaCC, a Java-based compiler generator, and uses the Saxon *XQuery* processor to parse and execute the *XQuery* triggers and rules in a DiscoSTEP specification.

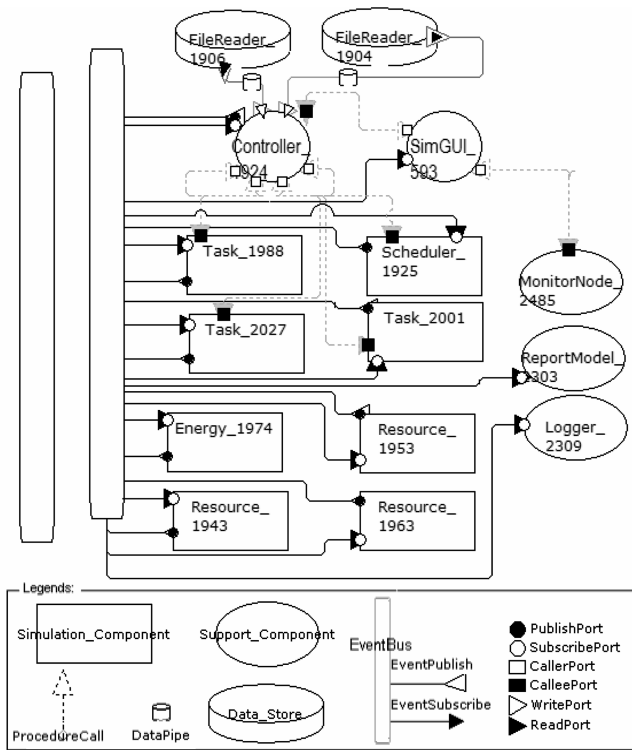


Fig. 7. The discovered architecture of AAMS.

Up to this point, we had not discovered anything about the publish/subscribe part of the architecture itself. The preliminary discovery results informed us that all the resource and task components interact with an object of the *PubSub* class using two procedure calls named *publish* and *subscribe*. We conjectured that the system implements publish/subscribe by creating a *PubSub* object and invoking its two methods. This led us to design a specification for this portion of the architecture. This specification creates an *EventBus* connector when it notices the instantiation of a *PubSub* object in the implementation. Once this has been done, an *EventTaker* role is created when DiscoTect notices a call to the *publish* method of the *PubSub* object and a *Publish* port on the component corresponding to the call and attaches them. Similarly, *PubSub.subscribe* leads to the creation of an *EventSender* role on the *EventBus* providing the method, the creation of a *Subscribe* port in the component requesting the method, and the creation of the attachment of port and role.

4.2 The Discovered Architecture

The final DiscoSTEP specification for AAMS consists of 10 rules. This specification took an expert about one week to develop, which included interacting with the original source code developer.

Applying the above DiscoSTEP specification to a running instance of AAMS yields the architectural model in Fig. 7. (We have manually laid out this model to enable easier comparison with the view in Fig. 6.) We uncovered four types of discrepancies between the documented architectural view and our discovered one.

1. *Isolated, extraneous components/connectors.* The result shows two *EventBus* connectors, one of which is isolated from the rest of the system. This indicates that one instance is instantiated but never used. Code optimization should resolve this discrepancy.
2. *Additional connections between components.* Fig. 7 does not show any connections between the controller component and simulation components such as tasks and schedulers. Nor does it inform us that some of the support components (e.g., Logger and Reporting) also subscribe to the event bus. Ignoring those “backdoor” connections makes the architectural view less accurate; moreover, it might compromise architectural analysis where all meaningful interactions between components should be considered. For example, in evaluating the performance of a publish/subscribe infrastructure, the existence of hidden communication channels could invalidate deadlock or throughput analyses.
3. *Misplaced connections between components.* The discovered architecture shows a different File I/O scheme: Instead of the GUI reading three files (cf., Fig. 6), the controller reads two files. This discrepancy could cause errors during evolution if the AAMS system was to work in a distributed environment. The evolution might require that the file reading components run on the same computer as that containing the files. The documented architecture would suggest that the Simulator GUI is the component that should be associated with the files, when, in fact, according to the implementation, this is incorrect.
4. *Missing components/connectors.* Two of the components—User and Env (Environment)—recorded in the document do not show up in the discovered architecture.

To confirm that DiscoTect discovered the actual architecture of the implementation and to understand the discrepancies, we conferred with the AAMS developers. They agreed that DiscoTect produced a more complete and correct architectural description than their diagram and that it had uncovered some errors in their coding. However, the missing User and Environment components are due to the fact that these represent user interaction and are not actual components in the implementation.

5 EJB CASE STUDY

In this section, we present a second case study to determine the runtime architecture of the Duke’s Bank Application—a simple EJB (Enterprise Java Beans) banking application created by Sun Microsystems as a demonstration of EJB functionality. Duke’s Bank allows bank customers to access their account information and to transfer balances from one account to another. It also provides an administration interface for managing customers and accounts. We use this case study to demonstrate how the architecture of an EJB application can be discovered using DiscoTect. We chose this system because its architecture is well documented in Sun Microsystems’ J2EE (Java2 Platform, Enterprise Edition) tutorial [16], which enables us to compare the

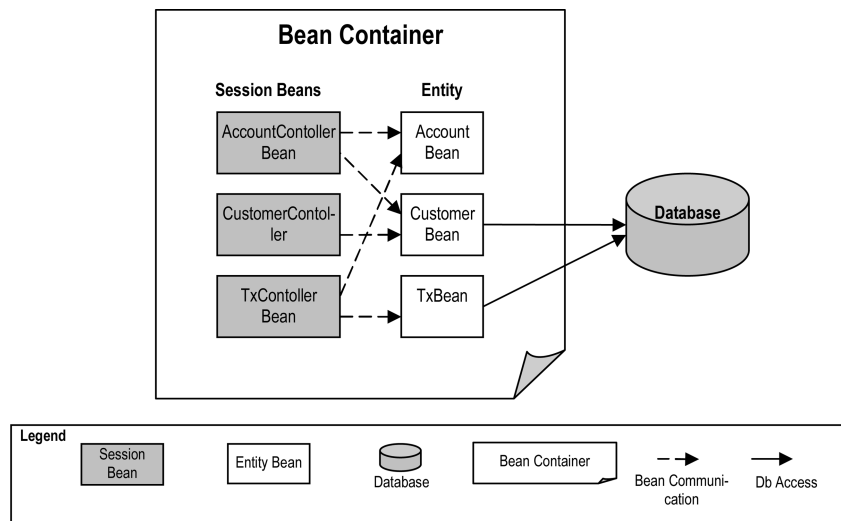


Fig. 8. Documented architectural view of the Duke's Bank application.

discovered architecture with the one presented in the documentation. For the case study, we used version 1.3 of the application provided by Sun and version 3.2.7 of the JBoss application server.

To instrument the system, we wrote an aspect that injected advice to object instantiations, method calls, and field modifications. We compiled the Duke's Bank application along with the aspect, using the AspectJ compiler, so that system execution events were traced as the application ran. In total, we developed 200 lines of AspectJ code to instrument the application and the server.

Fig. 8 gives a view of how the components interact in the Duke's Bank system as presented in [16]. The EJB application has three session beans: AccountControllerBean, CustomerControllerBean, and TxControllerBean (Tx stands for a business transaction, such as transferring funds). These session beans provide a client's view of the application's business logic. For each business entity represented in the simplified banking model, the application has a matching entity bean: AccountBean, Customer Bean, and TxBean. The business methods of the Account ControllerBean session bean manage the account-customer relationship and get the account information using the AccountBean and CustomerBean entity beans. The CustomerControllerBean provides methods for creating, removing, and updating customers through CustomerBean entity beans. The TxControllerBean session bean handles bank transactions. It accesses AccountBean entity beans to verify the account type and to set the new balance and accesses TxBean to keep records of the transactions.

5.1 Design of the EJB DiscoSTEP Program

We now present the steps taken to produce the DiscoSTEP specification to discover the Duke's Bank architecture. We used a specialization of an EJB style that distinguishes participating components as entity beans, session beans, bean containers, database, etc. These component types are based on the EJB specification found in [10].

As we did in the previous case study, we first produced primitive rules that merely observed object interaction and

creation (through procedure calls and object instantiations). We then refined these rules to classify objects into their architectural counterparts (e.g., Beans, Bean Containers, Database, etc.) by checking the class constructor names. For example, we created a SessionContainer object when its constructor had the name of "SessionContainer." The relationships between the beans, the bean containers, and the database were captured in the following way: According to the EJB specification, the beans are maintained by their corresponding containers, so we connected the beans with the containers controlling them by observing the procedure calls made by the containers to manage the life cycles of the beans. Knowing that database access was implemented using JDBC (Java Database Connectivity) [18], we monitored the standard JDBC APIs to uncover the connections between the beans and the database; the interactions between the beans were also monitored and represented as connectors linking them together.

5.2 The Discovered Architecture

The final DiscoSTEP specification code for Duke's Bank consists of eight rules. The process of developing the code iteratively took a DiscoSTEP expert approximately three days, which included understanding the working of the application and the server, developing the aspects, and writing the DiscoSTEP specification.

Applying the specification just described to a running instance of Duke's Bank yields the architectural model in Fig. 9. We have manually organized the layout of this model for better comprehensibility. We can make the following observations based on this process:

1. *Reflection of runtime instances.* Besides showing the bean and the containers, the discovered result also details each bean and container instance created at runtime. Tracing the individual bean and container instances is useful for performance analysis and fault diagnosis. In addition, the relatively complex m-to-n relationships between beans and bean containers are revealed.

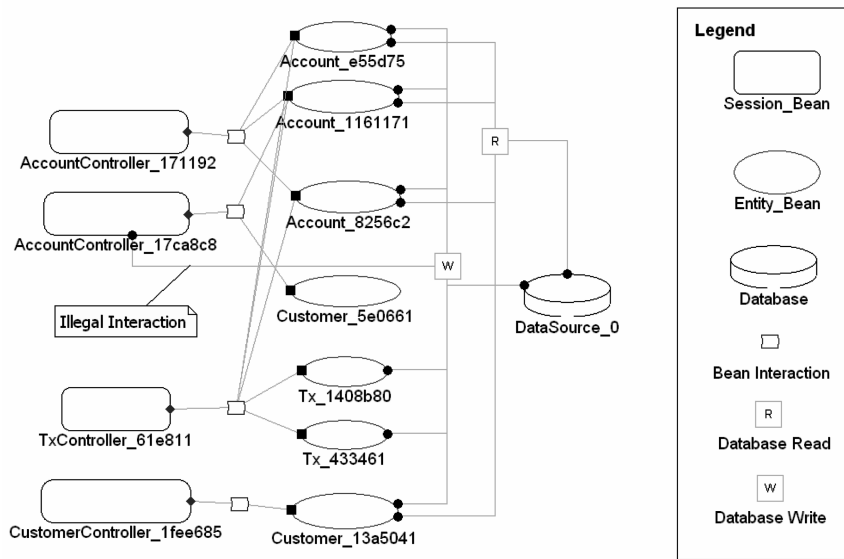


Fig. 9. Discovered architecture of Duke's Bank.

2. *Verification of Bean Interplay.* The interactions between the beans shown in Fig. 9 are consistent with those described in the architecture shown in Fig. 8: There are communication channels between AccountControllerBean and AccountBean, AccountControllerBean and CustomerBean, CustomerControllerBean and CustomerBean, TxControllerBean and TxBean, and TxControllerBean and AccountBean.
3. *Discrepancies in Database Access.* Fig. 8 does not show any connections between the session beans and the database, which implies that all database access goes through the entity beans. This is consistent with Sun's EJB specification [18]. However, a "database write" connector did appear in the discovered architecture. Further (manual) source code analysis confirmed that AccountControllerBean does directly write to the database, a violation of the architecture. As discussed in the previous section, identifying communication "backdoor" connections like this is useful for architectural analysis and to ensure architectural conformance.²

6 RELATED WORK

To analyze a running system, it is first necessary to get information from it. There are many technologies available for monitoring systems and we build on those. For example, Balzer and Goldman [4] and Wells and Pazandak [43] provide systems to instrument the source code to produce trace information and manipulate runtime artifacts to get the information; aspect-oriented systems such as AspectJ allow use to instrument binary code when the source is not available. We can use any of these approaches to get information out of the system.

Monitoring mechanisms do not by themselves solve the hard problem of mapping from code to more abstract

models; event correlation and abstraction is also needed. Work by Dias and Richardson [9] uses an XML-based language to describe runtime events and use patterns to map these events into high-level events. However, analyzing these abstracted events to determine architectural structure is not addressed. In addition, a simple static mapping from low-level system events to high-level events has limited expressiveness. For example, it cannot handle the case where the event analyzer initially has an interest in one set of events, but then changes its interest after the initial events have occurred. Also, it doesn't provide a way of specifying event correlations or mapping a series of correlated low-level events to a single high-level event—a crucial capability needed when discovering the architecture of a system. Approaches to correlating events, for example, GEM [29] and SEL [50] and work by Kaiser et al. [20], provide sophisticated event correlation languages that can detect when certain combinations of events have occurred and, in some cases, what actions to perform based on those correlations. Our approach is similar, but it makes architectural styles or patterns explicit; our mappings can be used to map low-level events to high-level architectural events that are amenable to checking for consistency against particular architectural styles.

In prior work, we developed an infrastructure for doing certain kinds of abstraction [15]. However, this approach was limited to observing the properties of a system and reflecting them in a preconstructed architectural model. Here, we show how to create such a model.

A number of researchers have investigated the problem of presenting dynamic information to an observer. For example, some researchers present information about variables, threads, activations, object interactions, and so forth [32], [40], [41], [49]. Ernst et al. show how to dynamically detect program invariants by examining values computed during a program execution and by looking for patterns and relationships among them [11]. This is somewhat different from detecting architectural structure.

2. Note that this error has been corrected in the most recent version of the Duke's Bank J2EE application.

Madhav [26] describes a system that allows Ada 95 programs to be monitored dynamically to check conformance to a Rapide architectural specification [25]. His approach requires the source code to be annotated so that it can be transformed to produce events to construct the architecture. In contrast, our approach does not require access to the source code and it does not rely on explicit architectural construction directives to be embedded in the code as does the approach used by Aldrich et al. [2].

There has been research in using Colored Petri Nets to dynamically determine software architectures from running systems using Colored Petri Nets. This work either focuses on a particular quality attribute (e.g., [1] for performance analysis), uses UML as the modeling language, which is very close to the actual code [48], or checks that interactions on preknown connections are of the correct protocol [6]. In contrast, our approach focuses on determining the component and connector architecture view of a running system where the particular architecture components and connections may not be known *a priori* and where the actual architectural view may be much more abstract than the programming language constructs used to implement the system. Moreover, our approach is agnostic about the specific architectural styles that may be used and the representation language.

A large body of research has investigated specification of the dynamic behavior of software architectures. Of the many approaches, some use explicit state machines (e.g., as described by Allen and Garlan [3] and Vieira et al. [39]). These approaches, however, do not link architecture to an executing system.

7 CONCLUSIONS AND FUTURE WORK

In this paper, we have described an approach to “discovering” the architecture of a running system that uses a set of pattern recognizers that convert monitored system observations into architecturally meaningful events. Our key idea and contribution is to exploit implementation regularities and knowledge of the architectural style that is being implemented to create a mapping that can be applied to any system that conforms to the implementation conventions, to yield a view in that architectural style. The mapping itself defines a novel form of behavior specification (realized as a Colored Petri Net) that relates low-level events to architecturally significant actions. The power of Petri Nets is used to model the concurrent threads of event recognition, allowing us to disentangle the interleaved sequences of low-level events that contribute to higher-level architectural behavior.

There are a number of advantages of this approach. First, it can be applied to any system that can be monitored at runtime. In our case, we have demonstrated two case studies written in Java, but we have recently successfully experimented with the use of AspectC to extract runtime information from C and C++ programs. In general, any monitoring environment that allows us to capture object creation, method invocation, and instance variable assignment will serve as a sufficient foundation for our runtime monitoring. Second, by simply substituting one mapping description for another, it is possible to accommodate

different implementation conventions for the same architectural style or, conversely, to accommodate different architectural styles that employ the same implementation conventions. For example, although not described here, we have been able to discover the Pipe/Filter architecture of systems implemented using different pipe conventions.

There are, however, several inherent weaknesses to the approach. The first is that it only works if an implementation obeys regular coding conventions. Completely ad hoc bodies of code are unlikely to benefit from this technique. Second, it only works if one can identify a target architectural style so that the mapping “knows” the output vocabulary. Third, as with any analysis based on runtime observations, it suffers from the problem that you can only analyze what is actually executed. Hence, questions like “is there *any* execution that might violate a set of style constraints” cannot be directly answered using this method. Fourth, the DiscoSTEP mapping needs to be created via an iterate-and-test paradigm and, hence, the results are somewhat dependent on the skill of the creator of the recognizer. Our techniques require some skill and judgment to be employed effectively. Thus, our techniques are best viewed as one of several technologies that an architect must have in his arsenal of architecture conformance checking tools. For example, we believe that DiscoTest can be effectively combined with static analysis tools such as Dali [21] or Armin [30] to provide complementary kinds of analysis, whereby runtime observations can be combined with statically extracted facts. In this way, we should be able to achieve a more complete and accurate picture of the as-built system. Finally, the level of effort required to write the specification suggests that this approach is best applied when there are many examples of a system using the same architectural styles and coding conventions. For example, the EJB Duke’s Bank case study requires understanding of the inner workings of a particular application server; however, we believe that applying the same specification to another EJB system running on the same server will require only a small incremental effort.

These potential defects also point the way to future research in this area. First is the area of system monitoring. As mentioned, we have experimented with a number of existing monitoring technologies for Java and C++. However, monitoring technologies for other kinds of implementations and system properties is a research area that should continue to provide increasing capabilities in the future that we can leverage.

Second is the area of codifying the ways in which architectural styles are implemented. As technology advances, implementation techniques will necessarily change, and it will be important to augment the set of mappings as that happens. We envision a large library of recognizers for common architectural frameworks available as open source libraries that would track the most common architectural frameworks in practical use.

Third is the area of architectural coverage metrics, similar to coverage metrics for testing. It would be good, for example, to have some confidence that, in running a system with various inputs, we have exercised a sufficiently

comprehensive part of the system to “know,” with some certainty, what its architecture is.

Fourth, we would like to find a way to make the definition of implementation-architecture mappings more declarative. While the operational definition of state machines as the carrier of those mappings is a good first step, we can imagine more abstract forms of characterization that will be easier to create and analyze.

Fifth, while the approach we have outlined focuses primarily on recognizing architectural *structure*, we believe it could be easily extended to *architectural behavior*. For example, we can imagine using the same runtime abstraction techniques to check that the observed interaction between two components conforms to the protocol expected over the corresponding architectural connector. Similarly, we might observe timing behavior, which could be compared with an architectural specification of expected performance.

Finally, we are still gaining experience with the use of these tools in practice. The process of developing the DiscoSTEP mapping is a long and iterative one that currently requires intimate knowledge of the implementation conventions. While we can never truly eliminate this need, we are investigating processes, conventions, and tools that could make this less time-consuming and error-prone.

ACKNOWLEDGMENTS

The research described in this paper was supported by the US Defense Advanced Research Projects Agency, under Grants N66001-99-2-8918 and F30602-00-2-0616, by a Software Engineering Institute (SEI) Internal R&D Grant, and by the US Army Research Office under Grants DAAD19-02-1-0389 (“Perpetually Available and Secure Information Systems”) and DAAD19-01-1-0485. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of any of the funding agencies.

REFERENCES

- [1] M. Ajmone Marsan, G. Balbo, G. Chiola, and G. Conte, “Modeling the Software Architecture of a Prototype Parallel Machine,” *Proc. SIGMETRICS Performance Evaluation Rev.*, vol. 15, no. 1, pp. 175-185, 1987.
- [2] J. Aldrich, C. Chambers, and D. Notkin, “ArchJava: Connecting Software Architecture to Implementation,” *Proc. 24th Int’l Conf. Software Eng.*, 2002.
- [3] R. Allen and D. Garlan, “Formalizing Architectural Connection,” *Proc. Int’l Conf. Software Eng. (ICSE)*, 1994.
- [4] R.M. Balzer and N.M. Goldman, “Mediating Connectors,” *Proc. 19th IEEE Int’l Conf. Distributed Computing Systems Workshop Electronic Commerce and Web-Based Applications*, 1999.
- [5] L. Bass, P. Clements, and R. Kazman, *Software Architecture in Practice*, second ed. Addison-Wesley, 2003.
- [6] Q. Bai, M. Zhang, and K.T. Win, “A Colored Petri Net Based Approach for Multi-Agent Interactions,” *Proc. Second Int’l Conf. Autonomous Robots and Agents*, Dec. 2004.
- [7] P. Clements, R. Kazman, and M. Klein, *Evaluating Software Architectures: Methods and Case Studies*. Addison-Wesley, 2001.
- [8] P. Clements, F. Bachmann, L. Bass, D. Garlan, J. Ivers, R. Little, R. Nord, and J. Stafford, *Documenting Software Architectures: Views and Beyond*. Addison Wesley, 2002.
- [9] M. Dias and D. Richardson, “The Role of Event Description on Architecting Dependable Systems (Extended Version from WADS),” *Lecture Notes in Computer Science—Book on Architecting Dependable Systems*, Springer-Verlag, 2003.
- [10] Sun Microsystems, <http://java.sun.com/products/ejb/docs.html>, 2006.
- [11] M.D. Ernst, J. Cockrell, W.G. Griswold, and D. Notkin, “Dynamically Discovering Likely Program Invariants to Support Program Evolution,” *IEEE Trans. Software Eng.*, vol. 27, no. 2, Feb. 2001.
- [12] D. Garlan, R.J. Allen, and J. Ockerbloom, “Exploiting Style in Architectural Design,” *Proc. ACM SIGSOFT ’94 Symp. Foundations of Software Eng.*, 1994.
- [13] D. Garlan, R.T. Monroe, and D. Wile, “Acme: Architectural Description of Component-Based Systems,” *Foundations of Component-Based Systems*, G. Leavens and M. Sitaraman, eds., Cambridge Univ. Press, 2000.
- [14] D. Garlan, A.J. Kompanek, and S.-W. Cheng, “Reconciling the Needs of Architectural Description with Object Modeling Notations,” *Science of Computer Programming*, vol. 44, no. 1, pp. 23-49, July 2002.
- [15] D. Garlan, S.-W. Cheng, and B. Schmerl, “Increasing System Dependability through Architecture-Based Self-Repair,” *Architecting Dependable Systems*, R. de Lemos, C. Gacek, and A. Romanovsky, eds., Springer-Verlag, 2003.
- [16] Sun Microsystems, <http://java.sun.com/docs/books/j2eetutorial/index.html>, 2006.
- [17] D. Jackson and A. Waingold, “Lightweight Extraction of Object Models from Bytecode,” *Proc. 1999 Int’l Conf. Software Eng.*, 1999.
- [18] Sun Microsystems, <http://java.sun.com/products/jdbc>, 2006.
- [19] K. Jensen, “An Introduction to the Theoretical Aspects of Coloured Petri Nets,” *A Decade of Concurrency*, J.W. de Bakker, W.-P. de Roever, and G. Rozenberg, eds., pp. 230-272, Springer-Verlag, 1994.
- [20] G. Kaiser, J. Parekh, P. Gross, and G. Veleto, “Kinesthetics eXtreme: An External Infrastructure for Monitoring Distributed Legacy Systems,” *Proc. Fifth Int’l Active Middleware Workshop*, 2003.
- [21] R. Kazman and S.J. Carriere, “Playing Detective: Reconstructing Software Architecture from Available Evidence,” *J. Automated Software Eng.*, vol. 6, no. 2, 1999.
- [22] R. Kazman, J. Asundi, J.S. Kim, and B. Sethananda, “A Simulation Testbed for Mobile Adaptive Architectures,” *Computer Standards and Interfaces*, 2003.
- [23] G. Kiczales, E. Hilsdale, J. Huginin, M. Kersten, J. Palm, and W. Griswold, “Getting Started with AspectJ,” *Comm. ACM*, vol. 4, no. 10, Oct. 2001.
- [24] P. Kruchten, “The 4+1 View Model of Architecture,” *IEEE Software*, vol. 12, no. 6, pp. 42-50, 1995.
- [25] D.C. Luckham, “Rapide: A Language and Toolset for Simulation of Distributed Systems by Partial Orderings of Events,” *Proc. DIMACS Partial Order Methods Workshop*, 1996.
- [26] N. Madhav, “Testing Ada 95 Programs for Conformance to Rapide Architectures,” *Proc. Reliable Software Technologies—Ada Europe ’96*, 1996.
- [27] N. Medvidovic and R.N. Taylor, “A Classification and Comparison Framework for Software Architecture Description Languages,” *IEEE Trans. Software Eng.*, vol. 26, no. 1, pp. 70-93, Jan. 2000.
- [28] G.C. Murphy, D. Notkin, and K.J. Sullivan, “Software Reflexion Models: Bridging the Gap between Source and High-Level Models,” *Proc. 1995 ACM SIGSOFT Symp. Foundations of Software Eng.*, 1995.
- [29] M. Mansouri-Samani and M. Sloman, “GEM: A Generalized Event Monitoring Language for Distributed Systems,” *IEEE/IOP/BSC Distributed Systems Eng. J.*, vol. 4, no. 2, June 1997.
- [30] L. O’Brien and C. Stoermer, “Architecture Reconstruction Case Study,” *Software Eng. Inst. Technical Note CMU/SEI-2003-TN-008*, 2003.
- [31] D. Perry and A. Wolf, “Foundations for the Study of Software Architecture,” *ACM SIGSOFT Software Eng. Notes*, vol. 17, no. 4, 1992.
- [32] S. Reiss, “JIVE: Visualizing Java in Action (Demonstration Description),” *Proc. 25th Int’l Conf. Software Eng.*, 2003.
- [33] B. Schmerl and D. Garlan, “AcmeStudio: Supporting Style-Centered Architecture Development (Demonstration Description),” *Proc. 26th Int’l Conf. Software Eng.*, May 2004.
- [34] M. Shaw, R. Deline, D. Klein, T.L. Ross, D.M. Young, and G. Zelesnik, “Abstractions for Software Architecture and Tools to Support Them,” *IEEE Trans. Software Eng.*, vol. 21, no. 4, Apr. 1995.
- [35] M. Shaw and D. Garlan, *Software Architecture: Perspectives on an Emerging Discipline*. Prentice Hall, 1996.

- [36] O. Spinczyk, A. Gal, and W. Schroder-Preikschat, "AspectC++: An Aspect-Oriented Extension to the C++ Programming Language," *Proc. 40th Int'l Conf. Tools Pacific: Objects for Internet, Mobile, and Embedded Applications*, vol. 10, 2002.
- [37] R.N. Taylor, N. Medvidovic, K.M. Anderson, E.J. Whitehead, J.E. Robbins, K.A. Nies, P. Oriezy, and D. Dubrow, "A Component- and Message-Based Architectural Style for GUI Software," *IEEE Trans. Software Eng.*, vol. 22, no. 6, June 1996.
- [38] S. Vestel, "MetaH Programmer's Manual, Version 1.09," technical report, Honeywell Technology Center, 1996.
- [39] M. Vieira, M. Dias, and D.J. Richardson, "Software Architecture Based on Statechart Semantics," *Proc. 10th Int'l Workshop Component Based Software Eng.*, 2001.
- [40] R.J. Walker, G.C. Murphy, B. Freeman-Benson, D. Wright, D. Swanson, and J. Isaak, "Visualizing Dynamic Software System Information through High-Level Models," *Proc. ACM SIGPLAN Int'l Conf. Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, 1998.
- [41] R.J. Walker, G.C. Murphy, J. Steinbok, and M.P. Robillard, "Efficient Mapping of Software System Traces to Architectural Views," *Proc. CASCON*, S.A. MacKay and J.H. Johnson, eds., 2000.
- [42] W3C Consortium, "XQuery 1.0 and XPath 2.0 Formal Semantics, W3C Working Draft 4 April 2005," <http://www.w3.org/TR/2005/WD-xquery-semantics-20050404/>, Apr. 2005.
- [43] D. Wells and P. Pazandak, "Taming Cyber Incognito: Surveying Dynamic/Reconfigurable Software Landscapes," *Proc. First Working Conf. Complex and Dynamic Systems Architectures*, 2001.
- [44] W3C, <http://www.w3.org/TR/xquery/>, 2006.
- [45] W3C, <http://www.w3.org/XML/>, 2006.
- [46] H. Yan, D. Garlan, B. Schmerl, J. Aldrich, and R. Kazman, "DiscoTect: A System for Discovering Architectures from Running Systems," *Proc. 26th Int'l Conf. Software Eng.*, May 2004.
- [47] H. Yan, J. Aldrich, D. Garlan, R. Kazman, and B. Schmerl, "Discovering Architectures from Running Systems: Lessons Learned," Technical Report CMU/SEI-2004-TR-016, Dec. 2004.
- [48] J. Xu and J. Kuusela, "Modeling Execution Architecture of Software System Using Colored Petri Nets," *Proc. First Int'l Workshop Software and Performance (WOSP '98)*, pp. 70-75, Oct. 1998.
- [49] A. Zeller, "Animating Data Structures in DDD," *Proc. SIGCSE/SIGCUE Program Visualization Workshop*, 2000.
- [50] D. Zhu and A.S. Sethi, "SEL, A New Event Pattern Specification Language for Event Correlation," *Proc. ICCCN-2001, 10th Int'l Conf. Computer Comm. and Networks*, Oct. 2001.



Bradley Schmerl received the PhD degree in computer science from Flinders University in South Australia. He is a senior systems scientist at Carnegie Mellon University. His research interests include dynamic adaptation, software architectures, and software engineering environments.



Jonathan Aldrich received the PhD degree from the University of Washington. He is an assistant professor of computer science at Carnegie Mellon University. His research interests include applying language and program analysis and verification techniques to software engineering problems.



David Garlan received the PhD degree in computer science from Carnegie Mellon University (CMU). He is a professor of computer science at CMU. His research interests include software architectures, formal methods, self-healing systems, and task-based computing.



Rick Kazman received the PhD degree in computational linguistics from Carnegie Mellon University (CMU). He is a visiting scientist at the Software Engineering Institute and CMU and a professor of information technology management at the University of Hawaii. His research interests include software architecture, software design, analysis, and reverse engineering tools, open source software, and software engineering economics.



Hong Yan is a doctoral candidate in computer science at Carnegie Mellon University. His research interests include design and analysis of software systems and computer networks.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.