Automating UI Generation by Model Composition

R. E. Kurt Stirewalt
Computer Science and Engineering Dept.
Michigan State University
East Lansing, Michigan 48824

Spencer Rugaber
College of Computing
Georgia Institute of Technology
Atlanta, Georgia 30332-0280

Copyright 1998 IEEE. Published in the 13th Conference on Automated Software Engineering (ASE'98) October 13-16, 1998 in Honolulu, Hawaii. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works, must be obtained from the IEEE. Contact: Hoes Lane/ PO Box 1331/ Piscataway, NJ 08855-1331, USA.

Abstract

Automated user-interface generation environments have been criticized for their failure to deliver rich and powerful interactive applications [18]. To specify more powerful systems, designers require multiple specialized modeling notations [12, 14]. The model-composition problem is concerned with automatically deriving powerful, correct, and efficient user interfaces from multiple models specified in different notations. Solutions balance the advantages of separating code generation into specialized code generators with deep, model-specific knowledge against the correctness and efficiency obstacles that result from such separation. We present a correct and efficient solution that maximizes the advantage of separation through run-time composition mechanisms.

1 Introduction

Building user interfaces (UIs) is time consuming and costly. In systems with graphical UIs (GUIs), nearly 50% of source code lines and development time could be attributed to the UI [11]. GUIs are usually built from a fixed set of modules composed in regular ways. Hence, GUI construction is a natural target for automation. Tools have been successful in supporting the presentation aspect of GUI functionality, but they provide only limited support for specifying behavior and the interaction of the UI with the underlying application functionality. The model-based approach to interactive system development addresses this deficiency by decomposing UI design into the construction of separate models, each of which is declaratively specified [5]. Once specified, automated tools integrate the models and generate an efficient system from them. The *model-composition problem* is the need to efficiently implement and automatically integrate interactive software from separate declarative models. This paper introduces the model-composition problem and presents a solution.

A *model* is a declarative specification of some single coherent aspect of a user interface, such as its appearance, or how it interfaces and interacts with the underlying application functionality. By focusing attention on a single aspect of a user interface, a model can be expressed in a highly-specialized notation. This property makes systems developed using the model-based approach easier to develop and maintain than systems produced using other approaches.

The MASTERMIND project [5, 12] is concerned with the automatic generation of user interfaces from three kinds of models: *Presentation models* represent the appearance of user interfaces in terms of their widgets and how the widgets behave; *Application models* represent which parts (functions and data) of applications are accessible from the user interface; and *Dialogue models* represent end-user interactions, how they are ordered, and how they affect the presentation and the application. The dialogue model acts as the glue between the presentation and application models by expressing constraints on the sequencing of behavior in these models. Model-specific compilers generate modules of code from each model, and these resulting modules are composed (**Figure 1**). A distinguishing characteristic of

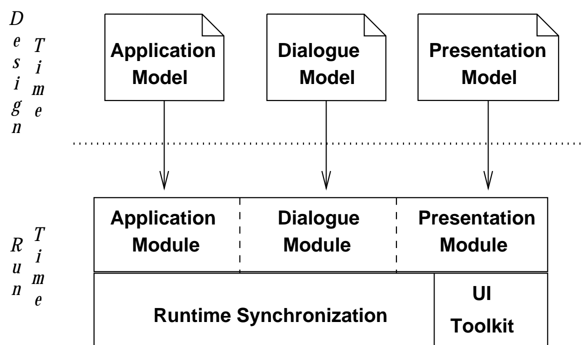

Figure 1: Model-based code generation

MASTERMIND is that the model-specific code generators work independently of one another.

Composing the code generated from multiple models is difficult. A model, by design, represents some aspects of a system and is neutral with respect to others [3]. Inevitably, however, functionality described in one model overlaps with or is dependent upon functionality described in another. A button, for example, is specified in the presentation model, but the behavior of the button influences behavior in other models, such as when pressing the button causes other widgets to be enabled or disabled. Such effects are described in the dialogue model. The effect of pressing a button might also cause some application method to be invoked. Such effects are described in the application model. When code generated from multiple models must cooperate, these redundancies and dependencies can be difficult to resolve. Resolving them automatically means that behavior in different models must be unified and the mechanism for this unification must be implemented efficiently.

The model-composition problem is concerned with automatically deriving powerful, correct, and efficient user-interfaces from separate presentation, dialogue, and application models. We present a two-fold solution. First, we formalize the three models as concurrent agents, which synchronize on common events (§ 3). Second, we make the synthesis efficient by implementing a run-time dialogue engine that actively synchronizes behavior in the code generated from presentation and application models (§ 4). We present the results of this approach on two examples and give evidence to show that it scales up (§ 5).

2 Background

We now present the issues that comprise the model-composition problem and the influences on our solution. Model-based approaches to user-interface generation use models that are specified in diverse and often incompatible notations. This characteristic complicates model composition because the composition mechanisms in one model may not exist in another (§ 2.1). Prior work in user-interface architectures suggests using communicating agents to structure user-interface code (§ 2.2). Formal models of communicating agents provide an called conjunction, which is useful for composing partial specifications of a system (§ 2.3). The contribution of this paper is an extension of conjunction from a specification-composition operator into a run-time–composition mechanism.

2.1 Model-based generation

The model-based approach to interactive system development bases system analysis, design, and implementation on a common repository of models. Unlike conventional software engineering, in which designers construct artifacts whose meaning and relevance can diverge from that of the delivered code, in the model-based approach, designers build models of critical system attributes and then analyze, refine, and synthesize these models into running systems. Model-based UI generation works on the premise that development and support environments may be built around declarative models of a system. Developers using this paradigm build interfaces by specifying models that describe the desired interface, rather than writing a program that exhibits the behavior [17].

One characteristic of model-based approaches is that, by restricting the focus of a model to a single attribute of the system, modeling notations can be specialized and highly declarative. The MASTERMIND Presentation Model [6], for example, combines concepts and terminology from graphic design with mechanisms for composing complex presentations through functional constraints. Dialogue models use state and event constructs to describe the user-computer conversation. Example notations include StateCharts [8] and Petri nets [13], which use a variety of composition mechanisms that include state hierarchy, concurrency, and communication. The MASTERMIND Application Model combines concepts and terminology from

object-oriented design techniques [15] with mechanisms for composing complex behavior based on method invocation.

These examples illustrate that composition mechanisms in one model may not exist in another model. It is not clear that any of these intra-model mechanisms are sufficient for inter-model composition, the subject of this paper. The model-composition problem can be restated as the need to unify behavior in multiple models without violating the rules of intra-model composition and while generating efficient code.

2.2 Multi-agent user-interface architectures

Our approach to model composition is based on prior work in multi-agent user-interface architectures, which provide design heuristics for structuring an interactive system. These architectures describe interactive systems as collections of independent communicating *agents*, which are independent computational units with identity and behavior. Two general frameworks—the Model-View-Controller (MVC) [9] and the Presentation-Abstraction-Control (PAC) [7]—define specific agent roles and provide guidance on how these agents should be connected.

The roles prescribed by the PAC framework most closely resemble those of the MASTERMIND models. In PAC, a *presentation* agent maintains the state of the display and accepts input from the user, an *abstraction* agent maintains a representation of the underlying application state, and a *controller* agent ensures that presentation and abstraction remain synchronized. The MASTERMIND Presentation, Application, and Dialogue models are descriptions of the roles to be played by presentation, abstraction, and controller agents.

Since MASTERMIND models describe PAC agents, we chose to make MASTERMIND models compose in the same manner that PAC agents compose. Specifically, the presentation and application models define actions, which are ordered by temporal constraints in the dialogue model. To make these ideas more formal, we built upon prior work in formal definitions of agent composition.

2.3 Formal models of agents

The PAC framework provides heuristic definitions of user-interface agent roles and connections. However, to generate code from models, we need a more formal definition of agents and agent composition. Process algebras, such as LOTOS [4], are formal notations used to describe concurrent, communicating agents. Process algebras are particularly useful for describing user-interface–agent composition [1, 2].

A *process* is an entity whose internal structure can only be discovered by observing the actions in which it participates. In this paper, we use the LOTOS [4] notation to specify processes. LOTOS is a process algebra that uses temporal operators to specify permissible orderings and dependences over actions.

A LOTOS process performs actions and interacts with other, concurrently executing, processes. Actions are built up out of atomic units called *events*. The set of events in which a process P may participate is called the *alphabet* of P (denoted $\alpha(P)$). If an event e is in the alphabet of two processes, then these processes can participate in actions that synchronize on e. When processes synchronize on an event, they simultaneously participate in actions over that event. During synchronization, an action can *offer* one or more values that can be *observed* by other actions participating in the same event.

Complex processes may be built by either combining sub-processes through an ordering operator (e.g., process P is the sequential composition of sub-processes P_1 and P_2) or by conjoining sub-processes so that they run independently but synchronize on a set of named events. In LOTOS, these synchronizing events must be specified alongside the conjunction operator ($\|$) in a process definition. For brevity, however, we use the conjunction operator without naming the event set. In our abbreviated LOTOS notation, the conjunction of P_1 and P_2 is written $P_1 \parallel P_2$.

Alexander uses conjunction to compose separately defined application and presentation agents [2]. Abowd uses agent-based separation to illuminate usability properties of interactive systems [1]. Both of these approaches rely on the use of conjunction to compose agents that are defined separately but that influence each other. In fact, conjunction is a general operator for composing partial specifications of a system [19]. The idea is that each partial specification imposes constraints upon variables (or, in the case of agents, events) that are mentioned in other partial specifications. When these specifications are conjoined, the common variables must satisfy each constraint.

We define the behavior of a system generated from MASTERMIND models to be any behavior that is consistent with the conjunction of constraints imposed by the dialogue, presentation, and application models. Our results extend conjunction from a specification tool into a mechanism for composing run-time modules.

2.4 Summary

Three issues that must be addressed to solve the model-composition problem: The solution must generate user-interfaces with rich dynamic behavior; the correctness of module composition must be demonstrated; and the generated modules must cooperate efficiently. In MASTERMIND, the rich expressive power is achieved through special-purpose modeling notations [12, 5]. The remainder of this paper addresses the generation of correct implementations with maximal efficiency while preserving the expressive power of MASTERMIND models.

3 Design requirements

Recall from **Figure 1** that each class of model has a code generator that synthesizes run-time modules for models in that class. The modules are generated without detailed knowledge of the other models. At run time, however, modules must cooperate as prescribed by the conjunction of the models that generated them. In this section, we present a detailed specification of the relationship between model composition and how the associated modules cooperate at run-time.

3.1 Notation

The subject of this paper is the automatic generation and composition of run-time modules from design-time models. A *module* is a unit of code generated from a single model. We use a third class of construct–the LOTOS process–to define composition correctness. In formal correctness arguments, we often refer to all three types of constructs and distinguish them by using different fonts. MASTERMIND models are written in the Sans Serif font (e.g., Presentation, Dialogue, and Application). LOTOS processes are written in capital italic letters (e.g., P, D, and A, respectively). Run-time modules are written in German letters (e.g., $\mathfrak{P}$, $\mathfrak{D}$, and $\mathfrak{A}$, respectively).

We now briefly define some process algebra notation. Suppose the behavior of an agent can be described by a LOTOS process B. If the agent can perform an action by synchronizing on event e (denoted $B(e)$), then its behavior from that point is be defined by another process $B' = B(e)$. The systems under study in this paper are deterministic, which means that $B(e)$ is always unique. Moreover, when processes are constructed as the conjunction of sub-processes, the structure is preserved through synchronization. That is, if $B \mathrel{\widehat{=}} B_1 \parallel \ldots \parallel B_n$, then $B(e) \mathrel{\widehat{=}} B'_1 \parallel \ldots \parallel B'_n$ where:

$$B'_i = \begin{cases} B_i(e) & \text{if } e \in \alpha(B_i) \\ B_i & \text{otherwise} \end{cases}$$

Any event that can be observed of a process P can be observed of any conjunction of P with other processes. This observation will be important when we define the Ω observer function (§ 3.4).

We will need to represent the architectural embedding of **Figure 1** equationally. In this case, we refer to the entire run-time system, $\mathfrak{U}$, as a module composed from $\mathfrak{P}$, $\mathfrak{D}$, and $\mathfrak{A}$, using the notation $\mathfrak{U}[\mathfrak{P}, \mathfrak{D}, \mathfrak{A}]$.

3.2 Inter-model composition

Model-based code generators construct a run-time module from a design-time model. The code generation strategy is model-specific, reflecting the specialization of models to a particular aspect of the system. At run time, however, modules must cooperate, and the cooperative behavior must not violate any of the constraints imposed by the models. There is an inherent distinction between behavior that is limited to the confines of a given model and behavior that affects or is affected by other models. Inter-model composition is concerned with managing this latter inter-model behavior.

Some behavior is highly model specific and neither influences nor is affected by behavior specified in other models. In the presentation model, for example, objects are implemented using graphical primitives in the Amulet toolkit [10], and attribute relations are implemented as declarative formulas that, at run-time, eagerly propagate attribute changes to dependent attributes. As long as changes in these attributes do not trigger behavior in the dialogue or application models, these aspects can be ignored when considering model composition. In the application model, object specifications are compiled into abstract classes under the assumption that the designer will later extend these

into subclasses and provide implementations for the abstract methods. As long as the details of these designer extensions do not trigger behavior in dialogue or presentation models, this behavior may also be ignored when defining model composition.

Within a module, entities compose according to a model-specific policy. In the presentation model, for example, objects compose by part-whole aggregation, and attributes compose by formula evaluation over dependent attributes. In the application model, objects compose using a combination of subclassing, aggregation, and polymorphism. When considering how models compose, some details of *intra-model* composition can be abstracted away, but not all of them. Models impose temporal sequencing constraints on the occurrence of inter-model actions, and models contribute to the values computed over the entire system. These constraints and contributions must be captured in some form and used to reason about model composition.

We chose to map this *inter-model* behavior into a semantic domain that is common across all of the models. This domain is described by the LOTOS notation, which specifies temporal constraints on actions and data values. We assume that LOTOS processes can be derived from the text of a model specification (§ 3.4). Designers may, for example, need to designate actions of interest to other models. LOTOS processes do not capture all of the behavior of models in composition, but they do express the essential constraining behavior.

3.3 Example

We now demonstrate an example of inter-model behavior expressed as a LOTOS process. The dialogue model being considered is for a Print/Save widget similar to those found in the user interfaces of drawing tools, web browsers, and word processors. Such widgets allow the user to print a document either to a printer or to a file on disk; we call the former option *printing* and the latter option *saving*. Options specific to printing, such as print orientation (e.g., portrait vs. landscape), and to saving, such as the file to save in, are typically enabled and disabled depending upon the user's choice of task. These ordering dependencies are reflected in the dialogue model for this widget. The inter-model behavior of this dialogue model can be described by the LOTOS process in **Figure 2**.

```
process PrintSave[ print, save, go, cancel, layout, kbd      (1)
                   lpr, write]                               (2)
    (lpdhost, filename : string,                             (3)
     doc               : doctype,                            (4)
     port              : bool) : exit :=                     (5)
  P[go, lpr, write, layout, kbd] [> (cancel; exit)           (6)
where                                                        (7)
  process P[go, lpr, write, layout, kbd] : exit :=           (8)
      Layout[go, lpr, layout]                                (9)
    [> (save; F[go, lpr, write, layout, kbd])                (10)
  endproc                                                    (11)
  process F[go, lpr, write, layout, kbd] : exit :=           (12)
      Edit[go, write, kbd]                                   (13)
    [> (print; P[go, lpr, write, layout, kbd])               (14)
  endproc                                                    (15)
  process Layout[go, lpr, layout] : exit :=                  (16)
       (layout ? port; Layout[go, lpr, layout])              (17)
    [] (go; lpr !lpdhost!port!doc;  exit)                    (18)
  endproc                                                    (19)
  process Edit[go, write, kbd] : exit :=                     (20)
       (kbd ? filename;  Edit[go, write, kbd])               (21)
    [] (go; write ! doc ! file; exit))                       (22)
  endproc                                                    (23)
endproc                                                      (24)
```

Figure 2: Print/save dialogue process.

The process *PrintSave* can synchronize on any of the events that follow in square brackets. In this example, the events *print*, *save*, *go*, *cancel*, *layout*, and *kbd* (line 1 in the figure) define points for synchronizing with the presentation; whereas the events *lpr* and *write* (line 2) define points for synchronizing with the underlying application. The parameters *lpdhost* and *filename* (line 3) store the name of the default printer and the user-selected filename respectively. The parameter *doc* (line 4) represents the document to be printed or saved, and the parameter *port* (line 5) represents the print orientation (portrait if true, landscape if false).

A separate presentation model defines buttons labeled **Send to Printer**, **Save to File**, **Print**, and **Cancel** which, when pressed, offer the events *print*, *save*, *go*, and *cancel* respectively. The presentation model also contains a pair of radio buttons that specify paper orientation. These buttons display graphics of a page in either portrait or landscape mode and, when selected, offer the event *port* with a value of *true* if the choice is for portrait orientation and *false* for landscape orientation. Finally, there is a text entry box in which the user can type in a file name. As the user edits

this name, the text box responds by offering the *kbd* event parameterized by the contents of the string typed so far. Note that the actual keys being pressed are not returned, as editing functionality is best handled in a text widget and is not what we would consider inter-model behavior. A separate application model defines procedures for issuing a print request and saving a file to disk. These procedures are responsive to the events *lpr* and *write* respectively. Actions that synchronize on these events offer a number of values including printer name (*lpdhost*) and filename (*filename*).

The temporal structure of dialogue, presentation, and application model composition is given in the behavior specification (line 6). The behavior of *PrintSave* is the behavior of the process *P* (defined on lines 8 through 11) with the caveat that it may be disabled (terminated) at any time by the observation of the *cancel* event. Process *P* represents what interactions and application invocations must happen in order to send a document to the printer. Most of this functionality is actually expressed in the sub-process *Layout* (defined on lines 16-19). *P* behaves like *Layout* in the normal case, but it can be disabled if the *save* event is observed. Recall that the *save* event is offered whenever the user alternates from the **Send to Printer** button to the **Save to File** button in the presentation model. The process *F* (defined on lines 12-15) likewise behaves like the process *Edit* (defined on lines 20-23) in the normal case, but is disabled if the event *print* is observed. Note that *F* and *P* are mutually disabling, which means that the user can switch back and forth between printing and saving as many times as he or she likes until hitting the **Go** button.

3.4 Models, modules, and processes

Processes like those in **Figure 2** are useful for understanding the relationship between models and modules. This relationship is complex, and so we describe it first for a single model and then for the three models in composition. We now formalize correctness conditions for the MASTERMIND dialogue model. A similar formalization exists for the other MASTERMIND models.

Figure 3 shows the relationship between dialogue models (members of the set $\mathsf{Dialogue}$), run-time modules generated by dialogue models (members of the set $\mathfrak{D}$), and the inter-model behavior of dialogue models (members of the set *Process*). The relationships between these sets are defined as functions that map members of one set into an-

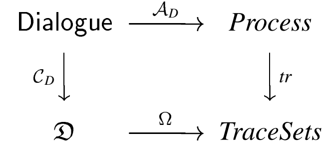

Figure 3: Dialogue compiler correctness.

other. The function $\mathcal{C}_D : \mathsf{Dialogue} \rightarrow \mathfrak{D}$ maps dialogue models to run-time implementation modules. Think of $\mathcal{C}_D$ as an abstract description of the dialogue model compiler. The function $\mathcal{A}_D : \mathsf{Dialogue} \rightarrow \mathit{Process}$ maps dialogue models into LOTOS processes describing their inter-model behavior. Think of $\mathcal{A}_D$ as an abstract interpretation of the dialogue model. The function $\Omega : \mathfrak{D} \rightarrow \mathit{TraceSets}$ maps run-time implementation modules (of any kind) into event traces of their observable behavior. Think of Ω as an observer of run-time behavior. Finally the function $tr : \mathit{Process} \rightarrow \mathit{TraceSets}$ maps a LOTOS process to the set of all possible action traces that can be observed of that process.

These sets and functions are related by the commutative diagram of **Figure 3**. Externally observable model behavior is mapped into a LOTOS process by $\mathcal{A}_D$, and the set of traces of a module's externally observable actions is recorded by Ω. We say that a dialogue model $d \in \mathsf{Dialogue}$ is consistent with the module $\mathcal{C}_D(d)$ if every trace $\sigma \in \Omega(\mathcal{C}_D(d))$ is in the set $tr(\mathcal{A}_D(d))$ and if there are no sequences $\varsigma \in tr(\mathcal{A}_D(d))$ such that $\varsigma \notin \Omega(\mathcal{C}_D(d))$. That is, the inter-model behavioral interpretation of d agrees exactly with the observable behavior of the run-time module generated from d. Commutativity of the diagram requires this property for any dialogue model expressible in the set $\mathsf{Dialogue}$.

3.5 Model-based synthesis

The correctness relationship between models and modules (**Figure 3**) can be extended to specify the correctness of module composition. We now have functions $\mathcal{A}_P$, $\mathcal{A}_D$, and $\mathcal{A}_A$ that map models into LOTOS processes. These processes should compose by conjunction. We also have a run-time module combinator $\mathfrak{U}$ that combines modules from $\mathfrak{P}$, $\mathfrak{D}$, and $\mathfrak{A}$ into a single module whose actions are observable by the Ω function. **Figure 4** shows the constraints on the behavior of these entities. Let $p \in$

$$
\begin{array}{l}
\forall p \in \textsf{Presentation} \\
\quad \forall d \in \textsf{Dialogue} \\
\quad\quad \forall a \in \textsf{Application} \\
\quad\quad\quad \Omega(\mathfrak{U}[\mathcal{C}_P(p), \mathcal{C}_D(d), \mathcal{C}_A(a)]) \\
\quad\quad\quad\quad = \\
\quad\quad\quad tr(\mathcal{A}_P(p) \parallel \mathcal{A}_D(d) \parallel \mathcal{A}_A(a))
\end{array}
$$

Figure 4: Module-composition correctness.

Presentation, $d \in$ Dialogue, and $a \in$ Application. Then the code generated from these models is correct if, for any observable behavior σ, σ is a legal trace in the conjunction of the models and vice-versa. This equation defines the conditions necessary for correct module composition without assuming any model-specific interpretation of these actions. It serves, therefore, as a specification of design requirements. In the next section, we present an implementation that satisfies these requirements.

4 Design

We now turn to the design of the run-time synchronization module and model-specific compilers of **Figure 1**. The correctness conditions of **Figure 4** impose constraints on these designs. Fortunately, these constraints do not require model-specific knowledge (e.g., graphical concepts in the presentation model or data layout in the application model). This allows us to design a generic infra-structure of inter-model cooperation and to assume this design when crafting model-specific code generation strategies. The design refines the notions of action and synchronization, which form the basis of inter-module communication, into run-time objects that implement these constraints.

4.1 Run-time control

One concern in designing a system is the implementation of software control [15]. Control can be implemented in many ways. In procedural systems, for example, control is synonymous with location in the code; whereas in concurrent systems, control is distributed and managed by multiple objects concurrently. User-interface software generally implements an event-driven, sequential control scheme, in which a single thread provides a facade of concurrency by dispatching small callback routines when input device activity is sensed. In the interest of providing a single style of control in our systems, we adopt the event-driven, sequential implementation.

The choice of software control implementation influences the design of actions and synchronization. Actions in the presentation module correspond to input device behavior like mouse and keyboard events. Since these events invoke callback procedures, we implement synchronization as a callback. This means that the temporal structure of the inter-model behavior must be implemented in such a way that all legal actions are enabled, all illegal actions are disabled, and after action synchronization, new actions are enabled or disabled. If one model can be made to represent this temporal structure, then functionality provided by the other models can be abstracted into context-independent actions and implemented using method callbacks.

By design, the temporal structure of the dialogue model represents the synchronization needs of the entire program. This makes it natural to treat the dialogue module as the arbiter of system control. In the architecture presented in **Figure 1**, the **Dialogue** module is a reactive component that computes the enabled/disabled status of actions in response to action synchronizations, and the other modules are collections of code that are invoked when actions synchronize. At run time, every action causes the dialogue module to compute the next state of the system. Based on this next state, actions embedded in other modules are enabled, disabled, or activated as appropriate.

4.2 Action synchronization

The dialogue module computes the set of enabled actions as a function of the observed actions. Actions can be thought of as entities that are enabled, disabled, and activated by an omniscient dialogue agent, and the model-specific interpretation of said actions can be structured to occur when the action is activated. `Action` objects are run-time entities that encapsulate the status (enabled or disabled) of observable actions with activation procedures that can be specialized by model-specific code generators to implement desired functionality. **Figure 5** shows our design as an OMT [15] object model.

The class `Action` in our design is abstract: It declares an operation `enable()` that must be supplied by a subclass. OMT denotes subclassing by a triangle, one point of which is connected to the superclass, with one or more lines emanating out to its subclasses. In MASTERMIND-

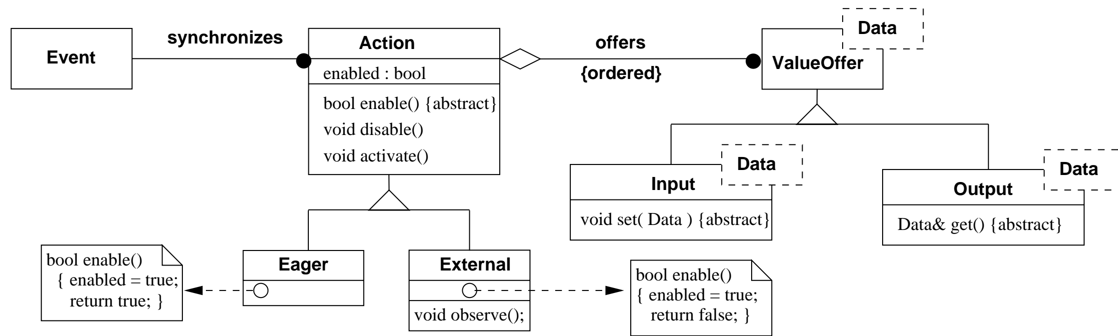

Figure 5: Object model for synchronization mechanism.

generated code, presentation model actions are associated with presentation module interaction objects, and application model actions are associated with the invocation of methods of objects in the application module. When a presentation model action can be offered, the graphical object associated with that action is enabled and made ready to accept user activity. When the graphical object detects such activity, it signals the rest of the system than an action synchronization is occurring. When an action in the application model can be offered, a method in the application module is invoked. The object model of **Figure 5** distinguishes between these interpretations of action enabling by subclassing `Action` into those that are `External` and those that are `Eager`. The synchronization requirements of an `Eager` action are met when the action is enabled; whereas `External` actions require both being enabled and observing activity generated by an external entity like a mouse. Such activity is posted by issuing the `observe()` message.

Class `Action` has a `synchronizes` association with class `Event`. Objects of class `Event` represent process events upon which multiple actions synchronize. `Event` objects encapsulate a unique name with the synchronization requirements of actions from multiple models. While there is an object of class *Action* for every action in any model, there is only one object of class `Event` for any distinct event.

LOTOS actions are often accompanied by one or more value inputs or outputs. We designed class `Action` to aggregate zero or more objects of the parameterized class `ValueOffer`. OMT denotes aggregation with the diamond operator. Zero or more objects of class `ValueOffer` are parts of every object of class `Action`. The parameter `Data` in **Figure 5** names a data type that *parameterizes* class `ValueOffer`. A parameterized class (denoted as a box with a dashed box in the upper right corner) can define local attributes or operations whose signatures vary with the parameter. Subclasses `Input` and `Output` use this parameter value to specialize `set` and `get` operations. These abstract operations must be supplied by model-specific code generators, which know how to supply and receive values in model-specific contexts.

LOTOS actions have the following syntax:

$$\begin{aligned} &\mathit{action} ::= \mathit{EventName}\ (\mathit{input} \mid \mathit{output})* \\ &\mathit{input}\ \ ::= '?'\ \mathit{Variable}\ ':'\ \ \mathit{Type}; \\ &\mathit{output} ::= '!'\ \mathit{Expr}; \end{aligned}$$

The only information a code generator has about an action is the event name, whether the value offers are inputs or outputs, the name and type of the variable in which to store the input, and the expression used to compute the output. This information sufficiently denotes an action object in our framework.

4.3 Run-time execution

`Event` objects internalize synchronization requirements of multiple actions and issue `activate` and `get/set()` messages to `Action` and `ValueOffer` objects as callbacks. To make this work, `Event` objects contain a pointer to all of the `Action` objects that synchronize and vice-versa. When an `Action` object is enabled by the dialogue module, the return value (true or false) is recorded in the corresponding `Event` object so that the synchronization requirements can be tabulated. A return value of false indicates that an action is `External`, in

which case the `Event` object records that the action is enabled but waits for external confirmation that the action has been chosen by the user. Once all of the synchronization requirements have been met, the `Event` issues the appropriate `activate`, `get`, and `set` operations and then instructs the dialogue module to compute the next state. This process is described in greater detail in [16].

5 Results and status

We evaluated our solution to the model-composition problem with respect to power, correctness, and efficiency.

Power We were able to express user interfaces in several case studies using our modeling notations. We tested the quality of user interfaces on two specific examples: the Print/Save widget described in § 3.3 and an airspace and runway executive that supports an air-traffic controller (ATC) [16]. The former demonstrates the ability to generate common, highly reusable tasks for standard graphical user interfaces. The latter demonstrates the ability to support a complex task using a direct-manipulation interface.

The ATC example testifies to the power of our approach. When flight numbers are keyed in to a text-entry box, an airplane graphic, augmented with the flight number, appears in the airspace. As more planes come into the airspace, the controller keys their flight number in a text-entry box. When the controller decides to change the position of a plane, he does so by dragging the airplane graphic to a new location on the canvas. As soon as he presses and holds the mouse button, a feedback object shaped like an airplane appears and follows the mouse to the new location. When the mouse is released, the plane icon moves to the newly selected location.

The presentation model of the ATC example is quite rich. It specifies gridding so that airplane graphics are always uniformly placed within the lanes, and it specifies feedback objects that give users information during an operation. In a real deployment, the location of the flights would probably change in response to asynchronous application signals from special hardware monitors. In such a deployment, these signals would be connected to `External` actions and would fit into the framework without change. For more details on this case study and the print/save dialogue, see Stirewalt [16].

Correctness In addition to being able to generate and manage powerful user interfaces, the composition of our modules is correct. Two aspects of our approach require justification on these grounds. First is the design of runtime action synchronization. This paper addresses the theoretical issues involved here. In practice, we have found the design to be quite robust. Second is the synthesis of the runtime dialogue component (member of the set $\mathfrak{D}$) from a dialogue model. As we mentioned earlier, the MASTERMIND Dialogue model notation can be thought of as a syntactic sugaring for a subset of Full LOTOS. We implemented a prototype dialogue model code generator whose correctness was validated in Stirewalt [16].

Efficiency We measured efficiency empirically by applying our prototype code generator on the ATC example. We generated dialogue modules and connected these with hand-coded presentation and application modules. On the examples we tried, we observed no time delays between interactions. We quantified these results by instrumenting the source code to measure the use of computation resources and wall-clock time. The maximum time taken during any interaction was 0.04 seconds. This compares well to the *de facto* HCI benchmark of response time, which is 0.1 seconds. We believe that more heavyweight, middle-ware solutions, such as implementing synchronization through object-request brokers, are not competitive with these results.

We are currently completing a new industrial-strength, dialogue code generator. This new code generator is incorporating state-space reduction technology and will improve interaction time that, in the prototype, is a function of the depth of a dialogue expression with constant time interaction. We are also working on adapting the presentation model code generator described in [6] to work within our infra-structure.

6 Conclusions

Generating code for a specialized modeling notation is easy. Integrating code generated from multiple models is difficult. Integration is much more complicated than mere linking of compiled object modules. For models to be declarative, they must assume that entities named in one

model have behavior that is elaborated in another model. Designers want to treat presentation, temporal context, and effect separately because each aspect in isolation can be expressed in a highly specialized language that would be less clear if it were required to express the other aspects as well. For interactive systems, composition by conjunction is essential to separating complex specifications into manageable pieces.

Unfortunately, programming languages like C++ and Java do not provide a conjunction operator. Such an operator is difficult to implement correctly and efficiently, and in fact, we did not try to implement it. Rather, by casting model composition into a formal framework that includes conjunction, we were able to express a correct solution and then refine the correct solution into an efficient design. This is a key difference between our approach and middle-ware solutions that implement object composition by general event registry and callback.

Our results contribute to the body of automated software engineering research in two ways. First, our framework is a practical solution that helps to automate the engineering of interactive systems. Second, our use of formal methods to identify design constraints and the subsequent refinement of these constraints into an object-oriented design may serve as a model for other researchers trying to deal with model composition in the context of code generation. The formality of the approach allowed us to minimize design constraints and was the key to arriving at a powerful, correct, and efficient solution.

References

[1] G. D. Abowd. *Formal Aspects of Human-Computer Interaction*. PhD thesis, University of Oxford, 1991.

[2] H. Alexander. Structuring dialogues using CSP. In M. Harrison and H. Thimbleby, editors, *Formal Methods in Human-Computer Interaction*. Cambridge University Press, 1990.

[3] L. Bass and J. Coutaz. *Developing Software for the User Interface*. SEI Series in Software Engineering. Addison-Wesley, 1991.

[4] T. Bolognesi and E. Brinksma. Introduction to the ISO specification language LOTOS. *Computer Network ISDN Systems*, 14(1), 1987.

[5] T. P. Browne et al. Using declarative descriptions to model user interfaces with MASTERMIND. In F. Paternò and P. Palanque, editors, *Formal Methods in Human Computer Interaction*. Springer-Verlag, 1997.

[6] P. Castells, P. Szekely, and E. Salcher. Declarative models of presentation. In *IUI'97: International Conference on Intelligent User Interfaces*, pages 137–144, 1997.

[7] J. Coutaz. PAC, an object-oriented model for dialog design. In *Human Computer Interaction - INTERACT'87*, pages 431–436, 1987.

[8] D. Harel. On visual formalisms. *Communications of the ACM*, 31(5), 1988.

[9] G. E. Krasner and S. T. Pope. A cookbook for using the model view controller user interface paradigm in smalltalk. *Journal of Object Oriented Programming*, 1(3), 1988.

[10] B. A. Myers et al. The Amulet environment: New models for effective user-interface software development. *IEEE Transactions on Software Engineering*, 23(6), 1997.

[11] B. A. Myers and M. B. Rosson. Survey on user interface programming. In *SIGCHI'92: Human Factors in Computing Systems*, May 1992.

[12] R. Neches et al. Knowledgeable development environments using shared design models. In *Intelligent Interfaces Workshop*, pages 63–70, 1993.

[13] P. Palanque, R. Bastide, and V. Sengès. Validating interactive system design through the verification of formal task and system models. In *Working Conference on Engineering for Human Computer Interaction*, 1995.

[14] A. Puerta. The Mecano project: Comprehensive and integrated support for model-based user interface development. In *Computer-Aided Design of User Interfaces*, 1996.

[15] J. Rumbaugh et al. *Object-Oriented Modeling and Design*. Prentice-Hall, 1991.

[16] R. E. K. Stirewalt. *Automatic Generation of Interactive Systems from Declarative Models*. PhD thesis, Georgia Institute of Technology, 1997.

[17] P. Szekely et al. Declarative models for user-interface construction tools: the MASTERMIND approach. In Bass and Unger, editors, *Engineering for Human-Computer Interaction*. Chapman & Hall, 1996.

[18] Pedro Szekely, Ping Luo, and Robert Neches. Beyond interface builders: Model-based interface tools. In *Bridges Between Worlds: Human Factors in Computing Systems: INTERCHI'93*, 1993.

[19] P. Zave and M. Jackson. Conjunction as composition. *ACM Transactions on Software Engineering and Methodology*, 2(4):371–411, 1993.