

FSATS: an Extensible Fire Support Simulator¹

Don Batory
Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712
batory@cs.utexas.edu

Clay Johnson, Bob MacDonald, Dale von Heeder
Applied Research Labs
University of Texas at Austin
Austin, Texas 78713
{clay, bob, drv}@arlut.utexas.edu

Abstract

This is a case study in the use of *product-line architectures (PLAs)* and *domain-specific languages (DSLs)* to design an extensible command-and-control simulator for Army fire support. The reusable components of our PLA are layers or “aspects” whose addition or removal simultaneously impacts the source code of multiple objects and multiple, distributed programs. The complexity of our component specifications is substantially reduced by using a DSL for defining and refining state machines, abstractions that are fundamental to simulators. We present preliminary results that show how our PLA and DSL synergistically produce a more flexible way of implementing state-machine-based simulators than is possible with a pure Java implementation.

Keywords: product-line architecture, domain-specific languages, GenVoca, aspects, state diagrams.

1 Introduction

Software evolution is a costly yet unavoidable consequence of a successful application. Evolution occurs when new features are added and existing capabilities are enhanced. Unfortunately, many applications suffer *design fatigue* — when further evolution is difficult and costly because of issues not addressed in the initial design [Gra97]. Software that is easily evolvable is a central problem today in software engineering.

There are three complimentary technologies that address software evolution: object-oriented design patterns, domain-specific languages, and product-line architectures. *Design patterns* are techniques for restructuring and generalizing object-oriented software [Gam95]. Evolution occurs by applying design patterns to an existing design; the effects of these changes are borne by programmers to manually transform an existing code base to match the updated design. Recent advances indicate that tool support for automating the applications of patterns is possible [Tok99]. *Domain-specific languages (DSLs)* raise the level of programming to allow customized applications to be specified compactly in terms of domain concepts; compilers translate DSL specifications into source code. Evolution is achieved by evolving DSL specifications [Deu97]. *Product-line architectures (PLAs)* are designs for families of related applications; application construction is accomplished by composing reusable components. Evolution occurs by plugging and unplugging components that encapsulate new and enhanced features [Bat98b, Bos99, Cza99, SEI99, Wei99]. Among PLA models, the GenVoca model is distinguished by an integration of ideas from aspect-oriented programming [Kic97], parameterized programming [Gog86], and program-construction by refinement [Bax92].

This paper presents a case study in the use GenVoca PLAs and DSLs to create an extensible command-and-control simulator for Army fire support. (Design patterns were also used, but they played a minor

1. This work was supported in part by the University of Texas Applied Research Labs and the U.S. Department of Defense Advanced Research Projects Agency in cooperation with the U.S. Wright Laboratory Avionics Directorate under contract F33615-91C-1788.

role). We discovered that components of distributed simulations are *not* conventional DCOM and CORBA components, but rather are layers or “aspects” whose addition or removal simultaneously impacts the source code of *multiple*, distributed programs. Further, we found that writing our components in a general-purpose programming language (Java) resulted in complex code that obscured a relatively simple, state-machine-based design. By extending Java with domain-specific abstractions (in our case, state-diagrams), our component specifications were readily understood by domain experts and knowledge engineers as well as application programmers. Thus, this case study is interesting not only because of the novelties introduced by PLAs and DSLs, but also their integration: using only one technology would be inadequate.

We begin by explaining the ideas and terminology of fire support. We review an existing simulator, called FSATS, and motivate its redesign. We present a GenVoca PLA for creating extensible fire-support simulators and introduce an extension to the Java language for defining and refining state-diagrams. Finally, based on simple measures of program complexity, we show how PLAs and DSLs individually simplify simulators, but only their combination provides practical extensibility.

2 Background

2.1 The Domain of Fire Support

Fire support is a command-and-control application that includes the detection of targets, assignment of weapons to attack the target, and coordination of the actual attack. The entities engaged in this process, called *operational facilities (OPFACs)*, are command posts that exchange tactical (theater-of-war) messages.

Forward observers (FO) are OPFACs that are stationed at intervals across the front-line of a battlefield (Figure 1). They are one of several kinds of sensors responsible for detecting potential targets. A hierarchy of *fire support elements (FSE)* is responsible for directing requests from FOs to the most appropriate weapon system to handle the attack. FOs report to their *fire support team (FIST)*; a FIST reports to a battalion FSE, a battalion FSE reports to a brigade FSE, and so on. Each FSE typically has one or more supporting *command posts (CPs)* with different weapon systems. For example, a brigade FSE might be supported by a *field artillery command post (FACP)*; a battalion FSE might be supported a mortar command post, and so on. In general, higher echelon FSEs are supported by higher echelon CPs with more powerful and/or longer range weapon systems.

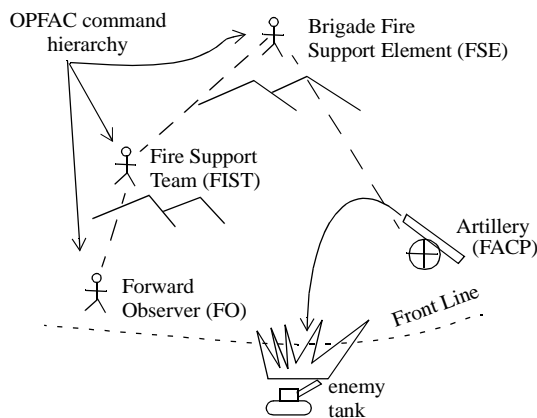


Figure 1: OPFAC Command Hierarchy

FOs, FISTs, and other FSEs are responsible for evaluating a target. An evaluation may result in (a) assigning the target to be attacked by a supporting weapon, (b) elevating the target to the next higher echelon FSE

for evaluation, or (c) denial — choosing not to attack the target. CPs are responsible for assigning targets to the best weapon or combination of weapons under their command. Once weapon(s) are assigned, messages are exchanged with the mission originator (usually an FO) to coordinate completion of the mission. The particular message sequence depends on the target and weapon. It is still generally the case that all messages are relayed along the chain of CPs and FSEs that were involved in initiating the mission, although newer systems permit messages to be exchanged directly between the weapon and observer. The message sequence for a particular mission is referred to as the *mission thread*. In general, an OPFAC can participate in any number of mission threads at a time.

A mission thread is an instance of a *mission type*. There are well over twenty mission types, including:

- *when-ready-fire-for-effect-mortars (WRFFE-mortars)* — a mortar CP is assigned to shoot at a target as soon as possible,
- *when-ready-fire-for-effect-artillery (WRFFE-artillery)* — one or more artillery CPs are assigned to shoot at a target as soon as possible,
- *when-ready-adjust-mortars (WRAdjust-mortars)* — a forward observer knows only approximately the location of the enemy and requests single rounds to be fired with the observer sending corrections between rounds until the target is hit, at which point it becomes a WRFFE-mortar mission, and
- *time-on-target-artillery (TOT-artillery)* — field artillery are requested to fire at a target so that all rounds land at the specified location at the specified time.

Each OPFAC (FO, FIST, FSE, etc.) performs different actions for each mission type. For example, the actions taken by a FO for a TOT-mortar mission are different than those for a WRFFE-artillery mission.

Clearly, the above description of fire support is highly simplified, e.g., the actions taken by specific OPFACs in a mission thread and the translation of messages into formats for tactical transmission were omitted. However, these details are unnecessary to understand contributions of this paper.

FSATS. Simulation plays a key role in U.S. Army testing and training. It avoids costs of mobilizing live forces, provides repeatability in testing, and allows force-on-force combat training without the liability. Simulation has been used to model virtual environments, weapons effects, outcome adjudication, and as computational resources increase, the fidelity has been refined to entity-level simulators.

Fire support is one of a number of domains that has been modernized by digital *Command, Control, Communications, Computer, and Intelligence (C4I)* systems that automate battlefield mission processing. *AFATDS (Advanced Field Artillery Tactical Data System)* is arguably the most sophisticated C4I system in existence, and provides the software backbone (message transmission, processing, etc.) for fire support for the Army [Mag99]. *FSATS (Fire Support Automated Test System)* is a system for testing AFATDS and other fire-support C4I systems. FSATS collects digital message traffic from command and control communication networks, interprets these messages, and stores them in a database for later analysis. FSATS can simulate any or all OPFACs used in AFATDS [ARL99]. The subject of a test can be overall system performance, individual OPFAC performance, or system operator performance. Thus, FSATS is used both in training Army personnel in fire support and debugging/testing AFATDS.

2.2 The Current FSATS Implementation

FSATS has been under development for almost ten years. It is typical of the systems mentioned in our introductory paragraph: it began with a clean design but as its capabilities were extended, limitations of that design became increasingly troublesome.

The implementation is a combination of decision rules encoded in database tables, a set of “common actions” written as Ada procedures, and a decision rule interpreter, also in Ada. One set of rules is associated with each (OPFAC type, message type) pair. When a tactical message is received by an OPFAC, the appropriate rule set is selected by the interpreter and each rule in the set is sequentially evaluated until one succeeds, at which point the action for that rule is executed and processing of that message terminates. There are from 200 to 1000 rules associated with each OPFAC type, divided among the various input message types. Each rule consists of a predicate, which is a conjunction of guards, and an action which is an index to a sequence of state and message common actions. Predicates typically contain five to ten guards (terms). The processing of rule sets is optimized, so that predicates can assume the failure of all previous predicates. Common actions range from simple (copy the target number field from the input to the output message) to complex (does there exist a supporting OPFAC of type mortar which is capable of shooting the target indicated by the current message?).

There are now obvious drawbacks to this design/implementation. While rule sets are used to express OPFAC behavior, OPFAC behavior is routinely understood and analyzed as mission threads. Figure 2 illustrates a mission thread, the horizontal execution path, that associates various rules spanning multiple OPFAC programs. This complicates the knowledge acquisition and engineering process to derive from an analysis of multiple mission threads the rules as they apply at each OPFAC. Conversely, it obfuscates analyzing and debugging system behavior where rules for multiple mission threads are merged into monolithic sets within each OPFAC program.




Figure 2: Rule Sets vs. Mission Threads

The contrast of the vertical nature of rule sets versus the horizontal or “cross-cutting” nature of mission threads in Figure 2 illustrates an encapsulation dichotomy that is not unique to FSATS [Ree92, Kic97]. In general, conventional OO approaches explore *use cases* (threads) for specification and analysis of system behavior. However, the concept of a use case is transient in a design process that identifies behavior (rules) with the actors (OPFACs) rather than the actions (missions). This trade-off is seemingly unavoidable given the need to produce objects that combine behaviors to react to a variety of situations. In FSATS, the transformation of mission threads into rule sets yields autonomous OPFACs at an increased cost to analysis and maintenance.

As FSATS evolved, rule sets quickly became large (>100) and unwieldy. Moreover, different missions might use the same message type at an OPFAC for slightly different purposes. Simpler rules that once sufficed often had to be factored to disambiguate their applicability to newer, more specialized missions. In worse cases, large subsets of rules had to be duplicated, resulting in a non-linear increase in rules and interactions. Moreover, the relationship between rules of different OPFACs, and the missions to which they applied, was lost. Modifying OPFAC rules became perilous without laborious analysis to rediscover and reassess those dependencies. The combinatorial effect of rule set interactions made analysis increasingly difficult and time-consuming.

FSATS management realized that the current implementation was not sustainable in the long term, and a new approach was sought. FSATS would continue to evolve through the addition of new mission types and by varying the behavior of an OPFAC or mission to accommodate doctrinal differences over time or

between different branches of the military. Thus, the need for extensible simulators was clearly envisioned. The primary goals of a redesign were to:

- disentangle the logic implementing different mission types to make implementation and testing of a mission independent of existing missions,
- reduce the “conceptual distance” from logic specification to its implementation so that implementations are easily traced back to requirements and verified, and
- allow convenient switching of mission implementations to accommodate requirements from different users and to experiment with new approaches.

2.3 GenVoca

The technology chosen to address problems identified in the first-generation FSATS simulation was the GenVoca methodology implemented using the *Jakarta Tool Suite (JTS)* [Bat98a]. At its core, GenVoca is a design methodology for creating product-lines and building architecturally-extensible software — i.e., software that is extensible via component additions and removals. GenVoca is a scalable outgrowth of an old and practitioner-ignored methodology called *step-wise refinement*, which advocates that efficient programs can be created by revealing implementation details in a progressive manner. Traditional work on step-wise refinement focussed on microscopic program refinements (e.g., $x+0 \rightarrow x$), for which one had to apply hundreds or thousands of refinements to yield admittedly small programs. While the approach is fundamental and industrial infrastructures are on the horizon [Sim95, Bax92], GenVoca extends step-wise refinement largely by scaling refinements to a component or layer (i.e., multi-class-modularization) granularity, so that applications of great complexity can be expressed as a composition of a few large-scale refinements [Bat92, Cza99].

Mixin-Layers. A GenVoca component typically encapsulates multiple classes. Figure 3a depicts component x with four classes A - D . Any number of relationships can exist among these classes; Figure 3a shows only inheritance relationships. That is, B and C are subclasses of A , while D has no inheritance relationship with A - C .




Figure 3: GenVoca Components and their Composition

A subclass is a *refinement* of a class: it adds new data members, methods, and/or overrides existing methods. A *GenVoca refinement* simultaneously refines multiple classes. Figure 3b depicts a GenVoca component y that encapsulates three refining classes (A , B , and D) and an additional class (E). Note that the refining classes (A , B , D) do not have their superclasses specified; this enables them to be “plugged” underneath their yet-to-be-determined superclasses.²

2. More accurately, a *refinement* of class A is a subclass of A with name A . Normally, subclasses must have distinct names from their superclass, but not so here. The idea is to graft on as many refinements to a class as necessary — forming a linear “refinement” chain — to synthesize the actual version of A that is to be used. Subclasses with names distinct from their superclass define entirely new classes (such as B and C above), which can subsequently be refined.

Figure 3c shows the result of composing \mathbf{y} with \mathbf{x} (denoted $\mathbf{y}\langle\mathbf{x}\rangle$). (The classes of \mathbf{y} are outlined in darker ovals to distinguish them from classes of \mathbf{x}). Note that the obvious thing happens to classes **A**, **B**, and **D** of component \mathbf{x} — they are refined by classes in \mathbf{y} as expected. That is, a linear inheritance *refinement chain* is created, with the original definition (from \mathbf{x}) at the top of the chain, and the most recent refinement (from \mathbf{y}) at the bottom. As more components are composed, the inheritance hierarchies that are produced get progressively broader (as new classes are added) and deeper (as existing classes are refined). As a rule, only the bottom-most class of a refinement chain is instantiated and subclassed to form other distinct chains. (These are indicated by the shaded classes of Figure 3c). The reason is that these classes contain all of the “features” or “aspects” that were added by “higher” classes in the chain. These “higher” classes simply represent intermediate derivations of the bottom class [Sma98, Bat98a, Fin98].

Representation. A GenVoca component/refinement is encoded in JTS as a class with nested classes. A representation of component \mathbf{x} of Figure 3a is shown below, where $\$TEqn.A$ denotes the most refined version of class **A** (e.g., classes $\mathbf{x.B}$ and $\mathbf{x.C}$ in Figure 3a have $\$TEqn.A$ as their superclass). We use the Java technique of defining properties via empty interfaces; interface **F** is used to indicate the “type” of component \mathbf{x} :

```
interface F { }           // empty

class X implements F {
  class A { ... }
  class B extends $TEqn.A { ... }
  class C extends $TEqn.A { ... }
  class D { ... }
}
```

Components like \mathbf{y} that encapsulate refinements are expressed as mixins — classes whose superclass is specified via a parameter. A representation of \mathbf{y} is a *mixin-layer* [Sma98, Fin98], where \mathbf{y} ’s parameter can be instantiated by any component that is of “type” **F**:

```
class Y <F S> extends S implements F {
  class A extends S.A { ... }
  class B extends S.B { ... }
  class D extends S.D { ... }
  class E { ... }
}
```

The composition of \mathbf{y} with \mathbf{x} , depicted in Figure 3c, is expressed by:

```
class MyExample extends Y<X>;
```

where $\$TEqn$ is replaced by **MyExample** in the instantiated bodies of \mathbf{x} and \mathbf{y} . Readers familiar with the GenVoca model will recognize that **F** corresponds to a realm interface, \mathbf{x} and \mathbf{y} are components of realm **F**, and **MyExample** is a type equation [Bat92]. Extensibility is achieved by adding and removing mixin-layers from applications; product-line applications are defined by different compositions of mixin layers.

3 The Implementation

The GenVoca-FSATS design was implemented using the *Jakarta Tool Suite (JTS)* [Bat98a], a set of Java-based tools for creating product-line architectures and compilers for extensible Java languages. The following sections outline the essential concepts of our JTS implementation.

3.1 A Design for an Extensible Fire-Support Simulator

The Design. The key idea behind the GenVoca-FSATS design is the encapsulation of individual mission types as components. That is, the central variabilities in FSATS throughout its history (and projected future) lie in the addition, enhancement, and removal of mission types. By encapsulating mission types as components, evolution of FSATS is greatly simplified.

We noted earlier that every mission type has a “cross-cutting effect”, because the addition or removal of a mission type impacts multiple OPFAC programs. A mission type is an example of a more general concept called a *collaboration* — a set of objects that work collectively to achieve a certain goal [Ree92, Van96, Sma98]. Collaborations have the desirable property that they can be defined largely in isolation from other collaborations, thereby simplifying application design. In the case of FSATS, a mission is a collaboration of objects (OPFACs) that work cooperatively to prosecute a particular mission. The actions taken by each OPFAC are defined by a protocol (state diagram) that it follows to do its part in processing a mission thread. Different OPFACs follow different protocols for different mission types.

An extensible, component-based design for FSATS follows directly from these observations. One component (**Basic**) defines an initial OPFAC class hierarchy and routines for sending and receiving messages, routing messages to appropriate missions, reading simulation scripts, etc. Figure 4 depicts the **Basic** component encapsulating multiple classes, one per OPFAC type. The OPFACs that are defined in **Basic** do not know how to react to external stimuli. Such reactions are encapsulated in mission components.

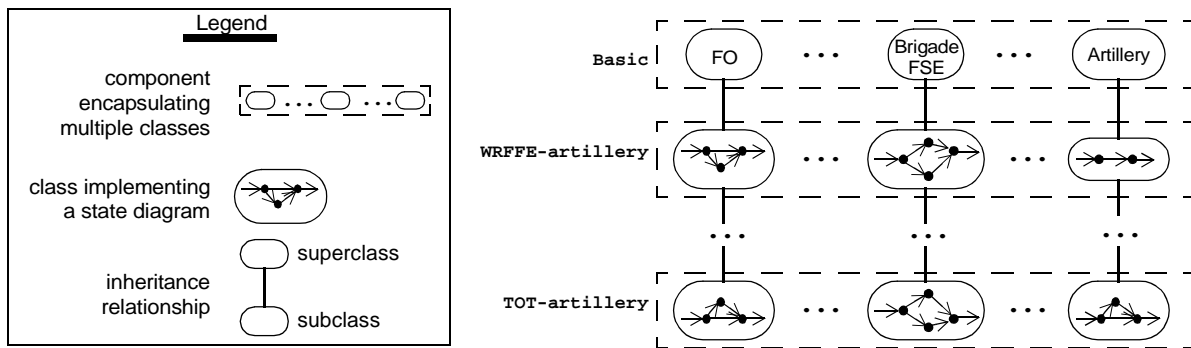


Figure 4: OPFAC Inheritance Refinement Hierarchy

Each mission component encapsulates protocols (expressed as state diagrams) that are added to each OPFAC that could participate in a thread of this mission type. Composing a mission component with **Basic** extends each OPFAC with knowledge of how to react to particular external stimuli and how to coordinate its response with other OPFACs. For example, when the **WRFFE-artillery** component is added, a forward observer now has a protocol that tells it how to react when it sees an enemy tank — it creates a **WRFFE-artillery** message which it relays to its FIST. The FIST commander, in turn, follows his **WRFFE-artillery** protocol to forward this message to his brigade FSE, and so on. Figure 4 depicts the **WRFFE-artillery** component encapsulating multiple classes, again one per OPFAC type. Each enclosed class encapsulates a protocol which is added to its appropriate OPFAC class. Component composition is accomplished via inheritance, and is shown by dark vertical lines between class ovals in Figure 4. The same holds for other mission components (e.g., **TOT-artillery**). Note that the classes that are instantiated are the bottom-most classes of these linear inheritance chains, because they embody all the protocols/features that have been grafted onto each OPFAC. Readers will recognize this is an example of the JTS paradigm of Section 2.3, where components are mixin-layers.

The GenVoca-FSATS design has distinct advantages:

- it is mission-type *extensible* (i.e., it is comparatively easy to add new mission types to an existing GenVoca-FSATS simulator),³
- each mission type is defined largely independent of others, thereby reducing the difficulties of specification, coding, and debugging, and
- understandability is improved: OPFAC behavior is routinely understood and analyzed as mission threads. Mission-type components directly capture this simplicity, avoiding the complications of knowledge acquisition and engineering of rule sets.

Implementation. There are presently 25 different mixin-layer components in GenVoca-FSATS, all of which we are composing now to form a “fully-loaded” simulator. There are individual components for each mission type, just like Figure 4. However, there is no monolithic `Basic` component. We discovered that `Basic` could be decomposed into ten largely independent layers (totalling 97 classes) that deal with different aspects of the FSATS infrastructure. For example, there are distinct components for:

- OPFACs reading from simulation scripts,
- OPFAC communication with local and remote processes,
- OPFAC proxies (objects that are used to evaluate whether OPFAC commanders are supported by desired weapons platforms),
- different weapon OPFACs (e.g., distinct components for mortar, artillery, etc.), and
- GUI displays for graphical depiction of ongoing simulations.

Packaging these capabilities as distinct components both simplifies specifications (because no extraneous details need to be included) and debugging (as components can largely be debugged in isolation). An important feature of our design is that all OPFACs are coded as threads executing within a single Java process. There is an “adaptor” component that refines (or maps) locally executing threads to execute in distributed Java processes [Bat99]. The advantage here is that it is substantially easier to debug layers and mission threads within a single process than debugging remote executions. Furthermore, only when distributed simulations are needed is the adaptor included in an FSATS simulator.

Perspective. It is worth comparing our notion of components with those that are common in today’s software industry. Event-based distributed architectures, where DCOM or CORBA components communicate via message exchanges, is likely to be a dominant architectural paradigm of the future [Tay99]. FSATS is a classic example: OPFAC programs are distributed DCOM/CORBA “components” that exchange messages. Yet the “components” common to distributed architectures are *orthogonal* to the components in the GenVoca-FSATS design. Our components (layers) encapsulate fragments of *many* OPFACs, instead of encapsulating an *individual* OPFAC. (This is typical of approaches based on collaboration-based or “aspect-based” designs).

Event-based architectures are clearly extensible by their ability to add and remove “component” instances (e.g., adding and removing OPFACs from a simulation). This is (*OPFAC*) *object population* extensibility, which FSATS definitely requires. But FSATS also needs *software* extensibility — OPFAC programs must

3. Although a product-line of different FSATS simulators is possible; presently the emphasis of FSATS is on extensibility. It is worth noting, however, that exponentially-large product-lines of FSATS simulators could be synthesized — i.e., if there are m mission components, there can be up to 2^m distinct compositions/simulators.

be mission-type extensible. While these distinctions seem obvious in hind-sight, they were not so prior to our work. FSATS clearly differentiates them.

3.2 A Domain-Specific Language for State Machines

We discovered that OPFAC rule sets were largely representations of state diagrams. We found that expressing OPFAC actions as state diagrams was a *substantial* improvement over rules; they are much easier to explain and understand, and require very little background to comprehend. One of the major goals of the redesign was to minimize the “conceptual distance” between architectural abstractions and their implementation. The problem we faced is that encodings of state diagrams are obscure, and given the situation that our specifications often *refined* previously created diagrams, expressing state diagrams in pure Java code was unattractive. To eliminate these problems, we used JTS to extend Java with a domain-specific language for declaring and refining state machines, so that our informal state diagrams (nodes, edges, etc.) had a direct expression as a formal, compilable document. This extended version of Java is called *JavaSM*.

Initial Declarations. A central idea of JavaSM is that a state diagram specification translates into the definition of a single class. There is a generated variable (`current_state`) whose value indicates the current state of the protocol (i.e., state-diagram-class instance). When a message is received by an OPFAC mission, a designated method is invoked with this message as an argument; depending on the state of the protocol, different transitions occur. Figure 5a shows a simple state diagram with three states and three transitions. When a message arrives in the `start` state, if method `booltest()` is true, the state advances to `stop`; otherwise the next state is `one`.

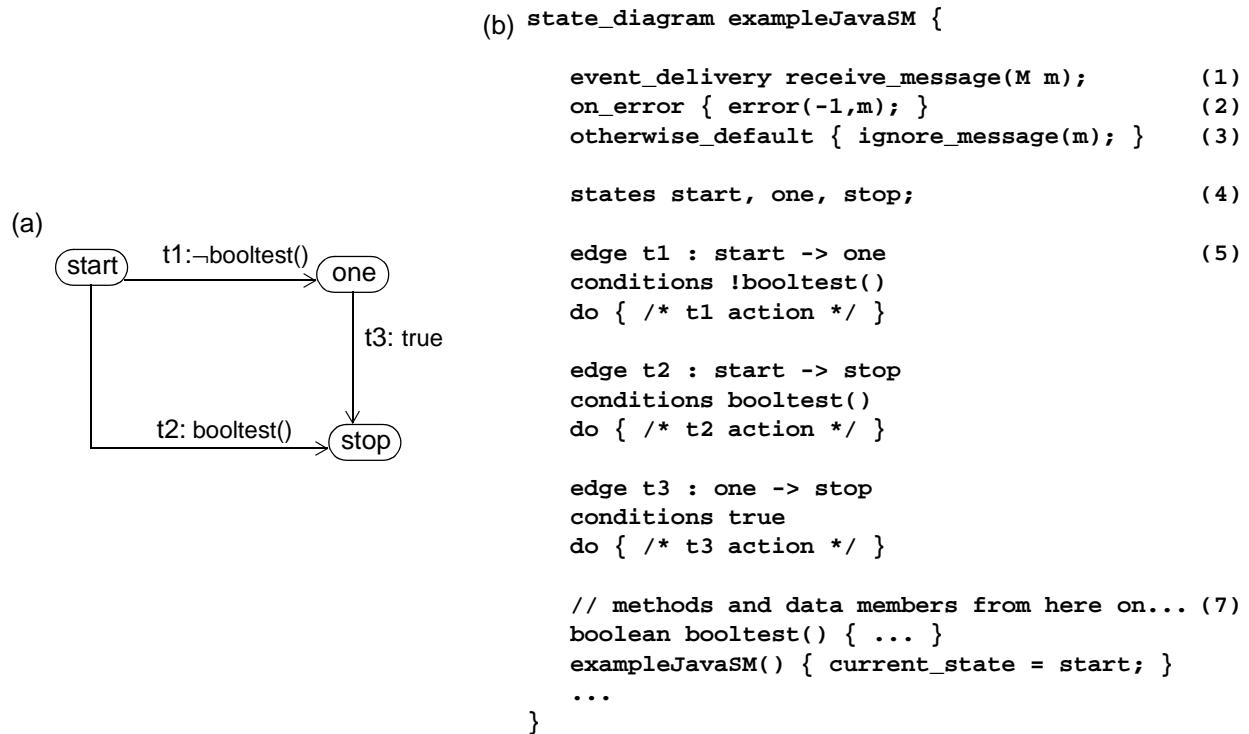


Figure 5: State Diagram Specification

Our model of FSATS required boolean conditions that triggered a transition to be arbitrary Java expressions with no side-effects, and the actions performed by a transition be arbitrary Java statements. Figure 5b shows a JavaSM specification of Figure 5a. (1) defines the name and formal parameters of the void method that delivers a message to the state machine. In the case that actions have corrupted the current

state, (2) defines the code that is to be executed upon error discovery. When a message is received and no transition is activated, (3) defines the code that is to be executed (in this case, ignore the message). The three states in Figure 5a are declared in (4). Edges are declared in (5): each edge has a name, start state, end state, transition condition, and transition action. Java data member declarations and methods are introduced after edge declarations. When the specification of Figure 5b is translated, the class `exampleJavaSM` is generated. Additional capabilities of JavaSM are discussed in [Bat98].

Refinement Declarations. State diagrams can be progressively refined in a layered manner. A *refinement* is the addition of states, actions, edges to an existing diagram. A common situation in FSATS is illustrated in Figure 6. Protocols for missions of the same general type (e.g., WRFPE) share the same protocol fragment for initialization (Figure 6a). A particular mission type (e.g., WRFPE-artillery) grafts on states and edges that are specific to it (Figure 6b). Additional missions contribute their own states and edges (Figure 6c), thus allowing complex state diagrams to be built in a step-wise manner.

The original state diagram and each refinement are expressed as separate JavaSM specifications that are encapsulated in distinct layers. When these layers are composed, their JavaSM specifications are translated into a Java class hierarchy. Figure 6d shows this hierarchy: the root class was generated from the JavaSM specification of Figure 6a; its immediate subclass was generated from the JavaSM refinement specification of Figure 6b; and the terminal subclass was generated from the JavaSM refinement specification of Figure 6c. Figure 7 sketches a JavaSM specification of this refinement chain.




Figure 6: Refining State Diagrams

```
state_diagram black {
    states one_black, two_black, three_black;

    edge a : one_black -> two_black ...
    edge b : one_black -> three_black ...
}

state_diagram shaded refines black {
    states one_shaded;

    edge c : one_black -> one_shaded ...
    edge d : one_shaded -> three_black ...
    edge e : two_black -> three_black ...
    edge f : two_black -> two_black ...
}

state_diagram white refines shaded {
    states one_white;

    edge g : two_black -> one_white ...
    edge h : one_white -> three_black ...
}
```

Figure 7: A JavaSM Refinement-and-Inheritance Hierarchy

Inheritance plays a central role in this implementation. All the states and edges in Figure 6a are inherited by the diagram refinements of Figure 6b, and these states, edges, etc. are inherited by the diagram refinements of Figure 6c. The diagram that is executed is created by instantiating the bottom-most class of the refinement chain of Figure 6d. Readers will again recognize this an example of the JTS paradigm of Section 2.3.

Perspective. Domain-specific languages for state diagrams are common (e.g., [Ber92, Har87-96b, Nei97, Ell97]). Our way of expressing state diagrams — namely as states with enter and exit methods, edges with conditions and actions — is an elementary subset of Harel’s Statecharts [Har87-96] and SDL extended finite state machines [Ell97]. The notion of “refinement” in Statecharts is the ability to “explode” individual nodes into complex state diagrams. This is very different than the notion of refinement explored in this paper. Our work is closer to the “refinement” of extended finite state machines in SDL where a process class (which encodes a state machine) can be refined via subclassing (i.e., new states and edges are added to extend the parent machine’s capabilities). While the idea of state machine refinements isn’t new, it is new in the context of a DSL-addition to a general-purpose programming language (Java) and it is fundamental in the context of component-based development of FSATS simulators.

4 Preliminary Results

Our preliminary findings are encouraging: the objectives of the redesign are met by the GenVoca-FSATS design:

- it is now possible to specify, add, verify, and test a mission type independent of other mission types (because layers/aspects encapsulate code by mission type, the same unit by which it is specified),
- it is now possible to remove and replace mission types to accommodate varying user requirements, and
- JavaSM allows a direct implementation of a specification, thereby reducing the “conceptual distance” between specification and implementation.

As is common in re-engineering projects, detailed statistics on the effort involved in the original implementation are not available. However, we can make some rough comparisons. We estimate the time to add a mission to the original FSATS simulator at about 1 month. A similar addition to GenVoca-FSATS was accomplished in about 3 days, including one iteration to identify and correct an initial misunderstanding of the protocols for that mission.

To evaluate the redesign in a less anecdotal fashion, we use simple measures of class complexity as indicators of overall program complexity. In particular, we use the number of methods (**nmeth**), the number of lines of code (**nloc**), and the number of tokens/symbols (**nsymb**) per class. (We originally used other metrics [Chi91], but found they provided no further insights). Because of our use of JTS, we have access to both component-specification code (i.e., layered JavaSM code written by FSATS engineers) and generated non-layered pure-Java code (which approximates code that would have been written by hand). By using metrics to compare pure-Java code vs. JavaSM code and layered vs. non-layered code, we can quantitatively evaluate the impact of layering and JavaSM on software simplification.

Complexity of Non-Layered Java Code. Consider a non-layered design of FSATS. Suppose all of our class refinement chains were “squashed” into single classes — these would be the classes that would be written by hand if a non-layered design were used. Consider the FSATS class hierarchy that is rooted by `MissionImpl`; this class encapsulates methods and an encoding of a state diagram that is shared by all OPFACS. (In our prototype, we have implemented different variants of WRFFE missions). Class `Formission`, a subclass of `MissionImpl`, encapsulates the additional methods and the Java-equivalent of state dia-

gram edges/states that define the actions that are specific to a Forward Observer. Other subclasses of `MissionImpl` encapsulate additions to that are specific to other OPFACs. The “Pure Java” columns of Table 1 present complexity statistics of the `FoMission` and `MissionImpl` classes. *Note that our statistics for subclasses, by definition, must be no less than those of their superclasses (because the complexity of superclasses is inherited).*

Class Name	Pure Java			JavaSM		
	nmeth	nloc	nsymb	nmeth	nloc	nsymb
<code>MissionImpl</code>	117	461	3452	54	133	1445
<code>FoMission</code>	119	490	3737	56	143	1615

Table 1. Statistics for Non-Layered Implementation of Class `FoMission`

One observation is immediately apparent: the number of methods (117) in `MissionImpl` is huge. Different encoding techniques for state diagrams might reduce the number, but the complexity would be shifted elsewhere (e.g., methods would become more complicated). Because our prototype presently includes missions for WRRFFE, we must expect that the number of methods in `MissionImpl` will increase. Consider the following: the generic WRRFFE mission contributes over 30 methods to `MissionImpl` alone; when WRRFFE is specialized for a particular weapon system (e.g., mortar), another 10 methods are added. Since WRRFFE is representative of mission complexity, as more mission types are added with their weapon specializations, it is not inconceivable that `MissionImpl` will have several hundred methods. Clearly, such a class would be both incomprehensible and unmaintainable.⁴

Now consider the effects of using JavaSM. The “JavaSM” columns of Table 1 show corresponding statistics, where state exit and enter declarations and edge declarations are treated as (equivalent in complexity as) method declarations. We call such declarations *method-equivalents*. Comparing the corresponding columns in Table 1, it is clear that coding in JavaSM reduces software complexity by a factor of 2. That is, the number of method-equivalents is reduced by a factor of 2 (from 119 to 56), the number of lines of code is reduced by a factor of 3 (from 490 to 143), and the number of symbols is reduced by a factor of 2 (from 3737 to 1615). However, the problem that we noted in the pure-Java implementation remains. Namely, the generic WRRFFE mission contributes over 10 method-equivalents to `MissionImpl` alone; when WRRFFE is specialized for a particular weapon system (e.g., mortar), another 3 method-equivalents are added. While this is substantially better than its non-layered pure-Java equivalent, it is not inconceivable that `MissionImpl` will have over a hundred method-equivalents in the future. *While the JavaSM DSL indeed simplifies specifications, it only delays the onset of design fatigue. Non-layered designs of FSATS may be difficult to scale and ultimately hard to maintain.*

Complexity of Layered Java Code. Now consider a layered design implemented in pure Java. The “Inherited Complexity” columns of Table 2 show the inheritance-cumulative statistics for each class of the `MissionImpl` and `FoMission` refinement chains. The rows where `MissionImpl` and `FoMission` data are listed in **bold** represent classes that are the terminals of their respective refinement chains. These rows correspond to the rows in Table 1. The “Isolated Complexity” columns of Table 2 show complexity statistics for individual classes of Table 2 (i.e., we are measuring class complexity and not including the complexity of superclasses). Note that most classes are rather simple. The `MissionAnyL.MissionImpl` class, for example, is the most complex, with 43 methods. (This class encapsulates “infrastructure” methods used by

4. It would be expected that programmers would introduce some other modularity, thereby decomposing a class with hundreds of methods into multiple classes with smaller numbers of methods. While this would indeed work, it would complicate the “white-board”-to-implementation mapping (which is what we want to avoid) and there would be no guarantee that the resulting design would be mission-type extensible.

Class Name	Inherited Complexity			Isolated Complexity		
	nmeth	nloc	nsymb	nmeth	nloc	nsymb
MissionL.MissionImpl	9	25	209	9	25	209
ProxyL.MissionImpl	11	30	261	2	5	52
MissionAnyL.MissionImpl	51	179	1431	43	149	1170
MissionWrffeL.MissionImpl	83	314	2342	35	135	911
MissionWrffeMortarL.MissionImpl	93	358	2677	13	44	335
MissionWrffeArtyL.MissionImpl	109	425	3187	19	67	510
MissionWrffeMlrsL.MissionImpl	117	461	3452	11	36	265
BasicL.FoMission	117	461	3468	0	0	16
MissionWrffeMortarL.FoMission	117	468	3547	4	7	79
MissionWrffeArtyL.FoMission	119	484	3687	7	16	140
MissionWrffeMlrs.FoMission	119	490	3737	3	6	50

Table 2. Statistics for a Layered Java Implementation of Class FoMission

all missions). Table 2 indicates that layering disentangles the logic of different aspects/features of the **FoMission** and **MissionImpl** classes into units that are small enough to be comprehensible and manageable by programmers. For example, instead of having to understand a class with 117 methods, the largest layered subclass has 43 methods; instead of 461 lines of code there are 149 lines, etc.

To gauge the impact of a layered design in JavaSM, consider the “Inherited Complexity” columns of Table 3 that show statistics for **MissionImpl** and **FoMission** refinement chains written in JavaSM. The “Isolated Complexity” columns of Table 3 show corresponding statistics for individual classes. These statistics show that layered JavaSM specifications are indeed compact: instead of 43 methods there are 24 method-equivalents, instead of 149 lines of code there are 65 lines, etc. Thus, a combination of domain-specific languages with layered designs noticeably reduces program complexity.

The reduction in program complexity is a key goal of our project; these tables support the observations of FSATS engineers: the mapping between a “white-board” design of FSATS protocols and an implementation is both direct and invertible with layered JavaSM specifications. That is, writing components in JavaSM matches the informal designs that domain experts use; it requires fewer mental transformations from design to implementation which simplifies maintenance and extensibility, and makes for a much less error-prone product. In contrast, mapping from the original FSATS implementation back to the design was not possible due to the lack of an association of any particular rule or set of rules with a specific mission.

Class Name	Inherited Complexity			Isolated Complexity		
	nmeth	nloc	nsymb	nmeth	nloc	nsymb
MissionL.MissionImpl	8	20	169	8	20	169
ProxyL.MissionImpl	10	25	221	2	5	52
MissionAnyL.MissionImpl	34	90	877	24	65	656
MissionWrffeL.MissionImpl	45	115	1132	11	25	255
MissionWrffeMortarL.MissionImpl	48	121	1231	3	6	99
MissionWrffeArtyL.MissionImpl	52	129	1383	4	8	152
MissionWrffeMlrsL.MissionImpl	54	133	1445	2	4	62
BasicL.FoMission	54	133	1461	0	0	16
MissionWrffeMortarL.FoMission	54	136	1518	2	3	57
MissionWrffeArtyL.FoMission	55	140	1586	3	4	68
MissionWrffeMlrs.FoMission	56	143	1615	2	3	29

Table 3. Statistics on a Layered JavaSM Implementation of Class FoMission

5 Conclusions

Extensibility is the property that simple changes to the design of a software artifact requires a proportionally simple effort to modify its source code. Extensibility is a result of premeditated engineering, whereby anticipated variabilities in a domain are made simple by design. Two complementary technologies are emerging that make extensibility possible: *product-line architectures (PLAs)* and *domain-specific languages (DSLs)*. Product-lines rely on components to encapsulate the implementation of basic features or “aspects” that are common to applications in a domain; applications are extensible through the addition and removal of components. Domain-specific languages enable applications to be programmed in high-level domain abstractions, thereby allowing compact, clear and machine-processable specifications to replace detailed and abstruse code. Extensibility is achieved through the evolution of specifications.

FSATS is a simulator for Army fire support and is representative of a complex domain of distributed command-and-control applications. The original implementation of FSATS had reached a state of design fatigue, where anticipated changes/enhancements to its capabilities would be very expensive to realize. We undertook the task of redesigning FSATS so that its inherent and projected variabilities — that of adding new mission types — would be easy to introduce. Another important goal was to minimize the “conceptual distance” from “white-board” designs of domain experts to actual program specifications; because of the complexity fire-support, these specifications had to closely match these designs to make the next-generation FSATS source understandable and maintainable.

We achieved the goals of extensibility and understandability through an integration of PLA and DSL technologies. We used a GenVoca PLA to express the building blocks of fire support simulators as layers or aspects, whose addition or removal simultaneously impacts the source code of multiple, distributed programs. But a layered design was insufficient, because our components could not be written easily in pure Java. The reason is that the code expressing state diagram abstractions was so low-level that it would be difficult to read and maintain. We addressed this problem by extending the Java language with a domain-specific language to express state diagrams and their refinements, and wrote our components in this extended language. Preliminary findings confirm that our component specifications are substantially simplified; “white-board” designs of domain experts have a direct and invertible expressions in our specifications. Thus, the combination of PLAs and DSLs was essential in creating extensible fire support simulators.

While fire support is admittedly a domain with very specific and unusual requirements, there is nothing domain-specific about the need for PLAs, DSLs, and their benefits. In this regard, FSATS is not unusual; it is a classical example of many domains where both technologies naturally complement each other to produce a result that is better than either technology could deliver in isolation. Research on PLA and DSL technologies should focus on infrastructures (such as IP and JTS) that support their integration; research on PLA and DSL methodologies must be more cognizant that synergy is not only possible, but desirable.

Acknowledgments. We thank Dewayne Perry (UTexas) for his insightful comments on an earlier draft and bringing SDL to our attention. We thank Frank Weil (Motorola) for clarifying discussions on SDL state machines.

6 References

- [ARL99] “System Segment Specification (SSS) for the Fire Support Automated Test System (FSATS)”, Applied Research Laboratories, The University of Texas, 1999. See also URL <http://www.arlut.utexas.edu/~fsatswww/fsats.shtml>.
- [Bat92] D. Batory and S. O’Malley, “The Design and Implementation of Hierarchical Software Systems with Reusable Components”, *ACM TOSEM*, October 1992.

- [Bat98a] D. Batory, B. Lofaso, and Y. Smaragdakis, "JTS: Tools for Implementing Domain-Specific Languages", *5th International Conference on Software Reuse*, Victoria, Canada, June 1998. Also URL <http://www.cs.utexas.edu/users/schwartz/JTS30Beta2.htm>.
- [Bat98b] D. Batory, "Product-Line Architectures", *Smalltalk and Java Conference*, Erfurt, Germany, October 1998.
- [Bat99] D. Batory, Y. Smaragdakis, and L. Coglianese, "Architectural Styles as Adaptors" *Software Architecture*, Kluwer Academic Publishers, Patrick Donohoe, ed., 1999.
- [Bax92] I. Baxter, "Design Maintenance Systems", *CACM*, April 1992.
- [Ber92] G. Berry and G Gonthier, "The Esterel Synchronous Programming language: Design, Semantics, and Implementation", *Science of Computer Programming*, 1992, 87-152.
- [Bos99] J. Bosch, "Product-Line Architectures in Industry: A Case Study", *ICSE 1999*, Los Angeles.
- [Cza99] K. Czarnecki and U.W. Eisenecker, "Components and Generative Programming", *ACM SIGSOFT*, 1999.
- [Chi91] S.R. Chidamber and C.F. Kemerer, "Towards a Metrics Suite for Object Oriented Design", *OOPSLA 1991*.
- [Deu97] A. van Deursen and P. Klint, "Little Languages: Little Maintenance?", *SIGPLAN Workshop on Domain-Specific Languages*, 1997.
- [Ell97] J. Ellsberger, D. Hogrefe, A. Sarma, *Formal Object-Oriented Language for Communicating Systems*, Prentice-Hall, 1997.
- [Fin98] R.B. Findler and M. Flatt, "Modular Object-Oriented Programming with Units and Mixins", *ICFP 98*.
- [Gam95] E. Gamma et.al. *Design Patterns Elements of Reusable Object-Oriented Software*. Addison-Wesley, Reading, Massachusetts, 1995.
- [Gra97] T. Graves, "Code Decay Project", URL <http://www.bell-labs.com/org/11359/projects/decay/>.
- [Gog86] J.A. Goguen, "Reusing and Interconnecting Software Components", *IEEE Computer*, February 1986.
- [Har87] D. Harel, "Statecharts: A Visual Formalism for Complex Systems", *Science of Computer Programming*, 1987, 231-274.
- [Har96] D. Harel and E. Gery, "Executable Object Modeling with Statecharts", *ICSE 1996*.
- [Kic97] G. Kiczales, J. Lamping, A. Mendhekar, C. Maeda, C. Lopes, J. Loingtier, and J. Irwin, "Aspect-Oriented Programming", *ECOOP 97*, 220-242.
- [Mag99] "System Segment Specification (SSS) for the Advanced Field Artillery Tactical Data System (AFATDS)", Magnavox, 1999.
- [Nei97] J. Neighbors, "DataXfer Protocol", BayFront Technologies, 1997, URL <http://bayfronttechnologies.com>.
- [Ree92] T. Reenskaug, et al., "OORASS: Seamless Support for the Creation and Maintenance of Object-Oriented Systems", *Journal of Object-Oriented Programming*, 5(6): October 1992, 27-41.
- [SEI99] Software Engineering Institute, "The Product Line Practice (PLP) Initiative", URL http://www.sei.cmu.edu/plp/plp_init.html.
- [Sim95] C. Simonyi, "The Death of Computer Languages, the Birth of Intentional Programming", *NATO Science Committee Conference*, 1995.
- [Sma98] Y. Smaragdakis and D. Batory, "Implementing Layered Designs with Mixin Layers", *ECOOP 1998*.
- [Tay99] R. Taylor, Panel on Software Reuse, *Motorola Software Engineering Symposium*, Ft. Lauderdale, 1999.
- [Tok99] L. Tokuda and D. Batory, "Evolving Object-Oriented Architectures with Refactorings", *Conf. on Automated Software Engineering*, Orlando, Florida 1999.
- [Van96] M. Van Hilst and D. Notkin, "Using Role Components to Implement Collaboration-Based Designs", *OOPSLA 1996*, 359-369.
- [Wei99] D.M. Weiss and C.T.R Lai, *Software Product-Line Engineering*, Addison-Wesley, 1999.