

Living it up with a Live Programming Language

Sean McDirmid

École Polytechnique Fédérale de Lausanne (EPFL)

sean.mcdirmid@epfl.ch

Abstract

A dynamic language improves programmer productivity through flexible typing, a focus on high-level programming, and by streamlining the edit-compile-debug cycle. **Live languages** go beyond dynamic languages with more programmer-centric features. A live language supports **live programming** that provides programmers with responsive and continuous feedback about how their edits affect program execution. A live language is also based on declarative programming constructs such as rules or data-flow connections so that programmers can write less code. A live language should also provide programmers with responsive semantic feedback to enable time-saving services such as code completion. This paper describes the design a textual live language known as SuperGlue. SuperGlue is based on reactive values known as **signals** that are supported with declarative data-flow connections and **dynamic inheritance**. Through signals and dynamic inheritance, SuperGlue supports live programming, declarative programming, and responsive semantic feedback. We demonstrate live programming in SuperGlue with a working prototype.

1. Introduction

Dynamic programming languages improve programmer productivity by reducing verbosity in both source code and the development process. Although labeled as “dynamic” because they often rely on dynamic typing, dynamic languages are best characterized by their emphasis of flexibility, conciseness, and ease of development over performance and generality. Dynamic typing promotes flexibility by enforcing fewer restrictions on type compatibility and conciseness by eliminating verbose type annotations. Dynamic languages also emphasize high-level programming models, often based on higher-order functions or objects, that reduce the amount of code needed to express a program. Together,

flexible typing and high-level programming models support component programming so that programmers can create feature-rich programs with little code by reusing existing components. Finally, dynamic languages streamline the edit-compile-debug cycle by either eliminating or hiding compilation and by allowing code to be changed while a program is running. For example, many Lisp [26], Smalltalk [12], and Erlang [1] environments support *hot swapping* where a program’s code can be updated without restarting.

Can we design languages that are more programmer-friendly than existing dynamic languages? For inspiration, look at visual languages such as Fabrik [16], LabVIEW [22], and Forms/3 [3], which are designed to be used by casual programmers, novices, and even end users. Visual languages are based on simple, often declarative, programming models that heavily emphasize component reuse. For example, Apple’s Quartz Composer [18] supports the assembly of animation components through reactive data-flow connections that encapsulate change-propagation details. The simple programming models of visual languages also enable *live programming*, which provides programmers with immediate and continuous feedback about how their edits affect program execution. For example, editing a component connection graph in Quartz Composer immediately changes an animation that is being composed. Compared to live programming, the hot swapping supported by dynamic languages is mushy: the effect of edited code on program behavior is not apparent until the code is eventually re-executed. For example, editing the imperative Smalltalk statement “`w fill: red`” to “`w fill: blue`” will not immediately re-color blue all widgets that have been bound to `w`.

In spite of their benefits, visual languages are unlikely to replace textual dynamic languages anytime soon. For one thing, text can be written more quickly than a visual notation [24]. More significantly, visual languages must deal with a *scaling-up problem* [4]: the ability to build large programs is hindered by visual notation that does not naturally support non-local naming and procedural or data type abstraction. Modern visual languages solve the scaling up problem in a variety of ways; e.g., Forms/3 [3] allow for name-based cell and form referencing and data abstraction through generic forms. Still, text’s natural support for naming and abstraction makes it more suitable for many pro-

programming tasks, especially when used by experienced programmers. Also, the simple programming models that enable live programming in visual languages can also be applied to textual languages. A *live language* is then a textual or visual language that is based on a simple programming model and supports live programming.

Support for textual naming in a live language is problematic: how are programmers provided with responsive semantic feedback about what names are valid or available? Semantic feedback speeds up the programming process by detecting typos early on and enabling services such as code completion, which allows programmers to browse available named entities. Programmers are deprived of responsive semantic feedback in dynamically-typed languages because valid names are unknown until run-time. Static typing that typically enables semantic feedback is too verbose and inflexible to be used in a live language. Implicit typing would work better in a live language where type inference could provide responsive semantic feedback without verbose and inflexible type annotations.

This paper explores the design of a textual live language that reuses the best ideas from dynamic and visual languages along with a few new ideas. The language, SuperGlue, is based on a simple reactive data-flow programming model that is suited to building interactive programs out of existing components. Components in SuperGlue are connected together through their *signals* into a data-flow graph. The values that flow across signal connections can change over time, which eases the expression of interactive programs. For scaling purposes, signals extend classes through *dynamic inheritance*, where what classes a signal extends can change at run time. Signals and dynamic inheritance seamlessly support live programming by masquerading edits as programmatic changes in the program's mutable execution state. Unlike Self [30] and like Cecil [6], dynamic inheritance in SuperGlue is based on declarative predicates [7], which is very compatible with a reactive data-flow model. Because class extensions are expressed declaratively, type inference can provide programmers with responsive semantic feedback. Dependent typing improves the expressiveness of type inference so that no type annotations are necessary.

Previous work [21] introduces SuperGlue's support for object-oriented signals, and shows how this support eases the construction of interactive programs. This paper expands on this work by augmenting SuperGlue with dynamic inheritance, which significantly simplifies the language, and live programming. Figure 1 shows how live programming occurs in a working Eclipse-based prototype. In this paper, we use a PacMan-like game as a running example. A program is edited as free-form text in the upper half of the programming environment while the execution is shown in the lower half. As soon as valid code is recognized during an edit, the effect of the edit is incorporated into the execution before the next edit is processed. So that live programming remains re-




Figure 1: An example of the SuperGlue programming environment being used in a live programming session; the editor is in the top half of the movie while the running program is shown in the lower half. Click on the movie to replay it.

sponsive, SuperGlue supports error recovery where syntax, semantic, and execution errors are duly noted in the background while program execution continues. The live programming experience that is shown as a movie¹ in Figure 1 is narrated as follows:

- (a) The programmer types “port pacMan > Figure” to define a signal that extends the game’s Figure class, meaning it is a figure that is drawn on the game’s canvas, which is shown on the bottom half of Figure 1. Because no shape for pacMan has been specified yet, it is drawn by default as a rectangle.
- (b) The programmer types the extension “pacMan > Pie,” which causes pacMan to extend the Pie class and be drawn as a circle. The shape of pacMan is updated on screen as soon as the programmer types the “e” in Pie.
- (c) The programmer types the connection “pacMan.fill = yellow,” causing this circle to be immediately colored yellow. Through type inference, code completion can complete typing for .fill in pacMan.
- (d) The programmer types “pacMan.extent = 330,” which causes pacMan’s mouth to immediately appear as a 30 degree slice. The programmer then finishes this line with “+ (time % 30),” which causes pacMan’s mouth to open and close continuously in a 30 degree motion.

The rest of this paper is organized as follows. Section 2 elaborates on live language design problems and shows how these problems are solved in SuperGlue. Section 3 describes how SuperGlue supports live programming. Section 4 describes our preliminary experience with how SuperGlue is used to develop a simple game. Related work is presented in Section 5 and we conclude in Section 6.

¹ All figures in this paper that show the SuperGlue programming environment are QuickTime movies that will play in Acrobat Reader 8 with QuickTime installed. A printable version of this document with frame-expanded movies is also available.

2. Language Design

SuperGlue is based on the data-flow programming model that was originally introduced by Karp and Miller [17] to deal with concurrency through a data-flow graph program representation. Even without considering concurrency, the data-flow model is appealing to live languages because all computational dependencies are explicit and lack implicit side effects. The data-flow model by itself does not scale: non-trivial programs involve a large or unbounded number of data-flow connections that are tedious or impossible to express individually. For this reason, SuperGlue augments its data-flow model with object-oriented constructs. As a brief overview, SuperGlue consists of the following constructs:

- **Signals**, which are data-flow values that facilitate inter-component communication. Signals are objects that extend classes, are connected to other signals, and contain signals defined in their extended classes. Signals are defined as class members using the `port` keyword.
- **Classes**, which are types that encapsulate behavior. Classes are not values and they can only be extended by signals. Classes are defined at the top-level of a program or as members in other classes using the `class` keyword.
- **Connections**, which relate signals to other signals so that their values are equivalent. Connections are rules that are similar to simple uni-directional constraints. Connections are expressed with the `=` operator.
- **Extensions**, which cause signals to extend classes. Extensions are rules that can target signals outside of their definitions. Extensions are expressed with the \triangleright^2 operator.
- **Conditions**, which guard when connection and extension rules can be applied. Conditions query existing connection and extension relationships through the same `=` and \triangleright operators that are used to create these relationships in rules. Conditions are expressed in `if` clauses that surround guarded connection and extension rules.

As an opening example, consider the following SuperGlue code that defines the `Ghost` class, three ghost signals, and connects the fill color of all ghost signals:

```
1 class Ghost { port fill  $\triangleright$  Color; }
2 port (blinky, clyde, bashful)  $\triangleright$  Ghost;
3 Ghost.fill = blue;
```

The first line of this code defines a `Ghost` class with a member `fill` signal that extends the `Color` class. The second line defines three signals: `blinky`, `clyde`, and `bashful` that extend the `Ghost` class. The third line is a rule that connects `fill` signal of all ghosts, including `blinky`, `clyde`, and `bashful`, to the color blue. When used in expressions, class identifiers are variables that quantify over all signals that extend the identified class; e.g., `Ghost` in `Ghost.fill` is a variable that quantifies over all signals that extend the `Ghost`

class. Unlike imperative assignments, connection rules have declarative meaning that is unrelated to statement execution order; e.g., on line three, blue is continuously connected to rather than discretely assigned to the ghost's fill color.

The nature of the values that flow across data-flow connections determines what programs can be concisely expressed in a data-flow language. In SuperGlue as well as many existing textual and visual data-flow languages, such as Lucid [33] and Quartz Composer [18], data flow is **reactive** in that time-varying values flow across connections. Reactive data-flow languages hide change-propagation details from programmers by dealing with the details generically at connection end points. By hiding change-propagation details, continuously changing state is treated in the same simple way that immutable state is. Consider the following code that causes `pacMan`'s mouth to close continuously:

```
1 class Pie { port extent  $\triangleright$  Number; }
2 port pacMan  $\triangleright$  Pie;
3 pacMan.extent = 330 + (time % 30);
```

The last line of this code connects `pacMan`'s `extent` signal to an expression that continuously increases from 330 until 360 degrees, causing `pacMan`'s mouth to shut continuously in a 30 degree motion. Such behavior occurs because the `time` signal is a periodically incrementing number; e.g., it is connected to 1042 and, after ten milliseconds, it is re-connected to 1043. At connection end-points, signals in SuperGlue are provided and consumed by components that express event handling details through a general-purpose language such as Java or Scala [23]. In our example, Scala code updates the `time` signal when time advances and redraws `pacMan` when its `extent` signal changes. We describe how this code interfaces with SuperGlue in Section 2.4. Because SuperGlue code is completely oblivious to what happens over signal connections, interactive programs, where continuous change is common, are expressed more easily. As we will see in Section 3, this obliviousness also supports live programming.

In our previous version of SuperGlue [21], class extensions could only be expressed in signal and class definitions. This restriction complicated SuperGlue's class system and prevented it from being very dynamic; i.e., signal types could not change at run-time. As a simplification and to support live programming, class extensions are now rules in the same way that connections are. For example, the statement "`pacMan \triangleright Figure`" is an extension rule that can be written down wherever `pacMan` is visible—not just where `pacMan` is defined. As rules, both connections and extensions can be guarded by conditions. Consider the following code that causes signals to extend the `PositionByKeyboard` class when they already extend both the `Position` and `Keyboard` classes:

```
1 class Position; class Keyboard;
2 class PositionByKeyboard;
3 if (Position  $\triangleright$  Keyboard)
4   Position  $\triangleright$  PositionByKeyboard;
```

The condition on line three of this code ensures that the `Position` identifier is only bound to signals that extend both the `Position` and `Keyboard` classes. Line four is an extension

²The user types `:`, which is then rendered in the editor as \triangleright .

that then causes these signals to extend the `PositionByKeyboard` class. As we will see in Section 2.2, such extensions are useful in adapting incompatible signals.

Extension rules in SuperGlue target signals and not classes. As syntactic sugar, an extension can be specified in a class definition, but the extension is applied to all signals that extend the defined class and not the class itself. No cyclic inheritance occurs in SuperGlue when the signals of two different classes extend the other class—instead, the classes are semantically equivalent in that a signal that extends one also extends the other. Such mutual extension constructions enable the refinement of existing classes. For example, the definition “`class RichNumber ▷ Number`” and the extension rule “`Number ▷ RichNumber`” does not create a cycle; instead they cause all `Number` signals to be `RichNumber` signals and vice versa. Any signal that is defined as a member of the `RichNumber` class is then a member of all `Number` signals, including numeric literals. Using such extensions, any class can be augmented with new signals after it has been defined.

Because conditions query signals whose values change over time, their truth values can also change over time. Because time-varying conditions can guard extension rules, SuperGlue supports *dynamic inheritance* where what classes a signal extends can vary over time. Dynamic inheritance enables the state-based classification of signals without deeply nested `if` statements, and, as described in Section 3, supports live programming without the need for special-case semantics. Unlike Self [30], where dynamic inheritance occurs through imperative slot assignment, dynamic inheritance in SuperGlue is predicated as in Cecil [6, 7]. Consider the following code that causes `pacMan` to extend the `Damaged` class when it overlaps with a ghost:

```
1 class Figure { port (overlaps, hit) ▷ Hit }
2 class Ghost ▷ Figure; class Damaged ▷ Figure;
3 port pacMan ▷ Figure;
4 if (pacMan.overlaps ▷ Ghost.Hit)
5   pacMan ▷ Damaged;
6 Damaged.fill = red;
```

For the duration that `pacMan` overlaps with some figure, `pacMan`’s `overlaps` signal is bound to the `hit` signal of this figure. The condition on line four is then only true for durations where the figure that `pacMan` overlaps with is a ghost. During these durations, `pacMan` extends the `Damaged` class, inheriting all the behavior of this class and being subject to rules that target signals of this class. For example, the connection of red to a damaged figure’s `fill` color would apply to `pacMan` while it overlaps with a ghost. As soon as `pacMan` no longer overlaps with a ghost, it unextends the `Damaged` class and the special damaged behavior no longer applies.

2.1 Implicit Typing

Many dynamic languages shun static typing because of its verbose type annotations, limited flexibility, and inability to deal with change. Unfortunately, the lack of static typing also deprives programmers of responsive semantic feedback

that is especially important in a live programming environment. Responsive semantic feedback allows programmers to converge on correct edits more quickly through code completion, where programmers can browse the names of available entities, and validation, which immediately informs programmers of typos. Static types can also convey programmer intent that is useful in resolving ambiguity. In SuperGlue, multiple connections can connect the same signal. To resolve ambiguity, connections that target more specific static types are considered before connections that target less specific static types if the guarding conditions of both connections are true. Such by-type prioritization is analogous method overriding in most object-oriented languages.

To provide semantic feedback without verbose and inflexible type annotations, types are computed in SuperGlue through type inference to support *implicit typing*. Implicit typing is viable in SuperGlue because extensions are declarative rules that can be accurately inferred by an analysis that is both flow and context-insensitive. The inferred type of a symbol is a set of classes that restricts the symbol from only being bound at run-time to signals that extend all of these classes. Sub-typing is then defined as a subset relationship: t_a is a sub-type of t_b if all classes in t_b are classes in t_a . Type inference builds up a symbol’s inferred type by analyzing surrounding extension conditions and known extension rules that target classes already contained in the inferred type.

As an example of how type inference occurs, consider the following code that overrides a position’s `x` signal when the position extends the `PositionCell` class:

```
1 class Position { port (x,y) ▷ Number; }
2 class PositionCell
3 { port (xc, yc) ▷ NumberCell; }
4 Position.x = 0;
5 if (Position ▷ PositionCell)
6   Position.x = Position.xc.result;
```

Overriding in SuperGlue is defined as follows: if the type of a is t_a , the type of c is t_c , and t_c is a strict sub-type of t_a , then the connection $a = b$ overrides the connection $c = d$. On line four, the inferred type of `Position` includes only the `Position` class because the type of a class identifier always contains itself. On line six, `Position`’s inferred type also contains the `PositionCell` class because of the surrounding extension-querying condition on line five. As a result, the connection on line six, which targets the type `(Position, PositionCell)`, overrides the connection on line four, which targets the type `(Position)`. Because `Position`’s inferred type on line six contains the `PositionCell` class, the member signal `xc` can be accessed through `Position` and is visible by code completion.

SuperGlue supports a form of dependent typing where nested classes are prefixed by their containing signals. When type inference computes a symbol’s type, prefixes are approximated as signal and class symbols. Consider the code in Figure 2 that creates ten ghosts. On line nine, the in-

```

1 class Duplicator
2 { size ▷ Number;
3   class T
4   class get ▷ T { port index ▷ Number; } }
5 port ghosts ▷ Duplicator;
6 ghosts.T ▷ (Ghost,Figure);
7 ghosts.size = 10;
8 if (ghosts.get.index = 3)
9   ghosts.get.fill = red;

```

Figure 2: SuperGlue code that creates ten ghosts, where the ghost at index three is colored red.

ferred type of `ghosts.get` contains the following classes: `ghosts.get` by definition, `Duplicator.get` because the `ghosts` prefix extends `Duplicator` on line five, `ghosts.T` because of the extension on line four and by a binding of `Duplicator` to `ghosts` on line five, `Duplicator.T` by line five again, and finally, `Ghost` and `Figure` by the extension on line six. On line nine, the `fill` signal can be accessed and connected because `ghosts.get` signals extend `Figure`.

2.2 Dynamic Typing and Adaptation

In SuperGlue, because both connections and extensions can be guarded by state-dependent conditions, when they are applied at run-time is not statically known. As a result, SuperGlue depends on dynamic typing to verify two connection properties: first, that primitive-extending signals are connected to concrete values of the appropriate primitive class; and second, that all signals are connected unambiguously. The first property ensures that component implementations correctly communicate through primitive-extending signals; e.g., a providing component pushes through a number that is consumed as a number by one or more other components. The second property detects multiple active connections to the same signal that do not having overriding relationships. In both cases, dynamic typing ensures that components communicate correct and unambiguous values at run-time.

Independently-developed or differently-purposed components will often not agree on the types of their non-primitive signals. For this reason, SuperGlue permits connections between type-incompatible non-primitive signals and allows compatibility to be defined through additional rules. Consider the following code where `pacMan`'s position is connected to a keyboard:

```

1 class Canvas
2 { class Figure { port position ▷ Position; }
3   port keyboard ▷ Keyboard;
4 port game ▷ Canvas, pacMan ▷ game.Figure;
5 pacMan.position = game.keyboard;

```

The signal types involved in the connection on line five are incompatible: the type of `pacMan.position` contains only the `Position` class while the type of `game.keyboard` contains only the `Keyboard` class. The programmer intends the position of `pacMan` to be controlled by the keyboard. Such intent can be realized with code that defines what it means

for a signal to extend the otherwise unrelated `Position` and `Keyboard` classes. At run-time, the above connection causes `pacMan`'s `position` signal to extend both the `Position` and `Keyboard` classes. These classes can then be combined to describe the target of an extension rule:

```

1 if (Position ▷ Keyboard)
2   Position ▷ PositionByKeyboard;

```

By this extension rule, any signal that extends both the `Keyboard` and `Position` classes also extends the `PositionByKeyboard` class, which we refer to as an *adapter mixin*. The application of an adapter mixin occurs automatically to adapt connections of incompatible types. The `PositionByKeyboard` class is then defined as follows:

```

1 class PositionByKeyboard▷ (PositionCell,Keyboard)
2 { go_up = key_up ; go_down = key_down ;
3   go_left = key_left; go_right = key_right; }

```

The `go_*` and `key_*` signals are respectively defined in the `PositionCell` and `Keyboard` classes whose definitions are not shown. The connections in `PositionByKeyboard` cause imperative position state that is maintained by `PositionCell` to be incremented according to keyboard input and a time signal that is not shown. This code is applied whenever positions are connected to keyboards. For example, if multiple figure positions are connected to the keyboard, these figures will move in unison by keyboard input.

2.3 Implementation

Our prototype implementation of SuperGlue conceptually builds and maintains a dependency graph that tracks how signal values are inter-dependent. So that dependencies are updated reactively, our prototype is based on a pattern of *damage and repair*³: when a signal changes, its dependencies are “damaged” and scheduled in a work-list for “repair” through re-evaluation. Damage is then propagated through the dependency graph until dependencies are either exhausted or are unaffected by the change. Maintaining an exact dependency graph is unrealistic because signals can depend on each other in cyclic ways that are difficult to track. Instead, our prototype is conservative: a dependency is recorded when the evaluation of signal *A* depends on signal *B*, but is never explicitly “unrecorded” when signal *A* no longer depends on signal *B*. Dependent signals are instead flushed after being damaged, and if dependent signals still depend on the changing signal, then this dependency will be re-recorded when they are repaired. Although it is possible for signals to be damaged by changes in other signals that they no longer depend on, flushing ensures that such cases only occur once per incident.

In our prototype, connection and extension rules are organized into *circuits* based on the inferred types they target. The *maximal inferred type* of a signal, which is computed by optimistically applying extension and connection rules to

³The pattern name was inspired by support in Eclipse [15] for presentation reconciling.

the signal value, determines what circuits are used to compute the signal's connection and dynamic type. As a result, run-time signal evaluation does not visit rules that are statically known to be completely unrelated to the signal.

2.4 Code Behind

SuperGlue code cannot express any kind of non-connective computations, such as arithmetic or drawing on a canvas. Instead, these computations are expressed as *code behind*⁴ connections and extensions that are written in Java or Scala [23]. A code behind connection expresses the connection's guard and result in Java or Scala code. The SuperGlue run-time automatically records all dependency on behalf of the code behind connections when they are evaluated. Code behind connections notify the run-time state changes; e.g., a `time` signal has advanced.

Code behind extensions are evaluated by the run-time when a signal extends or unextends their associated classes. Code behind extensions can implement behavior that reacts to a signal extending a class. For example, the code behind extension for the `Figure` class member of a `Canvas` class adds the extending signal to the list of figures that are drawn in the qualifying canvas signal, and then redraws the canvas. Code behind extensions are also responsible for preserving the class-specific state that the signal maintains as long as it extends the specified class. In our example, the code behind extension for `Figure` allocates state for each figure signal that is drawn on the containing canvas.

3. Live Programming

Live programming in SuperGlue is enabled by its simple reactive data-flow model and a programming environment that takes advantage of this model. By adhering to the declarative data-flow model, the code of a SuperGlue program has a direct relationship with its execution state. As a result, the programming environment can responsively update program execution state according to programmer edits. The object-oriented features that improve SuperGlue's scalability are declarative and do not interfere with this property; e.g., dynamic inheritance occurs via declarative predicates rather than imperative slot assignment. SuperGlue is tightly integrated into a programming environment that includes a source code editor that provides the usual modern conveniences such as auto-indenting, semantic highlighting, code completion, hover help, and hyperlink navigation. These services enhance live programming by allowing programmers to write, read, and navigate code more quickly; e.g., the programmer can use code completion to conveniently browse what signals and classes are available through a signal. As mentioned before, many of these services, such as code completion, would be difficult to implement reliably and responsively without implicit typing.

⁴“Code behind” is often used to describe imperative code that implements complicated behavior in an otherwise declaratively-specified user interface.

The SuperGlue code is processed by a runtime in four phases: scanning, parsing, type inference, and execution. All four phases are incrementally performed after every edit. As described in Section 2.3, the execution phase is additionally performed when program state changes; e.g., when the `time` signal advances. All four phases are based on a layered version of the damage and repair pattern that was described in Section 2.3: a keyboard edit damages the token stream, a change in the token stream damages the parse tree, a change in the parse tree damages inferred type information, and finally, a change in inferred type information can damage the program's execution.

The live programming experience is sensitive to the **error recovery** and **edit latency** characteristics of each code-processing phase. Syntactic, semantic, and run-time errors are inevitable while the programmer is typing. In the presence of these errors, a programmer should still receive useful feedback about the other parts of the program. Forcing programmers to immediately fix errors slows them down by forcing them off their preferred development plan. Edit latency is the amount of time that it takes for an edit to percolate through all four phases. Long edit latency reduces the responsive of the feedback provided to programmers, which leads to a “mushy” live programming experience. Based on personal experience, a programmer will notice edit latency of 100 milliseconds and become distracted by latency that is greater than 500 milliseconds.

Editing in the SuperGlue programming environment is syntax-aware and not syntax-directed [29]; i.e., syntax errors are allowed. Scanning and parsing in our prototype is fairly fast and supports very robust error recovery: in the presence of syntax errors, the program still executes, other parts of the program can be processed reliably, and editor services will still work correctly. We achieve this through precedence parsing [11], which as a bottom-up parsing method, is easily made incremental and very error tolerant. In fact, precedence parsers are not sensitive to syntax errors at all. The flip-side of precedence parsing is that it is not very expressive and cannot detect syntax errors as conventional BNF-based parsers can. To deal these problems, we augment precedence parsing with enhanced scanning, a brace matcher, and a type checker that also checks for syntax errors. A detailed discussion of these techniques is beyond the scope of this paper and will be reported on in the future.

3.1 Type Inference

Incremental type inference is achieved by storing computed type information and typing context at each node in the parse tree. The parse node's type information is damaged when the node's parse structure, typing context, or child-tree types have changed. Error recovery involves reporting detected errors in the programming environment and then reverting to a previously computed or default types for use in typing parent parse nodes. Also, symbol tables support change operations. When the symbol table is used by a parse

```

class MyFigure { port position; }
class Ghost ▷ MyFigure {
  this.position ▷ Position;
}
port clyde ▷ Ghost;
clyde.position,

```




Figure 3: An example of how type inference is used to provide semantic feedback in the SuperGlue programming environment.

node to lookup some name, a dependency with the parse node is recorded in the table. When a symbol is changed, all parse nodes that depend on the symbol’s new and old names are damaged. As with scanning and parsing, type errors, such as the failure to find a symbol bound to a name, are reported without stopping the executing program. At worst, a type error acts to disable the current rule being edited.

All connection and extension rules that are encoded in a SuperGlue program are organized by our prototype according to the inferred types that they target. Inference involves traversing rules that are compatible with a symbol’s growing inferred type. When a rule is changed in source code, all symbols whose inferred types are affected by the old and new version of the rule are damaged so that repair will recompute their inferred types using the new rule. Because a rule can be used in computing types for a large number of symbols, changing a rule can incur noticeable latency. For example, changing a rule that targets the `Number` class will not only damage all user-defined symbols whose type contains `Number`, but will also damage all numeric literals. In a 280 line SuperGlue program, which we describe in Section 4, changing a rule that targets the `Number` class incurs a type inference delay of around three and a half seconds⁵. However, because most classes are only extended a few times in a program, changes to most rules incur no noticeable latency.

Various techniques are used in our prototype to speed up incremental type inference. Because of dependent typing, the number of possible types in a program is exponential with respect to the number of declared classes—although in practice only a few types are significant. When editing changes an extension rule that targets an inferred type, all symbols whose inferred types are subtypes of this type must be damaged. All significant inferred types are hashed in our prototype so that the significant sub-types of a type can be quickly computed and its symbols damaged; otherwise the editing of extension rules would be intractable.

As an example of live type inference, consider the movie in Figure 3 that is narrated as follows:

- (a) The programmer types `port clyde ▷ Ghost`, where `Ghost` is underlined with red jaggies because it cannot be resolved to a symbol.
- (b) As soon as the programmer defines the `Ghost` class, the reference to `Ghost` in `clyde`’s definition becomes valid. The programmer then tries to access `clyde`’s `position` signal, which is underlined with red jaggies because `position` is not a member of `clyde`.
- (c) As soon as the programmer causes the `Ghost` class to extend the `MyFigure` class, the access to `clyde`’s `position` signal becomes valid.
- (d) Because the `position` signal in all `Ghost` signals extends the `Position` class, the `x` and `y` position signals show up in code completion on `clyde`’s `position` signal.

3.2 Execution

Hot swapping in a dynamic language replaces code so that its re-execution occurs according to the new version of the code as opposed to the old version. Changing class members or inheritance relationships can also cause updates to objects in the heap, as is supported in most Smalltalk [12] environments. By being based on a simple data-flow model, live programming in SuperGlue goes much more farther than this: code edits immediately change the program’s data-flow graph and the observable execution state of the program. Because of SuperGlue’s support for reactivity and dynamic inheritance, live programming can be supported with very few special-case semantics. Edits that change connection and extension rules semantically masquerade as programmatic change. As we described in Section 2, connection and extension rules are guarded by arbitrary conditions whose truth values can change over time. The editing of a rule can then resemble changing the state of an implicit guarding condition. For example, if the programmer types the code “`pacMan ▷ Pie`,” the added extension rule can be realized as a pre-existing extension “`if (P) pacMan ▷ Pie`,” where condition `P` changes from false to true. In our prototype, rule changes cause dependent run-time signal values to be damaged, where signal re-evaluation occurs according to an up-to-date set of rules. Through this masquerade, only edits that add or remove signal and class definitions need to be special-cased in our prototype.

Because variables can be bound to an unbounded number of signal values, rule changes can result in higher latency than what can occur during type inference. Additionally, edits that add new signals to a program incur latency in the initialization of these signals and all of their member signals. Because edit latencies are linear in the number of created signals, so it is easy to imagine how an edit that increases the size of something can take an arbitrary amount of time to process. We given an example of this problem Section 4.2. Initialization latency is currently one of the weakest parts of our prototype and we hope to improve on it in the future through tuning and multi-threading. Note

⁵ All times are measured in this paper on a 2 GHz Macbook Pro with 2 GB of RAM, running Java 1.5.0.07 and Eclipse 3.3M5.

```

ghosts.size = 8;
ghosts.get > Ghost;
class Ghost > (ghosts.get, Figure, GhostShape) {
  this.position > Random;
  this.fill > Random;
  this.position.time = time / 20
}

```



Figure 4: An example of how execution occurs during live editing in the SuperGlue programming environment.

that signal initialize time is needed in whether a change occurs programmatically or through editing. However, while editing, this latency acts to distract the programmer.

As with the other phases, all errors that occur during the execution phase are duly reported and then shunted off to the background as program execution continues. Errors that can occur during the execution phase include ambiguously connected signals, unconnected primitive signals, and custom errors that are detected in code behind, such as divide by zero. Unlike errors in the previous phases, which can be reported against the program’s source code, errors in the execution phase are reported against signal values and are recorded in a separate view. Beyond reporting errors, our prototype does not currently provide facilities for debugging why errors have occurred, which is left to future work. After an error is reported, a default value is used if necessary to allow program execution to continue.

For code behind to play well in SuperGlue’s live programming environment, it must report errors on signal values rather than throw exceptions, which would cause the program to stop executing. Code behind should also fall back on default behavior that indicates errors in the program’s execution, if possible. For example, the figure draw routine we have coded to draw 2D shapes will report an error and draw a figure as a rectangle if the figure is not a concrete shape. Such error recovery allows the programmer to ignore errors as convenient and focus on other parts of the program.

The movie in Figure 4 demonstrates how live programming updates execution. It is narrated as follows:

- (a) `pacMan`’s position is bound to the `keyboard` of the canvas it is displayed in. So that `pacMan`’s orientation corresponds to the direction it is moving in, its `rotation` signal is bound to its `position`. `pacMan` then moves as expected in a video game.
- (b) Ghosts for the game are created by defining a `ghosts` signal that extends the `Duplicator` class. Initially two

ghosts are created, which are drawn as ghost shapes on the canvas by extending the `Figure` and `Ghost` classes. Because both ghosts occupy the same default position, only one ghost is visible.

- (c) The `Ghost` class is defined as an open extension to the `ghosts.get` class. In this class, the `position` and `fill` signals of each ghost extend the `Random` class, causing each ghost to have a random color and motion.
- (d) Finally, the number of ghosts is changed from two to eight, causing six more ghosts to appear after around three seconds of edit latency.

4. Preliminary Experience

Although our SuperGlue prototype is still immature, it can already be used to write complete programs. Previous work [21] demonstrates how a previous prototype was used to build rich user interfaces with complicated widgets such as tables and trees. This section describes our experience with implementing a simple game in SuperGlue using a library that wraps Java2D components. The game is a simplified version of PacMan with the following features:

- The game consists of a player controlled `pacMan` protagonist and multiple computer controlled `ghost` antagonists on a canvas.
- Normally, if `pacMan` is touched by a ghost, it is transported to the northwest corner of the canvas.
- In the middle of the canvas is a `powerPellet`. If `pacMan` eats/touches the `powerPellet`, the `powerPellet` disappears for some amount of time. During this time, if `pacMan` touches a ghost, the ghost is transported to the northwest corner of the canvas.

Aspects of this game have been used in the examples of this paper. We divide our discussion into three parts: code complexity and design (Section 4.1), live programming experience (Section 4.2), and performance (Section 4.3).

4.1 Code Complexity and Design

The code for our game can be divided into three categories: 58 lines of game-specific SuperGlue code that is listed in Figure 5, 225 lines of library SuperGlue code, and 5,750 lines of library-specific Scala code. No game-specific Scala code was required, and we do not in general expect SuperGlue programs to require custom Java or Scala code. To be complete in our discussion, our prototype is implemented in 23,000 lines of Scala code. Our prototype also depends on an IDE framework, under in-house development at the EPFL, that handles most scanning, parsing, and IDE service details in Eclipse, which consists of 28,000 lines of Scala code.

A large portion of the code in our animation library wraps Java2D (2,600 lines of Scala code). This animation library supports the rectangle and ellipse-like arc (Pie) shapes, Affine transforms (scaling, rotation, translation), and shape-


```

class GhostShape ▷ Shape {
  port gs ▷ Shape;
  gs = (Pie + Rectangle.scale(height = 0.495).translate(y = 25) -
    Eye.translate(x = 10) - Eye.translate(x = 30));
  this.shape = gs.shape;
  class Eye ▷ Shape;
  Eye.shape = Pie.scale(width = 0.2, height = 0.2).translate(y = 10).shape;
}
class PacManShape ▷ Shape {
  port gs ▷ Shape;
  gs = Pie(start = open / 2, extent = 360 - (open)).rotate(by = rotation.by);
  this.shape = gs.shape;
  port open ▷ Number;
  open = 60;
  port rotation ▷ Rotation;
}
class PowerPellet ▷ (Shape, Timed) {
  port pie ▷ Rectangle(arc.width = 10, arc.height = 10);
  this.shape = pie.scale(width = 0.3, height = 0.3).rotate(by = (time * 4)).shape;
  if (this ▷ Canvas.Figure) {
    this.position.x = 500;
    this.position.y = 100;
    this.fill = cyan;
    this.enable = invisible.on!.result;
  }
  port invisible ▷ Alarm;
  invisible.duration = 40;
}
class Execution ▷ (Canvas, Clock) {
  this.period = 100;
  port pacMan ▷ (PacManShape, Figure, Entity);
  pacMan.fill = yellow;
  pacMan.open = time % 30;
  pacMan.position = keyboard;
  pacMan.rotation = pacMan.position;
  port ghosts ▷ Duplicator;
  ghosts.size = 2;
  ghosts.get ▷ Ghost;
  class GhostHit;
  class Ghost ▷ (ghosts.get, GhostShape, Figure, Entity) {
    position ▷ Random;
    position.time = time / 20;
    fill ▷ Random;
    hit ▷ GhostHit;
    position.step = 1.9;
  }
  class Pellet ▷ (PowerPellet, Figure) {
    if (overlaps = pacMan.hit)
      invisible.enable = true;
  }
  if (pacMan.overlaps ▷ GhostHit) {
    pacMan.reset_pos = pellet.invisible.on!.result;
    pacMan.fill ▷ red;
  }
  if (Ghost.overlaps = pacMan.hit)
    Ghost.reset_pos = pellet.invisible.on;
  if (pellet.invisible.on = true) Ghost.fill.color = blue.color;
  class Entity ▷ Figure {
    position ▷ PositionCell;
    port reset_x ▷ position.p_x.set(to = 0, enable = reset_pos);
    port reset_y ▷ position.p_y.set(to = 0, enable = reset_pos);
    port reset_pos ▷ Boolean;
    reset_pos = false;
  }
  port pellet ▷ Pellet;
}
program ▷ Execution;

```

Figure 5: The SuperGlue code for a PacMan-like game; this code has been copied directly from the running version of the game.

area operations (union, intersection, subtraction). On top of this functionality, we can derive more complicated shapes; e.g., the `GhostShape` in Figure 5 is a union of an ellipse and translated circle, with two small ellipses cut out to form its eyes. Our Java2D wrapper also supports fill and line colors, keyboard input, and rudimentary collision detection. Many Java2D features have not yet been wrapped, so a more complete wrapping of Java2D will require more code. The other Scala-implemented parts of the library are devoted to arithmetic, timers, interpolation, random number generation, and number cells that can be updated. Number cells are especially important because they are currently the only form of memory that our library supports. All other library classes are built on top of these classes using only SuperGlue code. For example, a `PositionCell` class composes two number cell classes and extends the `Position` class to provide a mutable position, and a `PositionByKeyboard` class that builds on `PositionCell` to update this mutable position according to keyboard input.

Our library leverages SuperGlue’s rule-based abstractions to provide default connections whenever possible, which reduces the number of explicit connections that must be expressed in a program. For example, many library classes depend on timers: number cells use time signals to control how often they are incremented (velocity), random number generators use time signals to control how often they generate new random numbers, and interpolation transforms a time signal into a wave. These classes extend the `Timed` class whose `time` signal is connected by default to the `time` signal of a containing `Clock` signal. Then, by having the `Game` class extend the `Clock` class, any time dependent behavior that is defined inside the `Game` class is by default connected from the game’s clock.

We have found mutual extensions (Section 2) and adapter mixins (Section 2.2) to be very useful and repeatedly applicable in the design of our animation library. The `PositionByKeyboard` adapter mixin has already been mentioned in Section 2.2. Additionally, the `PositionByRandom` class adapts the `Position` and `Random` classes to move figures in random patterns, the `ColorByRandom` class adapts the `Color` and `Random` classes to generate random colors, while the `RotationByPositionCell` class adapts the `Rotation` and `PositionCell` classes to compute the rotation of a shape according to the direction it is moving in. Mutual extensions allow us to create rich versions of classes that are only applied in restricted contexts. For example, numbers only support basic arithmetic operations by default. However, when used in the context of a `time` signal, it is often useful that numbers (including literals) also support interpolation and random number generation operators. By defining a `Number` extended `RichNumber` class inside the `Clock` class, programmers can write “`0 <> 360`” to express a value that oscillates between 0 and 360 according to the containing `Clock` signal, or “`0 # 256`” to express a random value between 0 and 256

that also changes according to the containing `time` signal. Type inference computes types that ensures these operators are only accessible on numbers (including literals) that are used in the body of a class that extends the `Clock` class.

Given the design of our animation library, the game-specific code in Figure 5 is very concise and limited to configuration tasks, e.g., how many ghosts, establishing core game relationships, e.g., that `pacMan` is controlled by the keyboard, and core game logic; e.g., resetting `pacMan` or a ghost’s position when they collide. All adaptation code between non game-specific classes has been counted as part of the library. Admittedly, programs might require adaptations between library classes that do not already exist, but these adaptations can be published or at least added to the developers own library for reuse in future programs.

4.2 Live Programming

The live programmer experience is difficult to describe on paper, which is why this paper has included movies that show live programming in action. However, edit latency can measurably affect the live programming experience. Most edits used in the construction of our game do not incur noticeable edit latency. However, as discussed in Section 3, there are situations where noticeable and sometimes very long pauses can occur; e.g., by causing many signals to extend a new class or creating many different signals at once. One informal benchmark that we use in our game is based on ghost creation latency: how long does it take to duplicate N ghosts? The good news: the edit latency of ghost creation is linear—creating $2N$ ghosts takes only twice as long as creating N ghosts. The bad news: it take around 500 milliseconds to duplicate one ghost. The ugly news: creating ten ghosts will waste five seconds of the programmer’s life. Significant edit latency can also occur when type inference and circuit caches are cold. For example, when a ghost is first created during a game, around two seconds are needed to build the ghost’s circuits. Fortunately, all future ghosts obtain these circuits from a cache.

Another part of the live programming experience to consider is error detection and debugging. Our current prototype only supports the reporting of errors, and does not provide any support for debugging beyond providing responsive feedback from the executing program. The value of responsive feedback as a debugging aid cannot be understated: being able to see at one moment on one screen how an edit affects the program’s execution allows the programmer to very quickly determine if the edit was good or not. Errors often do not accumulate as they would if the programmer waited until after many edits to test the program. Additionally, because they lack control flow, programs have simpler execution models that do not involve call stacks, instruction stepping, and so on. However, because of SuperGlue’s support for object-oriented dispatch of rules, code is substantially more abstract than execution; i.e., execution errors need not be directly related to one piece of code. For example, if a

signal is connected ambiguously, the involved rules could be scattered throughout the program.

4.3 Performance

More than most dynamic languages, SuperGlue eschews performance in favor of being higher-level, more flexible, and supporting live programming. Although edit latency is important to live programming, overall program performance is less so. However, there is the danger that SuperGlue could be too slow for many kinds of applications. One area where our prototype's performance could be improved is in application start-up. Because of our prototype's reliance on incremental work list processing, much work is done and thrown away because many nodes are processed before their dependencies are "ready." For this reason, on start up, our game requires around 32 seconds to start-up: two seconds for parsing, ten seconds for type inference, and 20 seconds to perform initial execution processing. Our focus on edit latency, which work lists address, has led us to poor batch performance, and future work should focus on fixing this.

Another way of measuring performance in our game is through frame rate. For our game, we can measure frame rate as related to ghost count: around 100 frames per second is possible for `pacMan` plus one or two ghosts, ten frames per second is possible for ten ghosts, and five frames per second for twenty ghosts. Again, although such performance might be acceptable for a simple game, it indicates that our implementation needs better performance so that SuperGlue can be used more widely.

5. Related Work

SuperGlue is strongly influenced by Self [30]. Every artifact of a Self program is an object whose every aspect, including what they extend, can be updated at run-time. Self supports no abstractions other than objects: class or method constructions are expressible as objects. SuperGlue is more conventional: behavior is expressed through distinct class constructs. Self supports code hot swapping and responsive live programming through a graphical "meta menu" that can directly edit specific objects [31, 19]. However, live programming does not apply to Self source code, and direct changes to objects are not copied back into the code of the program. In SuperGlue, all program behavior is derived from code that can undergo live programming.

Because dynamic inheritance in Self occurs through imperative slot assignment, the type of an object can change in any way at any time without the benefit of static typing. Dynamic object inheritance is supported with static typing in Cecil [6, 7] through declarative predicates, which are also used in SuperGlue. However in SuperGlue, reactive signals make predicated class extensions more responsive: because predicates refer to signals, when a signal extends or unextends a class is known immediately. In Cecil, program be-

havior cannot responsively react to changes in class extensions because predicates can only be checked on demand.

One of the main contributions of this paper is showing how visual language liveness can be realized in a textual language. Ever since SketchPad [28], visual languages have emphasized simple interactive programming models that accommodate end users or novices. SuperGlue is designed for programmers who typically use dynamic languages, and end-user programming is not a goal. Many modern visual languages solve the scaling up problem in ways that do not degrade live programming; e.g., assemblies in Fabrik [16] are reused as objects, and data types in Forms/3 [3] can be elegantly specified as generic spreadsheet forms. SuperGlue solves the scaling-up problem using conventional class constructs that are more suited to experienced programmers who are using textual languages. AgentSheets [25] is a visual spreadsheet language that is based on graphical rewrite rules and implicit component assembly through spatial adjacency. AgentSheets demonstrates how live languages can be based on models other than reactive data-flow, and how many visual language concepts (spatial adjacency) may not transfer easily to text.

Many existing textual languages are based on reactive data-flow programming models. Signal [2], Lucid [33], and LUSTRE [5] are textual reactive data-flow languages that emphasize the formal verification of concurrency in real-time systems. As a result, they are not designed to support the rapid development capabilities of a dynamic or live language. SuperGlue is very similar to functional-reactive languages such as Fran [10], Yampa [14], and FrTime [8]. Such languages augment the functional paradigm with reactive signals; e.g., the result of a function call $f(x)$ forms a signal that changes whenever the x signal changes. In contrast, SuperGlue is based on rules, which are less flexible but have the advantage that they do not need to be applied explicitly. It remains an open question whether functional-reactive languages can support responsive live programming in the presence of recursive function applications.

Subtext [9] bridges the gap between code and execution by allowing programmers to edit code that directly represents program executions that lacks textual names. Because edits operate directly on execution, Subtext supports responsive live programming. On the other hand, Subtext supports scaling through copying, which has yet to be shown effective in practice. In SuperGlue, code and execution remain separate and their gap is only bridged through live programming.

Flogo II [13] is another textual live language that focuses on end-user robot programming. Flogo II supports both reactive data-flow and procedural code, where the latter does not support responsive live programming. Given its focus on end users, Flogo II scaling constructs (actor-like processes) are more concrete and do not scale as well as SuperGlue's classes. An intriguing idea is Flogo II's support for *live text*: the state of an executing program is presented as graphical

annotations in the program's textual source code. For example, statements are grayed out when they are guarded by a condition that is currently false. SuperGlue does not currently support similar capabilities.

SuperGlue's static type system relies on type inference and dependent typing to provide programmers with semantic feedback while avoiding verbose type annotations. Scala [23] is a statically typed object-oriented language that utilizes dependent typing and type inference to reduce verbosity, and therefore is a viable alternative to dynamic languages. SuperGlue's type system is influenced by Scala's; however, dynamic typing is still used to validate connections. An alternative way of obtaining semantic feedback for a dynamic language is through program analysis; e.g., Olin Shivers describes various analyses for inferring types in Scheme [27]. However, global program analysis is a hard problem and so is often not very accurate or responsive. Type inference in SuperGlue relies on a very simple program analysis that remains accurate while being both flow and context insensitive, because SuperGlue code has neither control flow nor procedures.

6. Conclusions and Future Work

Programmer-centric dynamic languages have been around for awhile: Lisp and Smalltalk are respectively more than 40 and 25 years old. New dynamic languages, such as Ruby [20] and Python [32], still strongly resemble Lisp and Smalltalk in their feature sets. It is now time for new live languages that are more programmer friendly by supporting declarative live programming with responsive semantic feedback. This paper demonstrates the viability of both the design and implementation of such a live language that also supports textual naming and abstraction. We have shown that a live language should be based on a simple declarative programming model, i.e., reactive data-flow, that is augmented with declarative scaling abstractions, i.e., predicated object-oriented inheritance and dispatch. We believe that a live textual languages could also be based on other simple declarative computational models such as constraints, first-order logic, and rewrite rules, for which live visual languages already exist.

Future work will focus on improving SuperGlue's design and implementation in the following areas:

- SuperGlue's support for multi-element data, such as arrays and databases, needs to be fleshed out. We plan to enhance signals so that they can represent time-varying tables of values that support relational operators.
- Our prototype does not currently provide enough debugging support. We plan to allow programmers to inspect, reason about errors in, and edit specific signals.
- We must improve on latency and performance if the widespread use of SuperGlue is to be viable. Building a live programming environment is an undocumented

black art that requires a complete re-think of how code is scanned, parsed, typed, and executed. The prototype described in this paper has already undergone some tuning and there are still many improvements that we can make.

- We should explore the incorporation of other visual and interactive features into the SuperGlue's programming environment. Features to be explored include support for direct manipulation, where specific signal values can be changed directly, and live text, where source code is annotated with execution details.

Finally, we plan to support SuperGlue with rich libraries that will make it a good platform for building interactive programs in an interactive way.

References

- [1] J. L. Armstrong and R. Viriding. An experimental telephony switching language. In *Proc. of Interantation Switching Symposium*, May 1991.
- [2] A. Benveniste, P. L. Geurnic, and C. Jacquemot. Synchronous programming with events and relations: the Signal language and its semantics. In *Science of Computer Programming*, 1991.
- [3] M. Burnett, J. Atwood, R. Walpole, H. Gottfried, J. Reichwein, and S. Yang. Forms/3: A first-order visual language to explore the boundaries of the spreadsheet paradigm. In *Journal of Functional Programming*, pages 155–206, Mar. 2001.
- [4] M. M. Burnett, M. J. Baker, C. Bohus, P. Carlson, S. Yang, and P. van Zee. Scaling up visual programming languages. In *IEEE Computer*, pages 45–54, Mar. 1995.
- [5] P. Caspi, D. Pilaud, N. Halbwachs, and J. A. Plaice. LUSTRE: a declarative language for programming synchronous systems. In *Proc. of POPL*, 1987.
- [6] C. Chambers. Object-oriented multi-methods in Cecil. In *Proc. of ECOOP*, pages 33–56, June 1992.
- [7] C. Chambers. Predicate classes. In *Proc. of ECOOP*, pages 268–296, July 1993.
- [8] G. H. Cooper and S. Krishnamurthi. Embedding dynamic dataflow in a call-by-value language. In the Proceedings of ESOP, 2006.
- [9] J. Edwards. Subtext: Uncovering the simplicity of program. In *Proc. of OOPSLA Onward*, 2005.
- [10] C. Elliott and P. Hudak. Functional reactive animation. In *Proc. of ICFP*, volume 32 (8) of *SIGPLAN Notices*, pages 263–273. ACM, 1997.
- [11] R. W. Floyd. Syntactic analysis and operator precedence. In *Journal of the ACM*, volume 10 (3), pages 316–333, 1963.
- [12] A. Goldberg and D. Robson. *SmallTalk-80: The Language and its Implementation*. Addison Wesley, Boston, MA, USA, 1983.
- [13] C. M. Hancock. *Real-time Programming and the Big Ideas of Computational Literacy*. PhD thesis, Massachusetts Institute of Technology, Sept. 2003.

- [14] P. Hudak, A. Courtney, H. Nilsson, and J. Peterson. Arrows, robots, and functional reactive programming. In *Advanced Functional Programming*, volume 2638 of *Lecture Notes in Computer Science*, pages 159–187. Springer, 2002.
- [15] IBM. *The Eclipse Project*. <http://www.eclipse.org/>.
- [16] D. Ingalls, S. Wallace, Y.-Y. Chow, F. Ludolph, and K. Doyle. Fabrik, a visual programming environment. In *Proc. of OOPSLA*, pages 176–190, Nov. 1988.
- [17] R. M. Karp and R. E. Miller. Properties of a model for parallel computations: Determinacy, termination queueing. *SIAM Journal for Applied Mathematics*, 14(6):1390–1410, Nov. 1966.
- [18] P. O. Latour. *Quartz Composer*. Apple Computer, 2005. <http://developer.apple.com/graphicsimaging/quartz/quartzcomposer.html>.
- [19] J. H. Maloney and R. B. Smith. Directness and liveness in the Morphic user interface construction environment. In *ACM Symposium on User Interface Software and Technology*, pages 21–28. ACM, 1995.
- [20] Y. Matsumoto. *Ruby: Programmers' Best Friend*. <http://www.ruby-lang.org/en/>.
- [21] S. McDirmid and W. C. Hsieh. SuperGlue: Component programming with object-oriented signals. In *Proc. of ECOOP*, June 2006.
- [22] National Instruments Corporation. *LabVIEW User Manual*, 1990.
- [23] M. Odersky and et. al. The Scala language specification. Technical report, EPFL, Lausanne, Switzerland, 2007. <http://scala.epfl.ch>.
- [24] M. Petre. Why looking isn't always seeing: Readership skills and graphical programming. *Comm. of the ACM*, pages 33–44, June 1995.
- [25] A. Repenning. AgentSheets: an interactive simulation environment with end-user programmable agents. In *Proc. of Interaction*, 2000.
- [26] E. Sandewall. Programming in an interactive environment: the lisp experience. *Computing Surveys*, 10(1), Mar. 1978.
- [27] O. Shivers. The semantics of Scheme control-flow analysis. In *Proc. of PEMP*, pages 190–198, June 1991.
- [28] I. B. Sutherland. SKETCHPAD, a man-machine graphical communication system. In *Proc. of the Spring Joint Computer Conference*, pages 329–346, 1963.
- [29] T. Teitelbaum and T. W. Reps. The Cornell program synthesizer: A syntax-directed programming environment. *Communications of the ACM*, 24(9):563–573, 1981.
- [30] D. Ungar and R. B. Smith. Self: the power of simplicity. In *Proc. of OOPSLA*, pages 227–242, Oct. 1987.
- [31] D. Ungar and R. B. Smith. Programming as an experience: the inspiration for Self. In *Proc. of ECOOP*, pages 227–242, Aug. 1995.
- [32] G. van Rossum and F. L. Drake. *The Python Language Reference Manual*, Sept. 2003.
- [33] W. W. Wadge and E. A. Ashcroft. *Lucid, the Dataflow Programming Language*. Academic Press, 1985.