

Technical Dimensions of Programming Systems

Joel Jakobovic^a, Jonathan Edwards^b, and Tomas Petricek^c

a University of Kent, Canterbury, UK jdj9@kent.ac.uk

b jonathanmedwards@gmail.com

c University of Kent, Canterbury, UK T.Petricek@kent.ac.uk

Abstract Many programming systems go beyond programming languages. Programming is usually done in the context of a stateful environment, beyond just writing code, by interacting with a system through a graphical user interface. Much research effort focuses on building programming systems that are easier to use, accessible to non-experts, moldable and/or powerful, but such efforts are often disconnected. They are informal, guided by the personal vision of the authors and thus are only evaluable and comparable on the basis of individual experience using them. They fail to form a coherent body of research, since it is unclear how to build on past work. In the research world, much has been said and done that allows comparison of *programming languages*, yet no similar theory exists for *programming systems*; we believe that programming systems deserve a theory too. We examine some influential past programming systems and review their stated design principles, technical capabilities, and styles of user interaction. We propose a framework of *technical dimensions* which capture the underlying system characteristics and provide a means for conceptualizing and comparing programming systems. Since these characteristics may be compared or advanced independently, it should be easier to talk about programming systems in a way that can be shared and constructively debated rather than relying solely on personal impressions. By providing foundations for more systematic research in this area, we can help the designers of future programming systems to stand, at last, on the shoulders of giants.

NOTE TO REVIEWERS: The main contribution of the paper is a comprehensive survey of 21 design dimensions of programming systems. Each one comes with detailed discussion, including known uses, to properly motivate it. We also elect to include the *full set* of these dimensions in the main text. While these two aspects do make the paper longer than the recommended page limit, we prefer to err on this side than to privilege certain dimensions and relegate others to an appendix. We invite the reader to focus on the ones that best suit their interest.

ACM CCS 2012

- **Software and its engineering** → **Integrated and visual development environments;**
- *Human-centered computing* → *Interactive systems and tools; HCI design and evaluation methods;*

Keywords Programming Systems, Dimensions, Design, Framework, Analysis

The Art, Science, and Engineering of Programming

Perspective The Art of Programming

Area of Submission Programming systems



© Joel Jakobovic, Jonathan Edwards, and Tomas Petricek
This work is licensed under a “CC BY 4.0” license.
Submitted to *The Art, Science, and Engineering of Programming*.

Technical Dimensions of Programming Systems

A systematic presentation removes ideas from the ground that made them grow and arranges them in an artificial pattern.

— Paul Feyerabend, *The Tyranny of Science*, Polity Press (2011)

Irony is said to allow the artist to continue his creative production while immersed in a sociocultural context of which he is critical.

— Emmanuel Petit, *Irony or, the Self-Critical Opacity of Postmodernist Architecture*, Yale (2013)

1 Introduction

Many forms of software have been developed to enable programming. The classic form consists of a *programming language*, a text editor to enter source code, and a compiler to turn it into an executable program. Instances of this form are differentiated by the syntax and semantics of the language, along with the implementation techniques in the compiler or runtime environment. Since the advent of graphical user interfaces (GUIs), programming languages can be found embedded within graphical environments that increasingly define how programmers work with the language (by directly supporting debugging or refactoring, for instance.) Beyond this, the rise of GUIs also permits diverse visual forms of programming, including visual languages and GUI-based end-user programming tools. This paper relies on, and encourages, a shift of attention from *programming languages* to the more general notion of “software that enables programming”—in other words, *programming systems*.

A *programming system* may include tools, protocols, notations, and languages. It is a software artifact that makes it possible to construct programs, debug them, and turn them into operational, maintained, and evolvable artifacts running on appropriate hardware. This notion covers classic programming languages together with their editors, debuggers, compilers, and other tools. Yet it is intentionally broad enough to accommodate image-based programming environments like Smalltalk, operating systems like UNIX, and hypermedia authoring systems like Hypercard, in addition to various other examples we will mention.

1.1 What is the problem?

There is a growing interest in broader forms of *programming systems*, both in the programming research community and in industry. On the one hand, researchers are increasingly studying topics such as *programming experience* and *live programming* that require considering not just the *language*, but further aspects of a given system. On the other hand, commercial companies are building new programming environments like Replit¹ or low-code programming tools like Dark² and Glide.³ Yet, such topics remain

¹ <https://replit.com/>

² <https://darklang.com/>

³ <https://www.glideapps.com/>

at the sidelines of mainstream programming research. While *programming languages* are a well-established concept, analysed and compared in a common vocabulary, no similar foundation exists for the wider range of *programming systems*.

The academic research on programming suffers from this lack of common vocabulary. While we can thoroughly assess programming languages, as soon as we add interaction or graphics into the picture, we often get stuck on how the resulting system is vaguely “cool” or “interesting”. Moreover, when designing new systems, inspiration is often drawn from the same few standalone sources of ideas. These might be influential past systems like Smalltalk, programmable end-user applications like spreadsheets, or motivational illustrations by thinkers like Victor [66].

Instead of forming a solid body of work, the ideas that emerge are difficult to relate to each other. Similarly, the research methods used to study programming systems lack the more rigorous structure of programming language research methods. They tend to rely on singleton examples, which demonstrate the author’s ideas, but are inadequate methods for comparing new ideas with the work of others. This makes it hard to build on top and thereby advance the state of the art.

Studying *programming systems* is not merely about taking a programming language and looking at the tools that surround it. It presents a *paradigm shift* to a perspective that is, at least partly, *incommensurable* with that of languages. When studying programming languages, everything that matters is in the program code; when studying programming systems, everything that matters is in the *interaction* between the programmer and the system. As documented by Gabriel [21], looking at a *system* from a *language* perspective makes it impossible to think about concepts that arise from interaction with a system, but are not reflected in the language. Thus, we must proceed with some caution. As we will see, when we talk about Lisp as a programming system, we mean something very different from a parenthesis-heavy programming language!

1.2 Contributions

We propose a new common language as an initial, *tentative* step towards more progressive research on programming systems. Our set of “Technical Dimensions for Programming Systems” seeks to break down the holistic view of systems along various specific “axes” inspired by the approach of the *Cognitive Dimensions of Notation* [24]. While not strictly quantitative, we have designed them to be narrow enough to be comparable, so that we may say one system has more or less of a property than another. Generally, we see the various possibilities as tradeoffs and are reluctant to assign them “good” or “bad” status. If the framework is to be useful, then it must encourage some sort of rough consensus on how to apply it; we expect it will be more helpful to agree on descriptions of systems first, and settle normative judgements later.

The set of dimensions can be understood as a map of the design space of programming systems (Figure 1). Past and present systems will serve as landmarks, and with enough of them, unexplored or overlooked possibilities will reveal themselves. So far, the field has not been able to establish a virtuous cycle of feedback; it is hard for practitioners to situate their work in the context of others’ so that subsequent work

Technical Dimensions of Programming Systems

can improve on it. Our aim is to provide foundations for the study of programming systems that would allow such development.

1. We present the dimensions in detail, organised into related clusters: *interaction*, *notation*, *conceptual structure*, *customizability*, *automation*, *errors*, and *adoptability*.
2. We define these dimensions by reference to landmark programming systems of the past, and discuss any relationships between them.
3. We demonstrate the salience of these dimensions by applying them to example systems from both the past and present. We situate some experimental systems as explorations at the frontier of certain dimensions.

2 Related work

While we do have new ideas to propose, part of our contribution is integrating a wide range of *existing* concepts under a common umbrella. This work is spread out across different domains, but each part connects to programming systems or focuses on a specific characteristic they may have.

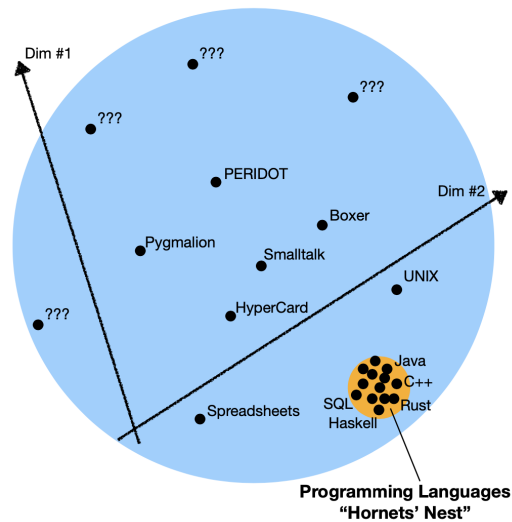
2.1 Which “systems” are we talking about?

The programming systems that shape our framework come from a few recognisable clusters:

- “Platforms” supporting arbitrary software ecosystems: UNIX, Lisp, Smalltalk, the Web
- “Applications” targeted to a specific domain: spreadsheets
- Mixed aspects of platform and application: HyperCard, Boxer, Flash, and programming language workflows

Richard Gabriel noted a “paradigm shift” [21] from the study of systems to the study of languages in computer science, which informs our distinction here. One consequence of the change is that a *language* is often formally specified apart from any specific implementations, while *systems* resist formal specification and are often *defined by* an implementation. We do, however, intend to recognize programming language implementations as a *small region* of the space of possible systems (Figure 1). Hence we refer to the *interactive programming system* aspects of languages, such as text editing and command-line workflow.

Our “system” concept is mostly technical in scope, with occasional excursions as in “Adoptability” (Section 4.7). This contrasts with the more socio-political focus found in [64]. It overlaps with Kell’s conceptualization of UNIX, Smalltalk, and Operating Systems generally [33], and we have ensured that UNIX has a place in our framework.



■ **Figure 1** A speculative sketch of one 2-dimensional slice of the space of possible systems.

2.2 Industry and research interest in programming systems

There is renewed interest in programming systems in both industry and research. In industry we see:

- Computational notebooks such as Jupyter⁴ that make data analysis more amenable to scientists by combining code snippets and their numerical or graphical output in a convenient document format.
- “Low code” end-user programming systems that present a simplified GUI for developing applications. One example is Coda,⁵ which combines tables, formulas, and scripts to enable non-technical people to build “applications as a document”.
- Specialized programming systems that augment a specific domain. For example Dark, which creates cloud API services with a “holistic” programming experience including a language and direct manipulation editor with near-instantaneous building and deployment.
- Even for general purpose programming with conventional tools, systems like Replit have demonstrated the benefits of integrating all needed languages, tools, and user interfaces into a seamless experience available from a browser with no setup.

In research, there are an increasing number of explorations of the possibilities of full programming systems:

- Subtext [12], which combines code with its live execution in a single editable representation.

⁴ <https://jupyter.org/>

⁵ <https://coda.io/welcome>

Technical Dimensions of Programming Systems

- Sketch-n-sketch [27], which can synthesize code by direct manipulation of its outputs.
- Hazel [48], a live functional programming environment featuring typed holes which enable execution of incomplete or type-erroneous programs.

Several research venues investigate programming systems:

- UIST (ACM Symposium on User Interface Software and Technology)
- VL/HCC (IEEE Symposium on Visual Languages and Human-Centric Computing)
- The LIVE programming workshop at SPLASH
- The PX (Programming eXperience) workshop at (Programming)

2.3 Characteristics already identified elsewhere

There are several existing projects identifying characteristics of programming systems. Some of these revolve around a single one, such as levels of liveness [63], or plurality and communicativity [34]. Others propose, as we do here, an entire collection:

- *Memory Models of Programming Languages* [58] identifies the “everything is an X” metaphors underlying many programming languages.
- The *Design Principles of Smalltalk* [30] documents the philosophical goals and dicta used in the design of Smalltalk.
- The “Gang of Four” *Design Patterns* [22] names and catalogues specific tactics within the *codebases* of software systems.
- The *Cognitive Dimensions of Notation* [24] introduces a common vocabulary for software’s *notational surface* and shows how they trade off and affect the performance of certain types of tasks.

Of these sources, the latter two bear the most obvious influence on our proposal. Our framework of “technical dimensions” continues the approach of the Cognitive Dimensions to the “rest” of a system beyond its notation. Our individual dimensions naturally fall under larger *clusters* that we present in a regular format, similar to the presentation of the classic Design Patterns. As for characteristics identified by others, part of our contribution is to integrate them under a common umbrella: liveness, pluralism, and uniformity metaphors (“everything is an X”) are incorporated as dimensions already identified by the related work.

We follow the attitude of *Evaluating Programming Systems* [13] in distinguishing our work from HCI methods and empirical evaluation. We are generally concerned with characteristics that are not obviously amenable to statistical analysis (e.g. mining software repositories) or experimental methods like controlled user studies, so numerical quantities are generally not featured.

Similar development seems to be taking place in HCI research focused on user interfaces. The UIST guidelines⁶ instruct authors to evaluate system contributions holistically, and the community has developed heuristics for such evaluation, such as

⁶ <https://uist.acm.org/uist2021/author-guide.html>

Evaluating User Interface Systems Research [47]. Our set of dimensions offers similar heuristics for identifying interesting aspects of programming systems, though they focus more on underlying technical properties than the surface interface.

Finally, we believe that the aforementioned paradigm shift from programming systems to programming languages has hidden many ideas about programming that are worthwhile recovering and developing further [53]. Thus our approach is related to the idea of *complementary science* developed by Chang [7] in the context of history and philosophy of science. Chang argues that even in disciplines like physics, superseded or falsified theories may still contain interesting ideas worth documenting. In the field of programming, where past systems are discarded for many reasons besides empirical failure, Chang's *complementary science* approach seems particularly suitable.

2.4 What we are trying to achieve

In short, while there is a theory for programming languages, programming *systems* deserve a theory too. It should apply from the small scale of language implementations to the vast scale of operating systems. It should be possible to analyse the common and unique features of different systems, to reveal new possibilities, and to build on past work in an effective manner. In Kuhnian terms, it should enable a body of “normal science”: filling in the map of the space of possible systems (Figure 1), thereby forming a knowledge repository for future designers.

3 Programming systems

We intentionally use the term *programming system* to refer to a broad range of systems which are programmable to a varying degree. This section highlights a number of example families of programming systems, following a roughly chronological order. This serves three purposes. First, looking at a number of examples from the past helps build an intuitive understanding of what we mean by a programming system. Second, it allows us to introduce example systems that we will use in the next section to illustrate the individual technical dimensions. Third, studying the difference between systems in individual families is one way of identifying and motivating interesting technical dimensions.

3.1 Interacting with computers

The key aspect of computers that enabled the rise of programming systems was the ability for a programmer to interact one-on-one with a computer. This was not possible in the 1950s when most computers were large and operated in a batch-processing mode. Two historical developments enabled such interactivity from the 1960s. First, time-sharing systems enabled interactive shared use of a computer via a teletype. Second, smaller computers such as the PDP-1 and PDP-8 provided similar direct interaction, while 1970s workstations such as the Alto and Lisp Machines added graphical displays and mouse input.

Technical Dimensions of Programming Systems

3.1.1 Lisp

The Lisp programming language, in the form of LISP 1.5 [40], arrived before the rise of interactive computers. Nevertheless, the existence of an interpreter and the absence of declarations made it natural to use Lisp interactively, with the first such implementations appearing in the early 1960s. Two branches of the Lisp family,⁷ MacLisp and the later Interlisp, fully embraced the so-called “conversational” way of working. Interaction occurred through the teletype at first, later giving way to the screen and keyboard. Even on the teletype, the system incorporated a number of ideas that remain popular with programming systems today.

Both MacLisp and Interlisp adopted the idea of *persistent address space*. Both program code and program state were preserved when powering off the system, and could be accessed and modified interactively as well as programmatically using the *same means*. This idea appeared on time-sharing systems and culminated with the development of Lisp Machines, which embraced the idea that the machine runs continually and saves the state to disk when needed. Today, while this is still not the default state for systems running “natively” on some hardware, it is widely seen in cloud-based services like Google Docs, online IDEs, or virtual machine and container images.

One idea not widely seen today, yet pioneered in MacLisp and Interlisp, was the use of *structure editors*. These let programmers work with Lisp data structures not as sequences of characters, but as nested lists. In Interlisp, for example, the programmer would use commands such as *P to print the current expression, or *(2 (X Y)) to replace its second element with the argument (X Y). The PILOT system [65], later integrated into Interlisp, offered even more sophisticated conversational features. For typographical errors and other slips, it would offer an automatic fix for the user to interactively accept, modifying the program in memory and resuming execution.

3.1.2 Smalltalk

Smalltalk came on the scene in the 1970s, with the ambition of providing “dynamic media which can be used by human beings of all ages” [31]. The authors saw computers as *meta-media* that could become a range of other media for education, discourse, creative arts, simulation and other applications not yet invented. Smalltalk was designed for single-user workstations with a graphical display, and pioneered this display not just for applications but also for programming itself. This evolved over the history of Smalltalk. In Smalltalk 72, one wrote code in the bottom half of the screen. When editing a definition, the window became a structure editor logically similar to that of Lisp, but controlled using a mouse and menus instead of a teletype. Smalltalk 76 completed the transition from a terminal-based interface to a graphical interface, and introduced the *class browser* for navigating through classes and modifying their code.

Smalltalk shared a number of other characteristics with Lisp, although its key concept was one of *objects* and *message passing* rather than *lists*. They both differ

⁷ The Lisp family consists of several branches, including MacLisp, InterLisp, ZetaLisp, Common Lisp, Scheme, Racket, and Clojure. see [62]

from most modern programming systems by adopting the aforementioned persistent address space model of programming, where all objects remain in memory. Any changes made to the system state by programming or execution are preserved when the computer is turned off. The fact that much of the Smalltalk environment was implemented in itself, a property shared with many later Lisp systems, made it possible to significantly modify the system from within.

3.1.3 UNIX

Both Lisp and Smalltalk worked, to some extent, as operating systems. The user started their machine directly in the Lisp or Smalltalk environment and was able to do everything they needed from *within* the system.⁸ This explains why it is worth considering (especially programmer-oriented) operating systems as programming systems too. A prime example of this is UNIX, a 1970s operating system for time-sharing computers.

Many aspects of programming systems are shaped by their intended target audience. UNIX was built for computer hackers themselves and, as such, has its interface close to the machine. Although UNIX is historically closely linked to the C programming language, it developed a language-agnostic set of abstractions that make it possible to use multiple programming languages in a single system. While everything is an object in Smalltalk, the ontology of the UNIX system consists of files, memory, executable programs, and running processes. Interestingly, there is an explicit stage distinction here, not present in Smalltalk or Lisp: UNIX distinguishes between volatile *memory* structures, which are lost when the system is shut down, and non-volatile *disk* structures that are preserved. The ontology of files, however, enables an open pluralistic environment.

3.2 Application platforms

The previously discussed programming systems were either universal, in that they did not focus on any particular kind of application, or they were focused on broad application areas. Lisp, for example, was designed for symbolic data manipulation in the context of Artificial Intelligence, while FORTRAN focused on scientific computing. However, as computers became more widely used, it became clear that there are more narrow typical kinds of applications that need to be built. For those, specialized programming systems began to appear. Although they are more focused, they also support programming based on rich interactions with specialized visual and textual notations.

⁸When the Lisp and Smalltalk systems were implemented on specialized computers—InterLisp and Smalltalk on the Alto and Xerox D, ZetaLisp and Common Lisp on Lisp machines—the user would start their computers directly in the programming system environment. When implemented on commodity hardware, the user would resume a saved image of the system.

Technical Dimensions of Programming Systems

3.2.1 Spreadsheets

Spreadsheets, along with word processors, were the application that turned personal computers from playthings for hackers into a business tool. The first system, VisiCalc, became available in 1979 and helped analysts perform budget calculations. Spreadsheets developed over time, acquiring features that made them into powerful programming systems in a way VisiCalc was not. The final step was the 1993 inclusion of *macros* in Excel, later extended with *Visual Basic for Applications*. As programming systems, spreadsheets are notable for their programming substrate (a grid) and evaluation model (automatic re-evaluation).

3.2.2 HyperCard

While spreadsheets were designed to solve problems in a specific application area, the next system we consider was designed around a particular application format. 1987 saw HyperCard [42], with programs as “stacks of cards” containing multimedia components and controls such as buttons. These controls could be programmed with pre-defined operations like “navigate to another card”, or via the HyperTalk scripting language for anything more sophisticated.

As a programming system, HyperCard is interesting for a couple of reasons. It effectively combines visual and textual notation. Programs appear the same way during editing as they do during execution. Most notably, HyperCard supports gradual progression from the “user” role to “developer”: a user may first use stacks, then go on to edit the visual aspects or choose pre-defined logic until, eventually, they learn to program in HyperTalk.

3.3 Developer platforms

Programming systems such as Smalltalk and HyperCard are relatively *self-contained*. It is clear what is *part of* the system, and what is on the *outside*. For many systems that began to appear in the late 1980s, this is not the case. To think about them, we have to consider a number of components, some of which may be conventional programming languages. The boundaries of these *developer platforms* are less well-defined and we acknowledge that the exact delineation we choose significantly affects our analysis.

3.3.1 Early and late Web

The Web appeared in 1989 as a way of sharing and organizing information, implementing the ideas of *hypertext*. The web gradually evolved from an *information sharing* system to a *developer platform* when client-side scripting using JavaScript became possible. The Web ecosystem started to consist of server-side and client-side programming tools. Today, the Web combines the notations of HTML, CSS, a wide range of server-side programming systems as well as JavaScript, and many languages that compile to JavaScript.

In the 1990s, the “early Web” became a widely used programming system. JavaScript code was distributed in a form that made it easy to copy and re-use existing scripts, which led to enthusiastic adoption by non-experts. This is comparable to the birth of microcomputers like Commodore 64 with BASIC a decade earlier.

The Web ecosystem continued to evolve. In the 2000s, multiple programming languages started to treat JavaScript as a compilation target, while JavaScript started to be used as a programming language on the server-side. This defines the “late Web” ecosystem, which is quite different from its early incarnation. JavaScript code was no longer simple enough to whimsically copy and paste, yet advanced developer tools provided functionality resembling early interactive programming systems like Lisp and Smalltalk. The *Document Object Model (DOM)* structure created by a web page is transparent, accessible to the user and modifiable through the built-in browser debugging tools, and third-party code to modify the structure can be injected via extensions. In this, the DOM resembles the *persistent image* model. The DOM also inspired further research on image-based programming: Webstrates [37] synchronizes DOM edits made in the browser to all other clients connected to a single server.

3.3.2 REPLs and notebooks

Another kind of developer ecosystem which evolved from simple scripting tools consists of modern data science tools, such as Jupyter, whose roots date back to “On-Line Systems” developed in the 1960s. The style was exemplified in conversational implementations of Lisp, where users could type commands to be evaluated and see the results printed; this interaction became known as the REPL (Read-Eval-Print Loop). In the late 1980s, Mathematica 1.0 combined the REPL interaction with a notebook document format that showed the commands alongside visual outputs, an idea pioneered in work on *literate programming* [38].

Today, REPLs exist for many programming languages. Unlike in Lisp, they are often separate from the running program. REPLs often maintain an execution state independent of a running program and there are many strategies for prototyping code in a REPL before making it a part of an ordinary compiled application.

Notebooks for data science are a particularly interesting example. Their primary output is the notebook itself, rather than a separate application to be compiled and run. The code lives in a document format, interleaved with other notations. Code is written in small parts that are executed (almost) immediately, offering the user more rapid feedback than in conventional programming. A notebook can be seen as a trace of how the result has been obtained, yet one often problematic feature of notebooks is that some allow the user to run code blocks out-of-order! Thus, while a Common Lisp REPL user could just dribble their session to a file, retracing one’s steps in a notebook can be rather more subtle.

3.3.3 Haskell and other languages

The aforementioned 1990s paradigm shift from thinking about *systems* to thinking about *languages* means that researchers tend to emphasize the language side of programming. However, all programming languages are a part of a richer ecosystem that consist of editors and other tools. In our analysis, we choose Haskell as our example of a clearly *language-focused* programming system.

Like most programming languages, Haskell code can be written in a wide range of text editors, some of which support assistance tools such as syntax highlighting and

Technical Dimensions of Programming Systems

auto-completion. These offer immediate feedback while editing code, such as when highlighting type errors.

Haskell is mathematically rooted and relies on mathematical intuition for understanding many of its concepts. This background is also reflected in the notations it uses. In addition to the concrete language syntax (used when writing code), the Haskell ecosystem also uses an informal mathematical notation, which is used when writing about Haskell (e.g. in academic papers or on the whiteboard). This provides an additional tool for manipulating Haskell programs and experimenting with them on paper *in vitro*, in ways that other systems may attempt to achieve through experimentation within the system *in vivo*.

4 Technical dimensions

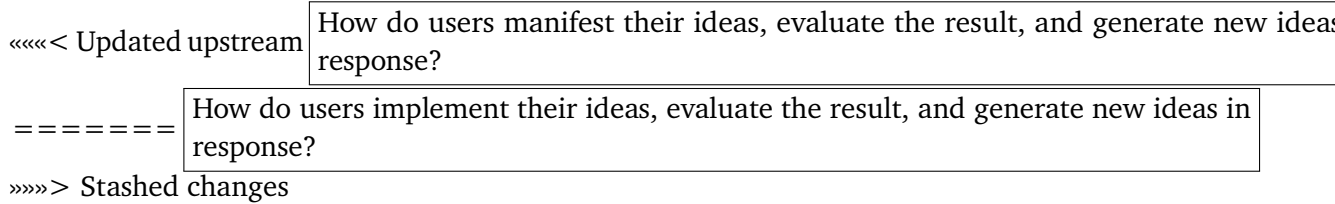
For the rest of this paper, we present our proposed technical dimensions grouped under *clusters*. The clusters may be regarded as “topics of interest” or “areas of inquiry” when studying a given system, grouping together related dimensions against which to measure it.

Each cluster is named and opens with a boxed *summary*, followed by a short *description*, and closes with a list of any *relations* to other clusters along with any *references* if applicable. Within the main description, individual *dimensions* are listed. Sometimes, a particular value along a dimension (or a combination of values along several dimensions) can be recognized as a familiar pattern—perhaps with a name already established in the literature. These are marked as *examples*. Finally, interspersed discussion that is neither a *dimension* nor an *example* is introduced as a *remark*. We include a concise reference sheet on the next page, though of course this cannot substitute for reading the relevant sections at least once.

Dimension (CLUSTER)	Summary	Range of key examples
INTERACTION	How do users manifest their ideas, evaluate the result, and generate new ideas in response?	
Feedback Loops	How wide are the various gulfs of <i>execution</i> and <i>evaluation</i> and how are they related?	Immediate Feedback (short) vs. batch mode (long) gulf of evaluation
Modes of Interaction	Which sets of feedback loops only occur together?	Setup vs. editing vs. debugging
Abstraction Construction	How do we go from abstractions to concrete examples and vice versa?	Programming by Example vs. first principles
NOTATION	How are the different textual / visual programming notations related?	
Notational Structure	What notations are used to program the system and how are they related?	Notations overlap and need sync vs. complement each other
Surface / Internal Notations	What is the connection between what a user sees and what a computer program sees?	Sequence Editing vs. Rendering, Structure Editing vs. Recovery
Primary / Secondary Notations	Is one notation more important than others?	Secondary build scripts vs. visual editor and code on equal footing in Flash
Expression Geography	Do similar expressions encode similar programs?	Concise yet error-prone vs. explicit yet verbose
Uniformity of Notations	Does the notation use a small or a large number of basic concepts?	Lisp S-expressions vs. English-like textual notations
CONCEPTUAL STRUCTURE	How is meaning constructed? How are internal and external incentives balanced?	
Conceptual Integrity vs. Openness	Does the system present as elegantly <i>designed</i> or pragmatically <i>improvised</i> ?	Integrity (Everything is a X) vs. openness (compatible mixtures)
Composability	What are the primitives? How can they be combined to achieve novel behaviors?	Sequence, selection, repetition, function abstraction, recursion, logical connectives
Convenience	Which wheels do users not need to reinvent?	Small vs. expansive standard libraries
Commonality	How much is <i>common structure</i> explicitly marked as such?	Common structure is redundantly flattened vs. factored out
CUSTOMIZABILITY	Once a program exists in the system, how can it be extended and modified?	
Staging of Customization	Must we customize <i>running</i> programs differently to <i>inert</i> ones? Do these changes last beyond termination?	Source code vs. config files, Developer Tools tab, auto image-based persistence, scripting language
Externalizability	Which portions of the system's state can be referenced and transferred to/from it?	None (state is private) vs. all state exposed as human-legible, CSS-like addressing
Additive Authoring	How far can the system's behavior be changed by <i>adding</i> expressions?	None (requires power to change original) vs. full (anything can be overridden repeatedly)
Self-Sustainability	How far can the system's behavior be changed from within?	None (rely on external tools) vs. self-sufficient (contains everything needed)
AUTOMATION	How far does the system remove the need to spell out implementation in minute detail?	
Degrees of Automation	What part of program logic does not need to be explicitly specified?	Garbage collection (low-tech) vs. Prolog engine (hi-tech)
ERRORS	What does the system consider to be an <i>error</i> ? How are they prevented and handled?	
Error Detection	What errors can be detected in which feedback loops, and how?	Human inspection in live coding vs. partial automation in static typing
Error Response	How does the system respond when an error is detected?	Does it stop, recover automatically, ignore the error or ask the user how to continue?
ADOPTABILITY	How does the system facilitate or obstruct adoption by both individuals and communities?	
Learnability	What is the attitude towards the <i>learning curve</i> and what is the target audience?	HyperCard for the general public vs. FORTRAN for scientists
Sociability	What are the social and economic factors that make the system the way it is?	Funding, volunteers, code sharing, Q/A sites vs. documentation, sense of belonging

Technical Dimensions of Programming Systems

4.1 Interaction



An essential aspect of programming systems is how the user interacts with them when creating programs. Take the standard form of statically typed, compiled languages with straightforward library linking: here, programmers write their code in a text editor, invoke the compiler, and read through error messages they get. After fixing the code to pass compilation, a similar process might happen with runtime errors.

Other forms are yet possible. On the one hand, some typical interactions like compilation or execution of a program may not be perceptible at all. On the other hand, the system may provide various interfaces to support the plethora of other interactions that are often important in programming, such as looking up documentation, managing dependencies, refactoring or pair programming.

We focus on the interactions where programmer interacts with the system to construct a program with a desired behavior. To analyze those, we use the concepts of *gulf of execution* and *gulf of evaluation* from *The Design of Everyday Things* [46].

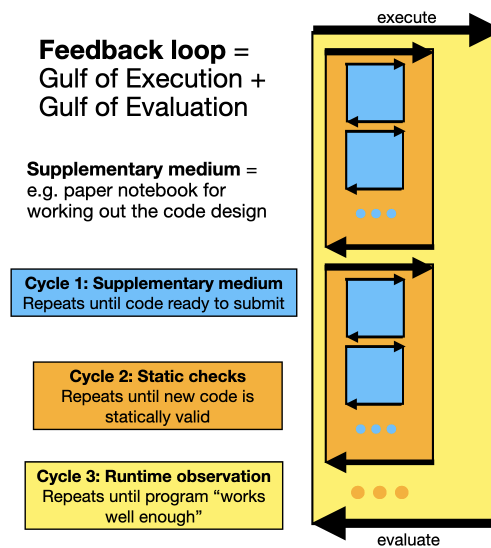
4.1.1 Dimension: feedback loops

In using a system, one first has some idea and attempts to make it exist in the software; the gap between the user's goal and the means to execute the goal is known as the *gulf of execution*. Then, one compares the result actually achieved to the original goal in mind; this crosses the *gulf of evaluation*. These two activities comprise the *feedback loop* through which a user gradually realises their desires in the imagination, or refines those desires to find out "what they actually want".

A system must contain at least one such feedback loop, but may contain several at different levels or specialized to certain domains. For each of them, we can separate the gulf of execution and evaluation as independent legs of the journey, with possibly different manners and speeds of crossing them.

For example, we can analyze statically checked *programming languages* (e.g. Java, Haskell) into several feedback loops (Figure 2):

1. Programmers often think about design details and calculations on a whiteboard or notebook, even before writing code. This *supplementary medium* has its own feedback loop, even though this is often not automatic.
2. The code is written and is then put through the static checker. An error sends the user back to writing code. In the case of success, they are "allowed" to run the program, leading into cycle 3.
 - The execution gulf comprises multiple cycles of the supplementary medium, plus whatever overhead is needed to invoke the compiler (such as build systems).



■ **Figure 2** The nested feedback loops of a statically-checked programming language.

- The evaluation gulf is essentially the waiting period before static errors or a successful termination are observed. Hence this is bounded by some function of the length of the code (the same cannot be said for the following cycle 3.)
3. With a runnable program, the user now evaluates the *runtime* behavior. Runtime errors can send the user back to writing code to be checked, or to tweak dynamically loaded data files in a similar cycle.
- The execution gulf here may include multiple iterations of cycle 2, each with its own nested cycle 1.
 - The *evaluation* gulf here is theoretically unbounded; one may have to wait a very long time, or create very specific conditions, to rule out certain bugs (like race conditions) or simply to consider the program as fit for purpose.
 - By imposing *static checks*, some bugs can be pushed earlier to the evaluation stage of cycle 2, reducing the likely size of the cycle 3 *evaluation* gulf.
 - On the other hand, this can make it harder to write statically valid code, which may increase the number of level-2 cycles, thus increasing the total *execution* gulf at level 3.
 - Depending on how these balance out, the total top-level feedback loop may grow longer or shorter.

4.1.2 Example: immediate feedback

The specific case where the *evaluation* gulf is minimized to be imperceptible is known as *immediate feedback*. Once the user has caused some change to the system, its effects (including errors) are immediately visible. This is a key ingredient of *liveness*, though it is not sufficient on its own. (See *Relations*)

The ease of achieving immediate feedback is obviously constrained by the computational load of the user's effects on the system, and the system's performance on such

Technical Dimensions of Programming Systems

tasks. However, such “loading time” is not the only way feedback can be delayed: a common situation is where the user has to manually ask for (or “poll”) the relevant state of the system after their actions, even if the system finished the task quickly. Here, the feedback could be described as *immediate upon demand* yet not *automatically demanded*. For convenience, we choose to include the latter criterion—automatic demand of result—in our definition of immediate feedback.

In a *REPL* or *shell*, there is a *main* cycle of typing commands and seeing their output, and a *secondary* cycle of typing and checking the command line itself. The output of commands can be immediate, but usually reflects only part of the total effects or even none at all. The user must manually issue further commands afterwards, to check the relevant state bit by bit. The secondary cycle, like all typing, provides immediate feedback in the form of character “echo”, but things like syntax errors generally only get reported *after* the entire line is submitted. This evaluation gulf has been reduced in the JavaScript console of web browsers, where the line is “run” in a limited manner on every keystroke. Simple commands without side-effects, such as calls to pure functions, can give instantly previewed results—though partially typed expressions and syntax errors will not trigger previews.

4.1.3 Example: direct manipulation

Direct manipulation [57] is a special case of an immediate feedback loop. The user sees and interacts with an artefact in a way that is as similar as possible to real life; this typically includes dragging with a cursor or finger in order to physically move a visual item, and is limited by the particular haptic technology in use.

Naturally, because moving real things with one’s hands does not involve any waiting for the object to “catch up”,⁹ direct manipulation is necessarily an immediate-feedback cycle. If, on the other hand, one were to move a figure on screen by typing new coordinates in a text box, then this could still give *immediate feedback* (if the update appears instant and automatic) but would *not* be an example of direct manipulation.

Spreadsheets contain a feedback loop for direct manipulation of values and formatting, as in any other WYSIWYG application. They also contain another loop for formula editing and formula invocation. Here, there is larger execution gulf for designing and typing formulas. This makes it an “immediate feedback” loop only *on-demand* as defined above.

4.1.4 Dimension: modes of interaction

The possible interactions in a programming system are typically structured so that interactions, and the associated feedback loops, are only available in certain *modes*. For example, when creating a new project, the user may be able to configure the project through a conversational interface like `npm init` in modern JavaScript. Such

⁹ In some situations, such as steering a boat with a rudder, there is a delay between input and effect. But on closer inspection, this delay is between the rudder and the boat; we do not see the hand pass through the wheel like a hologram, followed by the wheel turning a second later. In real life, objects touched directly give immediate feedback; objects controlled further down the line might not!

interactions are no longer available once the project is created. This idea of interaction modes goes beyond just programming systems, appearing in software engineering methodologies. In particular, having a separate *implementation* and *maintenance* phase would be an example of two modes.

Editing vs debugging. A good example is the distinction between *editing* and *debugging* mode. When debugging a program, the user can modify the program state and get (more) immediate feedback on what individual operations do. In some systems, one can even modify the program itself during debugging. Such feedback loops are not available outside of debugging mode.

Lisp systems sometimes distinguish between *interpreted* and *compiled* mode. The two modes do not differ just in the efficiency of code execution, but also in the interactions they enable. In the interpreted mode, code can be tested interactively and errors may be corrected during the code execution (see *Error response*). In the compiled mode, the program can only be tested as a whole. The same two modes also exist, for example, in some Haskell systems where the REPL uses an interpreter (GHCi) distinct from the compiler (GHC).

Jupyter notebooks. A programming system may also unify modes that are typically distinct. The Jupyter notebook environment does not have a distinct debugging mode; the user runs blocks of code and receives the result. The single mode can be used to quickly try things out, and to generate the final result, partly playing the role of both debugging and editing modes. However, even Jupyter notebooks distinguish between editing a document and running code.

4.1.5 Dimension: abstraction construction

A necessary activity in programming is going between abstract schemas and concrete instances. Abstractions can be constructed from concrete examples, first principles or through other methods. A part of the process may happen in the programmer's mind: they think of concrete cases and come up with an abstract concept, which they then directly encode in the system. Alternatively, a system can support these different methods directly.

One option is to construct abstractions *from first principles*. Here, the programmer starts by defining an abstract entity such as an interface in object-oriented programming languages. To do this, they have to think what the required abstraction will be (in the mind) and then encode it (in the system).

Another option is to construct abstractions *from concrete cases*. Here, the programmer uses the system to solve one or more concrete problems and, when they are satisfied, the system guides them in creating an abstraction based on their concrete case(s). In a programming language IDE this manifests as the “extract function” refactor, whereas in other systems we see approaches like macro recording.

Pygmalion. In Pygmalion [61], all programming is done by manipulating concrete icons that represent concrete things. To create an abstraction, you can use “Remember mode”, which records the operations done on icons and makes it possible to bind this recording to a new icon.

Jupyter notebook. In Jupyter notebooks, you are inclined to work with concrete things, because you see previews after individual cells. This discourages creating

Technical Dimensions of Programming Systems

abstractions, because then you would not be able to look inside at such a fine grained level.

Spreadsheets. Up until the recent introduction of lambda expressions into Excel, spreadsheets have been relentlessly concrete, without any way to abstract and reuse patterns of computation other than copy-and-paste.

4.1.6 Relations

- *Errors* (Section 4.6) A longer evaluation gulf delays the detection of errors. A longer execution gulf can increase the *likelihood* of errors (e.g. writing a lot of code or taking a long time to write it). By turning runtime bugs into statically detected bugs, the combined evaluation gulfs can be reduced.
- *Adoptability* (Section 4.7): The *execution* gulf is concerned with software using and programming in general. The time taken to realize an idea in software is affected by the user's familiarity and the system's *learnability*.
- *Notation* (Section 4.2): Feedback loops are related to *notational structures*. In a system with multiple notations, each notation may have different associated feedback loops. The motto "The thing on the screen is supposed to be the actual thing" [50], adopted in the context of live programming, relates *liveness* to a direct connection between surface and internal notations. The idea is that interactable objects should be equipped with faithful behavior, instead of being intangible shadows cast by the hidden *real* object.

4.2 Notation

How are the different textual / visual programming notations related?

Programming is always done through some form of notation. We consider notations in the most general sense and include any structured gesture using textual or visual notation. Textual notations primarily include programming languages, but also things like configuration files. Visual notations include graphical programming languages. Other kinds of structured gestures include user interfaces for constructing visual elements used in the system.

4.2.1 Dimension: notational structure

In practice, most programming systems use multiple notations. Different notations can play different roles in the system. On the one hand, multiple *overlapping notations* can be provided as different ways of programming the same aspects of the system. In this case, each notation may be more suitable to different kinds of users, but may have certain limitations (for example, a visual notation may have a limited expressive power). On the other hand, multiple *complementing notations* may be used as the means for programming different aspects of the system. In this case, programming the system requires using multiple notations, but each notation may be more suitable for the task at hand; think of how HTML describes document structure while JavaScript specifies its behavior.

4.2.2 Example: overlapping notations

A programming system may provide multiple notations for programming the same aspect of the system. This is typically motivated by an attempt to offer easy ways of completing different tasks: say, a textual notation for defining abstractions and a visual notation for specifying concrete structures. The crucial issue in this kind of arrangement is *synchronizing* the different notations; if they have different characteristics, this may not be a straightforward mapping. For example, source code may allow more elaborate abstraction mechanisms like loops, which will appear as visible repetition in the visual notation. What should such a system do when the user edits a single object that resulted from such repetition? Similarly, textual notation may allow incomplete expressions that do not have an equivalent in the visual notation. For programming systems that use *overlapping notations*, we need to describe how the notations are synchronized.

Sketch-n-Sketch [27] employs overlapping notations for creating and editing SVG and HTML documents. The user edits documents in an interface with a split-screen structure that shows source code on the left and displayed visual output on the right. They can edit both of these and changes are propagated to the other view. The code can use abstraction mechanisms (such as functions) which are not completely visible in the visual editor (an issue we return to in *expression geography* below). Sketch-n-Sketch can be seen as an example of a *projectional editor*.¹⁰

UML Round-tripping. Another example of a programming system that utilizes the *overlapping notations* structure are UML design tools that display the program both as source code and as a UML diagram. Edits in one result in automatic update of the other. An example is the Together/J¹¹ system. To solve the issue of notation synchronization, such systems often need to store additional information in the textual notation, typically using a special kind of code comment. In this example, after the user re-arranges classes in UML diagrams, the new locations need to be updated in the code.

4.2.3 Example: complementing notations

A programming system may also provide multiple complementing notations for programming different aspects of its world. Again, this is typically motivated by the aim to make specifying certain aspects of programming easier, but it is more suitable when the different aspects can be more clearly separated. The key issue for systems with complementing notations is how the different notations are connected. The user may need to use both notations at the same time, or they may need to progress from one to the next level when solving increasingly complex problems. In the latter case, the learnability of progressing from one level to the next is a major concern.

Spreadsheets and HyperCard. In Excel, there are three different complementing notations that allow users to specify aspects of increasing complexity: (i) the visual grid,

¹⁰ Technically, traditional projectional editors usually work more directly with the abstract syntax tree of a programming language.

¹¹ <https://www.mindprod.com/jgloss/togetherj.html>

Technical Dimensions of Programming Systems

(ii) formula language and (iii) a macro language such as Visual Basic for Applications. The notations are largely independent and have different degrees of expressive power. Entering values in a grid cannot be used for specifying new computations, but it can be used to adapt or run a computation, for example when entering different alternatives in What-If Scenario Analysis. More complex tasks can be achieved using formulas and macros. A user gradually learns more advanced notations, but experience with a previous notation does not help with mastering the next one. The approach optimizes for easy learnability at one level, but introduces a hurdle for users to surmount in order to get to the second level. The notational structure of *HyperCard* is similar and consists of (i) visual design of cards, (ii) visual programming (via the GUI) with a limited number of operations and (iii) HyperTalk for arbitrary scripting.

Boxer and Jupyter. *Boxer* [10] uses *complementing notations* in that it combines a visual notation (the layout of the document and the boxes of which it consists) with textual notation (the code in the boxes). Here, the textual notation is always nested within the visual. The case of Jupyter notebooks is similar. The document structure is graphical; code and visual outputs are nested as editable cells in the document. This arrangement is common in many other systems such as Flash or Visual Basic, which both combine visual notation with textual code, although one is not nested in the other.

4.2.4 Dimensions: surface notation and internal notation

All programming systems build up structures in memory, which we can consider as an *internal notation* not usually visible to the user. Even though such structures might be revealed in a debugger, they are hidden during normal operation. What the user interacts with instead is the *surface notation*, typically one of text or shapes on a screen. Every interaction with the surface notation alters the internal notation in some way, and the nature of this connection is worth examining in more detail. To do this, we illustrate with a simplified binary choice for the form of these notations.

4.2.5 Examples: implicit vs. explicit structure

Let us partition notations into two families. Notations with *implicit structure* present as a sequence of items, such as textual characters or audio signal amplitudes. Those with *explicit structure* present as a tree or graph without an obvious order, such as shapes in a vector graphics editor. These two types of notations can be transformed into each other: the implicit structure contained in a string can be *parsed* into an explicit syntax tree, and an explicit document structure might be *rendered* into a sequence of characters with the same implicit structure.

Now consider an interface to enter a personal name made up of a forename and a surname. For the surface notation, there could be a single text field to hold the names separated with a space; here, the sub-structure is implicit in the string. Alternatively, there could be two fields where the names are entered separately, and their separation is explicit. A similar choice exists for the internal notation built up in memory: is it a single string, or two separate strings?

We can see that these choices give four combinations. More interestingly, they exhibit unique characters owing to two key asymmetries. Firstly, surface notation is

mostly used by humans, while the internal notation is mostly used by the computer. Secondly, and most significantly, computer programs can only work with explicit structure, while humans can understand both explicit and implicit structure. Because of the practical consequences of this asymmetry, we will examine the combinations with emphasis on the *internal* notation first.

4.2.6 Examples: one string in memory (implicitly structured internal notation)

The simplest case here would be with implicit structure in the surface notation, i.e. a single text box for the full name. Edits to the surface are straightforwardly mirrored internally and persisted to disk. This corresponds to *text editing*. We can generalize this to an idea of *sequence editing* if we view the fundamental act as *recording* events to a list over time. For text, these are key presses; for an audio editing interface they would be samples of sound amplitude.

In the other case, with two text boxes, we have *sequence rendering*. The information about the separation of the two strings, present in the interface, is not quite “thrown away” but is made *implicit* as a space character in the string. This combination corresponds to Visual Basic generating code from GUI forms, video editors combining multiple clips and effects into a single stream, and 3D renderers turning scene graphs into pixels. Another example is line-based diff tools, which provide side-by-side views and related interfaces, yet must ultimately forward the user’s changes to the underlying text file.

Critically, in both of these cases, a computer program can only manipulate the stored sequences *as* sequences; that is, by inserting, removing, or serially reading. The appealing feature here is that these operations are simple to implement and may be re-usable across many types of sequences. However, any further structure is implicit and, to work with it programmatically, a user must write a program to *parse* it into something explicit. Furthermore, errors introduced at this stage may simply be *recorded* into the sequence, only to be discovered much later in an attempt to use the data.

4.2.7 Examples: two strings in memory (explicitly structured internal notation)

With two text boxes, both notations match, so there is not much work to do. As with sequence editing, edits on the surface can be mirrored to the internal notation. This corresponds to vector graphics editors and 3D modelling tools, as well as *structure editors* for programming languages. For this reason we call this combination *structure editing*.

With a single text field, we have *structure recovery*. Parsing needs to happen each time the input changes. This style is found in the DOM inspector in browser developer tools, where HTML can be edited as text to make changes to the document tree structure. More generally, this is the mode found in compilers and interpreters which accept program source text yet internally work on tree and graph structures. It is also possible to do a sort of structure editing this way, where the experience is made to resemble text editing but the output is explicitly structured.

In both of these cases, in order to write programs to transform, analyze, or otherwise work with the digital artefact the user has created, one can trivially navigate the

Technical Dimensions of Programming Systems

stored structure instead of parsing it for every use. Parsing is either done away with altogether or is reduced to a transient process that happens during editing; this means errors can be caught at the moment they are introduced instead of remaining latent.

4.2.8 Dimension: primary and secondary notations

In practice, most programming systems use multiple notations. Even in systems based on traditional programming languages, the *primary notation* of the language is often supported by *secondary notations* such as annotations encoded in comments and build tool configuration files. However, it is possible for multiple notations to be primary, especially if they are *overlapping* as defined earlier.

Programming languages. Programming systems built around traditional programming languages typically have further notations or structured gestures associated with them. The primary notation in UNIX is the C programming language. Yet this is enclosed in a programming *system* providing a multi-step mechanism for running C code via the terminal, assisted by secondary notations such as shell scripts. Some programming systems attempt to integrate tools that normally rely on secondary notations into the system itself, reducing the number of secondary notations that the programmer needs to master. For example, in the Smalltalk descendant Pharo, versioning and package management is done from within Pharo, removing the need for secondary notation such as git commands and dependency configuration files.¹²

Haskell. In Haskell, the primary notation is the programming language, but there are also a number of secondary notations. Those include package managers (e.g. the cabal.project file) or configuration files for Haskell build tools. More interestingly, there is also an informal mathematical notation associated with Haskell that is used when programmers discuss programs on a whiteboard or in academic publications. The idea of having such a mathematical notation dates back to the *Report on Algol 58* [51], which explicitly defined a “publication language” for “stating and communicating problems” using Greek letters and subscripts.

4.2.9 Dimension: expression geography

A crucial feature of a notation is the relationship between the structure of the notation and the structure of the behavior it encodes. Most importantly, do *similar expressions* in a particular notation represent *similar behavior*?¹³ Visual notations may provide a more or less direct mapping. On the one hand, similar-looking code in a block language may mean very different things. On the other hand, similar looking design of two HyperCard cards will result in similar looking cards—the mapping between the notation and the logic is much more direct.

C/C++ expression language. In textual notations, this may easily not be the case. Consider the two C conditionals:

¹² The tool for versioning and package management in Pharo can still be seen as an *internal* domain-specific language and thus as a secondary notation, but its basic structure is *shared* with other notations in the Pharo system.

¹³ See Basman’s [1] similar discussion of “density”.

- if $(x==1) \{ \dots \}$ evaluates the Boolean expression $x==1$ to determine whether x equals 1, running the code block if the condition holds.
- if $(x=1) \{ \dots \}$ assigns 1 to the variable x . In C, assignment is an expression *returning* the assigned value, so the result 1 is interpreted as true and the block of code is *always* executed.

A notation can be designed to map better to the logic behind it, for example, by requiring the user to write $1=x$. This solves the above problem as 1 is a literal rather than a variable, so it cannot be assigned to ($1=x$ is a compile error).

4.2.10 Dimension: uniformity of notations

One common concern with notations is the extent to which they are uniform. A uniform notation can express a wide range of things using just a small number of concepts. The primary example here is S-expressions from Lisp. An S-expression is either an atom or a pair of S-expressions written $(s1 . s2)$. By convention, an S-expression $(s1 . (s2 . (s3 . nil)))$ represents a list, written as $(s1 s2 s3)$. In Lisp, uniformity of notations is closely linked to uniformity of representation.¹⁴ In the idealized model of LISP 1.5, the data structures represented by an S-expression are what exists in memory. In real-world Lisp systems, the representation in memory is more complex. A programming system can also take a very different approach and fully separate the notation from the in-memory representation.

Lisp systems. In Lisp, source code is represented in memory as S-expressions, which can be manipulated by Lisp primitives. In addition, Lisp systems have robust macro processing as part of their semantics: expanding a macro revises the list structure of the code that uses the macro. Combining these makes it possible to define extensions to the system in Lisp, with syntax indistinguishable from Lisp. Moreover, it is possible to write a program that constructs another Lisp program and not only run it interpretively (using the eval function) but compile it at runtime (using the compile function) and execute it. Many domain-specific languages, as well as prototypes of new programming languages (such as Scheme), were implemented this way. Lisp the language is, in this sense, a “programmable programming language”. [16, 14]

4.2.11 References

Cognitive Dimensions of Notation [24] provide a comprehensive framework for analysing individual notations, while our focus here is on how multiple notations are related and how they are structured. It is worth noting that the Cognitive Dimensions also define *secondary notation*, but in a different sense to ours. For them, secondary notation refers to whether a notation allows including redundant information such as color or comments for readability purposes.

The importance of notations in the practice of science, more generally, has been studied by [36] as “paper tools”. These are formula-like entities which can be manipu-

¹⁴ Notations generally are closely linked to representation in that the notation may mirror the structures used for program representation. Basman et al. [3] refer to this as a distinction between “dead” notation and “live” representation forms).

Technical Dimensions of Programming Systems

lated by humans in lieu of experimentation, such as the aforementioned mathematical notation in Haskell: a “paper tool” for experimentation on a whiteboard. Programming notations are similar, but they are a way of communicating with a machine; the experimentation does not happen on paper alone.

4.2.12 Relations

- *Interaction* (Section 4.1): The feedback loops that exist in a programming system are typically associated with individual notations. Different notations may also have different feedback loops.
- *Adoptability* (Section 4.7): Notational structure can affect learnability. In particular, complementing notations may require (possibly different) users to master multiple notations. Overlapping notations may improve learnability by allowing the user to edit the program in one way (perhaps visually) and see the effect in the other notation (such as code.)
- *Errors* (Section 4.6). A process that merely records user actions in a sequence (such as text editing) will, in particular, record any *errors* the user makes and defer their handling to later use of the data, keeping the errors *latent*. A process which instead treats user actions as edits to a structure, with constraints and correctness rules, will be able to catch errors at the moment they are introduced and ensure the data coming out is error-free.

4.3 Conceptual Structure

How is meaning constructed? How are internal and external incentives balanced?

4.3.1 Dimension: conceptual integrity vs. openness

The evolution of programming systems has led away from *conceptual integrity* towards an intricate ecosystem of specialized technologies and industry standards. Any attempt to unify parts of this ecosystem into a coherent whole will create *incompatibility* with the remaining parts, which becomes a major barrier to adoption. Designers seeking adoption are pushed to focus on localized incremental improvements that stay within the boundaries established by existing practice. This creates a tension between how highly they can afford to value conceptual elegance, and how open they are to the pressures imposed by society. We will turn to both of these opposite ends—*integrity* and *openness*—in more detail.

4.3.2 Example: conceptual integrity

I will contend that Conceptual Integrity is the most important consideration in system design. It is better to have a system omit certain anomalous features and improvements, but to reflect one set of design ideas, than to have one that contains many good but independent and uncoordinated ideas. [6]

Conceptual integrity arises not (simply) from one mind or from a small number of agreeing resonant minds, but from sometimes hidden co-authors and the thing designed itself. [20]

Conceptual integrity strives to reduce complexity at the source; it employs *unified concepts* that may *compose orthogonally* to generate diversity. Perhaps the apotheosis of this approach can be found in early Smalltalk and Lisp machines, which were complete programming systems built around a single language. They incorporated capabilities commonly provided *outside* the programming language by operating systems and databases. Everything was done in one language, and so everything was represented with the datatypes of that language. Likewise the libraries and idioms of the language were applicable in all contexts. Having a *lingua franca* avoided much of the friction and impedance mismatches inherent to multi-language systems. A similar drive exists in the Python programming language, which follows the principle that “There should be one—and preferably only one—obvious way to do it” in order to promote community consensus on a single coherent style.

In addition to Smalltalk and Lisp, many programming languages focus on one kind of data structure [58]:

- In COBOL, data consists of nested records as in a business form.
- In Fortran, data consists of parallel arrays.
- In SQL, data is a set of relations with key constraints.
- In scripting languages like Python, Ruby, and Lua, much data takes the form of string-indexed hash tables.

Finally, many languages are *imperative*, staying close to the hardware model of addressable memory, lightly abstracted into primitive values and references into mutable arrays and structures. On the other hand, *functional* languages hide references and treat everything as immutable structured values. This conceptual simplification benefits certain kinds of programming, but can be counterproductive when an imperative approach is more natural, such as in external input/output.

4.3.3 Example: conceptual openness

Perl, contra Python. In contrast to Python’s outlook, Perl proclaims “There is more than one way to do it” and considers itself “the first postmodern programming language” [67]. “Perl doesn’t have any agenda at all, other than to be maximally useful to the maximal number of people. To be the duct tape of the Internet, and of everything else.” The Perl way is to accept the status quo of evolved chaos and build upon it using duct tape and ingenuity. Taken to the extreme, a programming system becomes no longer a *system*, properly speaking, but rather a *toolkit for improvising* assemblages of *found* software. Perl can be seen as championing the values of *pluralism, compatibility, or conceptual openness* over conceptual integrity. This philosophy has been called *Postmodern Programming* [45].

C++ , contra Smalltalk. Another case is that of C++ , which added to C the Object-Oriented concepts developed by Smalltalk while remaining 100% compatible with C, down to the level of ABI and performance. This strategy was enormously successful for adoption, but came with the tradeoff of enormous complexity compared to languages designed from scratch for OO, like Smalltalk, Ruby, and Java.

Worse, contra Better. Richard Gabriel first described this dilemma in his influential 1991 essay *Worse is Better* [19] analyzing the defeat of Lisp by UNIX and C. Because

Technical Dimensions of Programming Systems

UNIX and C were so easy to port to new hardware, they were “the ultimate computer viruses” despite providing only “about 50%–80% of what you want from an operating system and programming language”. Their conceptual openness meant that they adapted easily to the evolving conditions of the external world. The tradeoff was decreased conceptual integrity, such as the undefined behaviours of C, the junkyard of working directories, and the proliferation of special purpose programming languages to provide a complete development environment.

UNIX and Files. Many programming languages and systems impose structure at a “fine granularity”: that of individual variables and other data and code structures. Conversely, systems like UNIX and the Web impose fewer restrictions on how programmers represent things. UNIX insists only on a basic infrastructure of “large objects” [33], delegating all fine-grained structure to client programs. This scores many points for conceptual openness. *Files* provide a universal API for reading and writing byte streams, a low-level construct containing so many degrees of freedom that it can support a wide variety of formats and ecosystems. *Processes* similarly provide a thin abstraction over machine-level memory and processors.

Conceptual integrity is necessarily sacrificed for such openness; while “everything is a file” gestures at integrity, in the vein of Smalltalk’s “everything is an object”, exceptions proliferate. Directories are special kinds of files with special operations, hardware device files require special ioctl operations, and many commands expect files containing newline separators. Additionally, because client programs must supply their *own* structure for fine-grained data and code, they are given little in the way of mutual compatibility. As a result, they tend to evolve into competing silos of duplicated infrastructure [33, 32].

The Web. Web HTTP endpoints, meanwhile, have proven to be an even more adaptable and viral abstraction than UNIX files. They operate at a similar level of abstraction as files, but support richer content and encompass internet-wide interactions between autonomous systems. In a sense, HTTP GET and PUT have become the “subroutine calls” of an internet-scale programming system. Perhaps the most salient thing about the Web is that its usefulness came as such a surprise to everyone involved in designing or competing with it. It is likely that, by staying close to the existing practice of transferring files, the Web gained a competitive edge over more ambitious and less familiar hypertext projects like Xanadu [44].

The choice between compatibility and integrity correlates with the personality traits of *pragmatism* and *idealism*. It is pragmatic to accept the status quo of technology and make the best of it. Conversely, idealists are willing to fight convention and risk rejection in order to attain higher goals. We can wonder which came first: the design decision or the personality trait? Do Lisp and Haskell teach people to think more abstractly and coherently, or do they filter for those with a pre-existing condition? Likewise, perhaps introverted developers prefer the cloisters of Smalltalk or Lisp to the adventurous “Wild West” of the Web.

4.3.4 Dimension: composability

In short, *you can get anywhere by putting together a number of smaller steps*. There exist building blocks which span a range of useful combinations. Composability is, in

a sense, key to the notion of “programmability” and every programmable system will have some level of composability (e.g. in the scripting language.)

UNIX shell commands are a standard example of composability. The base set of primitive commands can be augmented by programming command executables in other languages. Given some primitives, one can “pipe” one’s output to another’s input (`|`), sequence (`;` or `&&`), select via conditions, and repeat with loop constructs, enabling full imperative programming. Furthermore, command compositions can be packaged into a named “script” which follows the same interface as primitive commands, and named subprograms within a script can also be defined.

In *HyperCard*, the *Authoring Environment* is *non-composable* for programming buttons: there is simply a set of predefined behaviors to choose from. Full scriptability is available only in the *Programming Environment*.

The *Haskell type system*, as well as that of other functional programming languages, exhibits high composability. New types can be defined in terms of existing ones in several ways. These include records, discriminated unions, function types and recursive constructs (e.g. to define a List as either a Nil or a combination of element plus other list.) The C programming language also has some means of composing types that are analogous in some ways, such as structs, unions, enums and indeed even function pointers. For every type, there is also a corresponding “pointer” type. It lacks, however, the recursive constructs permitted in Haskell types.

4.3.5 Dimension: convenience

In short, *you can get to X, Y or Z via one single step*. There are ready-made solutions to specific problems, not necessarily generalizable or composable. Convenience often manifests as “canonical” solutions and utilities in the form of an expansive standard library.

Composability without convenience is a set of atoms or gears; theoretically, anything one wants could be built out of them, but one must do that work. This situation has been criticized as the *Lisp Curse*[68].

Composability *with* convenience is a set of convenient specific tools *along with* enough components to construct new ones. The specific tools themselves could be transparently composed of these building blocks, but this is not essential. They save users the time and effort it would take to “roll their own” solutions to common tasks.

For example, let us turn to a convenience factor of *UNIX* shell commands, having already discussed their composability above. Observe that it would be possible, in principle, to pass all information to a program via standard input. Yet in actual practice, for convenience, there is a standard interface of *command-line arguments* instead, separate from anything the program takes through standard input. Most programming systems similarly exhibit both composability and convenience, providing templates, standard libraries, or otherwise pre-packaged solutions, which can nevertheless be used programmatically as part of larger operations.

4.3.6 Dimension: commonality

Humans can see Arrays, Strings, Dicts and Sets all have a “size”, but the software needs to be *told* that they are the “same”. Commonality like this can be factored out

Technical Dimensions of Programming Systems

into an explicit structure (a “Collection” class), analogous to database *normalization*. This way, an entity’s size can be queried without reference to its particular details: if `c` is declared to be a Collection, then one can straightforwardly access `c.size`.

Alternatively, it can be left implicit. This is less upfront work, but permits instances to *diverge*, analogous to *redundancy* in databases. For example, Arrays and Strings might end up with “length”, while Dict and Set call it “size”. This means that, to query the size of an entity, it is necessary to perform a case split according to its concrete type, solely to funnel the diverging paths back to the commonality they represent:

```
if (entity is Array or String) size := entity.length
else if (entity is Dict or Set) size := entity.size
```

4.3.7 Examples: flattening and factoring

Data structures usually have several “moving parts” that can vary independently. For example, a simple pair of “vehicle type” and “color” might have all combinations of (Car, Van, Train) and (Red, Blue). In this *factored* representation, we can programmatically change the color directly: `pair.second = Red` or `vehicle.colour = Red`.

In some contexts, such as class names, a system might only permit such multi-dimensional structure as an *exhaustive enumeration*: RedCar, BlueCar, RedVan, BlueVan, RedTrain, BlueTrain, etc. The system sees a flat list of atoms, even though a human can see the sub-structure encoded in the string. In this world, we cannot simply “change the color to Red” programmatically; we would need to case-split as follows:

```
if (type is BlueCar) type := RedCar
else if (type is BlueVan) type := RedVan
else if (type is BlueTrain) type := RedTrain
...
```

The *commonality* between RedCar, RedVan, BlueCar, and so on has been *flattened*. There is implicit structure here that remains *un-factored*, similar to how numbers can be expressed as singular expressions (16) or as factor products (2,2,2,2). *Factoring* this commonality gives us the original design, where there is a pair of values from different sets.

In *relational databases*, there is an opposition between *normalization* and *redundancy*. In order to fit multi-table data into a *flat* table structure, data needs to be duplicated into redundant copies. When data is *factored* into small tables as much as possible, such that there is only one place each piece of data “lives”, the database is in *normal form* or *normalized*. Redundancy is useful for read-only processes, because there is no need to join different tables together based on common keys. Writing, however, becomes risky; in order to modify one thing, it must be synchronized to the multiple places it is stored. This makes highly normalized databases optimized for writes over reads.

4.3.8 Remark: the end of history?

Today we live in a highly developed world of software technology. It is estimated that 41,000 person years have been invested into Linux. We describe software development technologies in terms of *stacks* of specialized tools, each of which might

capitalize over 100 person-years of development. Programming systems have become programming ecosystems: not designed, but evolved. How can we noticeably improve programming in the face of the overwhelming edifice of existing technology? There are strong incentives to focus on localized incremental improvements that don't cross the established boundaries.

The history of computing is one of cycles of evolution and revolution. Successive cycles were dominated in turn by mainframes, minicomputers, workstations, personal computers, and the Web. Each transition built a whole new technology ecosystem replacing or on top of the previous. The last revolution, the Web, was 25 years ago, with the result that many people have never experienced a disruptive platform transition. Has history stopped, or are we just stuck in a long cycle, with increasingly pent-up pressures for change? If it is the latter, then incompatible ideas now spurned may yet flourish.

4.3.9 References

- How to Design a Good API and Why it Matters [5]

4.4 Customizability

Once a program exists in the system, how can it be extended and modified?

Programming is a gradual process. We start either from nothing, or from an existing program, and gradually extend and refine it until it serves a given purpose. Programs created using different programming systems can be refined to different extents, in different ways, at different stages of their existence.

Consider three examples. First, a program in a conventional programming language like Java can be refined only by modifying its source code. However, you may be able to do so by just adding new code, such as a new interface implementation. Second, a spreadsheet can be modified at any time by modifying the formulas or data it contains. There is no separate programming phase. However, you have to modify the formulas directly in the cell—there is no way of modifying it by specifying a change in a way that is external to the cell. Third, a *self-sustaining* programming system, such as Smalltalk, does not make an explicit distinction between “programming” and “using” phases, and it can be modified and extended via itself. It gives developers the power to experiment with the system and, in principle, replace it with a better system from within.

4.4.1 Dimension: staging of customization

For systems that distinguish between different stages, such as writing source code versus running a program, customization methods may be different for each stage. In traditional programming languages, customization is done by modifying or adding source code at the programming stage, but there is no (automatically provided) way of customizing the created programs once they are running.

There are a number of interesting questions related to staging of customization. First, what is the notation used for customization? This may be the notation in which a program was initially created, but a system may also use a secondary notation for

Technical Dimensions of Programming Systems

customization (consider Emacs using Emacs Lisp). For systems with a stage distinction, an important question is whether such changes are *persistent*.

Smalltalk, Interlisp and similar. In image-based programming systems, there is generally no strict distinction between stages and so a program can be customized during execution in the same way as during development. The program image includes the programming environment. Users of a program can open this, navigate to a suitable object or a class (which serve as the *addressable extension points*) and modify that. Lisp-based systems such as *Interlisp* follow a similar model. Changes made directly to the image are persistent. The PILOT system for Lisp [65] offers an interactive way of correcting errors when a program fails during execution. Such corrections are then applied to the image and are thus persistent.

Document Object Model (DOM) and Webstrates: In the context of Web programming, there is traditionally a stage distinction between programming (writing the code and markup) and running (displaying a page). However, the DOM can be also modified by browser Developer Tools—either manually, by running scripts in a console, or by using a userscript manager such as Greasemonkey. Such changes are not persistent in the default browser state, but are made so by Webstrates [37] which synchronize the DOM between the server and the client. This makes the DOM collaborative, but not (automatically) *live* because of the complexities this implies for event handling.

4.4.2 Dimension: addressing and externalizability

Programs in all programming systems have a representation that may be exposed through notation such as source code. When customizing a program, an interesting question is whether a customization needs to be done by modifying the original representation, or whether it can be done by *adding* something alongside the original structure.

In order to support customization through addition, a programming system needs a number of characteristics introduced by Basman et al. [3, 4]. First, the system needs to support *addressing*: the ability to refer to a part of the program representation from the outside. Next, *externalizability* means that a piece of addressed state can be exhaustively transferred between the system and the outside world. Finally, *open authoring* requires that system behaviours can be *changed* by simply *adding* a new expression containing addresses—in other words, anything can be *overridden* without being *erased*. Of particular importance is how addresses are specified and what extension points in the program they can refer to. The system may offer an automatic mechanism that makes certain parts of a program addressable, or this task may be delegated to the programmer.

Cascading Style Sheets (CSS): CSS is a prime example of a system that offers open authoring with rich addressability mechanisms that are partly automatic (when referring to tag names) and partly manual (when using element IDs and class names). Given a web page, it is possible to modify (almost) any aspects of its appearance by simply *adding* additional rules to a CSS file. The Infusion project [2] offers similar customizability mechanisms, but for behaviour rather than just styling.

Object Oriented Programming and Aspect Oriented Programming: in conventional programming languages, customization is done by modifying the code itself. OOP

and AOP make it possible to do so by adding code independently of existing program code. In OOP, this requires manual definition of extension points, i.e. interfaces and abstract methods. Functionality can then be added to a system by defining a new class (although injecting the new class into existing code without modification requires some form of configuration such as a dependency injection container). AOP systems such as AspectJ [35] provides a richer addressing mechanism. In particular, it makes it possible to add functionality to the invocation of a specific method (among other options) by using the *method call pointcut*. This functionality is similar to *advising* in Pilot [65].

4.4.3 Dimension: self-sustainability

For most programming languages, programming systems, and ordinary software applications, if one wants to customize beyond a certain point, one must go beyond the facilities provided in the system itself. Most programming systems maintain a clear distinction between the *user level*, where the system is used, and *implementation level*, where the source code of the system itself resides. If the user level does not expose control over some property or feature, then one is forced to go to the implementation level. In the common case this will be a completely different language or system, with an associated learning cost. It is also likely to be lower-level—lacking expressive functions, features or abstractions of the user level—which makes for a more tedious programming experience.

It is possible, however, to carefully design systems to expose deeper aspects of their implementation *at the user level*, relaxing the formerly strict division between these levels. For example, in the research system *3-Lisp* [60], ordinarily built-in functions like the conditional *if* and error handling *catch* are implemented in 3-Lisp code at the user level.

The degree to which a system's inner workings are accessible to the user level, we call *self-sustainability*. At the maximal degree of this dimension would reside “stem cell”-like systems: those which can be progressively evolved to arbitrary behavior without having to “step outside” of the system to a lower implementation level. In a sense, any difference between these systems would be merely a difference in initial state, since any could be turned into any other.

The other end, of minimal self-sustainability, corresponds to minimal customizability: beyond the transient run-time state changes that make up the user level of any piece of software, the user cannot change anything without dropping down to the means of implementation of the system. This would resemble a traditional end-user “application” focused on a narrow domain with no means to do anything else.

The terms “self-describing” or “self-implementing” have been used for this property, but they can invite confusion: how can a thing describe itself? Instead, a system that can *sustain itself* is an easier concept to grasp. The examples that we see of high self-sustainability all tend to be *Operating System-like*. UNIX is widely established as an operating system, while Smalltalk and Lisp have been branded differently. Nevertheless, all three have shipped as the operating systems of custom hardware, and have similar responsibilities. Specifically: they support the execution of “programs”; they define an interface for accessing and modifying state; they provide standard

Technical Dimensions of Programming Systems

libraries of common functionality; they define how programs can communicate with each other; they provide a user interface.

UNIX: Self-sustainability of UNIX is owed to the combination of two factors. First, the system is implemented in binary files (via ELF¹⁵) and text files (for configuration). Second, these files are part of the user-facing filesystem, so users can replace and modify parts of the system using UNIX file interfaces.

Smalltalk and COLAs: Self-sustainability in Smalltalk is similar to UNIX, but at a finer granularity and with less emphasis on whether things reside in volatile (process) or non-volatile (file) storage. The analogous points are that (1) the system is implemented as objects with methods containing Smalltalk code, and (2) these are modifiable using the class browser and code editor. Combined Object Lambda Architectures, or COLAs [54], are a theoretical system design to improve on the self-sustainability of Smalltalk. This is achieved by generalizing the object model to support relationships beyond classes.

4.4.4 References

In addition to the examples discussed above, the proceedings of self-sustaining systems workshops [29, 28] provides numerous examples of systems and languages that are able to bootstrap, implement, modify, and maintain themselves; Gabriel's analysis of programming language revolutions [21] uses *advising* in PILOT, related Lisp mechanisms, and "mixins" in OOP to illustrate the difference between the "languages" and "systems" paradigms.

4.4.5 Relations

- *Flattening and factoring* (Section 4.3.7): related in that "customizability" is a form of creating new programs from existing ones; factoring repetitive aspects into a reusable standard component library facilitates the same thing.
- *Interaction* (Section 4.1): this determines whether there are separate stages for running and writing programs and may thus influence what kind of customization is possible.

4.5 Automation

How far does the system remove the need to spell out implementation in minute detail?

Ultimately, at the hardware level, computers are primitive calculating machines. They require a full and exact specification of the instructions to run. Ever since the 1940s, programmers have envisioned that some form of "automatic programming" will alleviate the need for tediously specifying details at this level. While this level still remains today, many aspects of the task of "programming" can and have been *automated*.

¹⁵ Executable and Linkable Format.

Automation can take a number of forms. Extracting common functionality into a library may be merely good use of *factoring* (Section 4.3.7), but to the user of the library, this may appear as automation. In high-level programming languages, many details are also omitted; those are filled in by the compiler. Finally, the program may also be executed by a more sophisticated runtime that fills in details not specified explicitly, such as when running an SQL query or using a logic programming system like Prolog.

4.5.1 Remark: notations

Even with high-level of automation, programming involves manipulating some program notation. In high-level functional or imperative programming languages, the programmer writes code that typically has clear operational meaning. When using more declarative programming like SQL, Prolog or Datalog, the meaning of a program is still unambiguous, but it is not defined operationally—there is a (more or less deterministic) inference engine that solves the problem based on the provided description. Finally, systems based on *programming by example* step even further away from having clear operational meaning—the program may be simply a collection of sample inputs and outputs, from which a (typically non-deterministic) engine infers the concrete steps of execution.

4.5.2 Dimension: degrees of automation

There are many degrees of automation in programming systems, but the basic mechanism is always the same—given a program, some logic is specified explicitly and some is left to a reusable component that can do the rest. In the case of library reuse, the reusable component is just the library. In the case of higher-level programming languages, the reusable component may include a memory allocator or a garbage collector. In case of declarative languages or programming by example, the reusable component is a general purpose inference engine.

Higher levels of automation require more complex *reusable components* than lower levels. This is a difference between *level of automation* and *factoring*—producing systems with higher level of automation requires more than simply extracting (factoring) existing code into a reusable component. Instead, it requires doing more work and introducing a higher level of indirection between the program and the reusable component.

There is also an interesting (and perhaps inevitable) trade-off. The higher the level of automation, the less explicit the operational meaning of a program. This has a wide range of implications. Smaragdakis [59] notes, for example, that this means the implementation can significantly change the performance of a program.

4.5.3 Example: programming by example

An interesting case of automation is *programming by example* [39]. In this case, the user does not provide even a declarative specification of the program behavior, but instead specifies sample inputs and outputs. A more or less sophisticated algorithm then attempts to infer the relationship between the inputs and the outputs. This may, for example, be done through program synthesis where an algorithm composes a

Technical Dimensions of Programming Systems

transformation using a (small) number of pre-defined operations. Programming by example is often very accessible and has been used in spreadsheet applications [25].

4.5.4 Example: next-level automation

Throughout history, programmers have always hoped for the next level of “automatic programming”. As observed by Parnas [49], “automatic programming has always been a euphemism for programming in a higher-level language than was then available to the programmer”.

We may speculate whether Deep Learning will enable the next step of automation. However, this would not be different in principle from existing developments. We can see any level of automation as using *artificial intelligence* methods. This is the case for declarative languages or constraint-based languages—where the inference engine implements a traditional AI method (GOFAI, i.e., Good Old Fashioned AI).

4.5.5 Relations

- *Flattening and factoring* (Section 4.3.7: One typically automates the thing at the lowest level in one’s factoring (by making the lowest level a thing that exists outside of the program—in a system or a library)

4.6 Errors

What does the system consider to be an *error*? How are they prevented and handled?

A computer system is not aware of human intentions. There will always be human mistakes that the system cannot recognize as errors. Despite this, there are many that it *can* recognize, and its design will determine *which* human mistakes can become detectable program errors. This revolves around several questions: What can cause an error? Which ones can be prevented from happening? How should the system react to errors?

Following the standard literature on errors [55], we distinguish four kinds of errors: slips, lapses, mistakes and failures. A *slip* is an error caused by transient human attention failure, such as a typo in the source code. A *lapse* is similar but caused by memory failure, such as an incorrectly remembered method name. A *mistake* is a logical error such as bad design of an algorithm. Finally, a *failure* is a system error caused by the system itself that the programmer has no control over, e.g. a hardware or a virtual machine failure.

4.6.1 Dimensions: error detection

Errors can be identified in any of the *feedback loops* that the system implements. This can be done either by a human or the system itself, depending on the nature of the feedback loop.

Consider three examples. First, in live programming systems, the programmer immediately sees the result of their code changes. Error detection is done by a human and the system can assist this by visualizing as many consequences of a code change as possible. Second, in a system with a static checking feedback loop (such as syntax checks, static type systems), potential errors are reported as the result of the analysis.

Third, errors can be detected when the developed software is run, either when it is tested by the programmer (manually or through automated testing) or when it is run by a user.

Error detection in different feedback loops is suitable for detecting different kinds of errors. Many slips and lapses can be detected by the static checking feedback loop, although this is not always the case. For example, consider a “compact” *expression geography* where small changes in code may result in large changes of behaviour. This makes it easier for slips and lapses to produce hard to detect errors. Mistakes are easier to detect through a live feedback loop, but they can also be partly detected by more advanced static checking.

4.6.2 Example: static typing

In statically typed programming languages like Haskell and Java, types are used to capture some information about the intent of the programmer. The type checker ensures code matches the lightweight specification given using types. In such systems, types and implementation serve as two descriptions of programmer’s intent that need to align; what varies is the extent to which types can capture intent and the way in which the two are constructed; that is, which of the two comes first.

4.6.3 Examples: TDD, REPL and live coding

Whereas static typing aims to detect errors without executing code, approaches based on immediate feedback typically aim to execute (a portion of) the code and let the programmer see the error immediately. This can be done in a variety of ways.

In case of *test-driven development*, tests play the role of specification (much like types) against which the implementation is checked. Such systems may provide more or less immediate feedback, depending on when tests are executed (automatically in the background, or manually). Systems equipped with a read-eval-print loop (REPL) let programmers run code on-the-fly and inspect results. For successful error detection, the results need to be easily observable: a printed output is more helpful than a hidden change of system state. Finally, in live coding systems, code is executed immediately and the programmer’s ability to recognize errors depends on the extent to which the system state is observable. In live coded music, for example, you *hear* that your code is not what you wanted, providing an easy-to-use immediate error detection mechanism.

4.6.4 Remark: eliminating latent errors

A common aim of error detection is to prevent *latent errors*, i.e. errors that occurred at some *earlier* point during execution, but only manifest themselves through an unexpected behaviour later on. For example, we might dereference the wrong memory address and store a junk value to a database; we will only find out upon accessing the database. Latent errors can be prevented differently in different feedback loops. In a live feedback loop, this can be done by visualizing effects that would normally remain hidden. When running software, latent errors can be prevented through a mechanism that detects errors as early as possible (e.g. initializing pointers to null and stopping if they are dereferenced.)

Technical Dimensions of Programming Systems

Elm and time-travel debugging. One notable mechanism for identifying latent errors is the concept of *time-travel debugging* popularized by the Elm programming language. In time-travel debugging, the programmer is able to step back through time and see what execution steps were taken prior to a certain point. This makes it possible to break execution when a latent error manifests, but then retrace the execution back to the actual source of the error.

4.6.5 Dimension: error response

When an error is detected, there are a number of typical ways in which the system can respond. The following applies to systems that provide some kind of error detection during execution.

- It may attempt to automatically recover from the error as best as possible. This may be feasible for simpler errors (slips and lapses), but also for certain mistakes (a mistake in an algorithm's concurrency logic may often be resolved by restarting the code.)
- It may proceed as if the error did not happen. This can eliminate expensive checks, but may lead to latent errors later.
- It may ask a human how to resolve the issue. This can be done interactively, by entering into a mode where the code can be corrected, or non-interactively by stopping the system.

Orthogonally to the above options, a system may also have a way to recover from latent errors by tracing back through the execution in order to find the root cause. It may also have a mechanism for undoing all actions that occurred in the meantime, e.g. through transactional processing.

Interlisp and Do What I Mean (DWIM). Interlisp's DWIM facility attempts to automatically correct slips and lapses, especially misspellings and unbalanced parentheses. When Interlisp encounters an error, such as a reference to an undefined symbol, it invokes DWIM. In this case, DWIM then searches for similarly named symbols frequently used by the current user. If it finds one, it invokes the symbol automatically, corrects the source code and notifies the user. In more complex cases where DWIM cannot correct the error automatically, it starts an interaction with the user and lets them correct it manually.

4.6.6 Relations

- *Feedback loops:* Error detection always happens as part of an individual feedback loop. The feedback loops thus determine the structure at which error detection can happen.
- *Automation:* A semi-automatic error recovery system (such as DWIM) implements a form of automation. The concept of antifragile software [43] is a more sophisticated example of error recovery through automation.
- *Expression geography:* In an expression geography where small changes in notation lead to valid but differently behaved programs, a slip or lapse is more likely to lead to an error that is difficult to detect through standard mechanisms.

4.6.7 References

The most common error handling mechanism in conventional programming languages is exception handling. The modern form of exception handling has been described in [23]; [56] documents the history and influences of Software Engineering on exception handling. The concept of *antifragile software* [43] goes further by suggesting that software could improve in response to errors. Work on Chaos Engineering [8] is a step in this direction.

[55] analyses errors in the context of human errors and develops a classification of errors that we adopt. In the context of computing, errors or *miscomputation* has been analysed from a philosophical perspective [17, 15]. Notably, attitudes and approaches to errors also differ for different programming subcultures [52].

4.7 Adoptability

How does the system facilitate or obstruct adoption by both individuals and communities?
--

We consider adoption by individuals as the dimension of *Learnability*, and adoption by communities as the dimension of *Sociability*.

4.7.1 Dimension: learnability

Mainstream software development technologies require substantial effort to learn. Systems can be made easier to learn in several ways:

- Specializing to a specific application domain.
- Specializing to simple small-scale needs.
- Leveraging the background knowledge, skills, and terminologies of specific communities.
- Supporting learning with staged levels of complexity and assistive development tools [18]. Better *Feedback Loops* can help (Section 4.1).
- Collapsing heterogeneous technology stacks into simpler unified systems. This relates to the dimensions under *Conceptual Structure* (Section 4.3).

FORTRAN was a breakthrough in programming because it specialized to scientific computing and leveraged the background knowledge of scientists about mathematical formulas. COBOL instead specialized to business data processing and embraced the business community by eschewing mathematics in favor of plain English.

LOGO was the first language explicitly designed for teaching children. Later BASIC and Pascal were designed for teaching then-standard programming concepts at the University level. BASIC and Pascal had second careers on microprocessors in the 90's. These microprocessor programming systems were notable for being complete solutions integrating everything necessary, and so became home schools for a generation of programmers. More recently languages like Racket, Pyret, and Grace have supported learning by revealing progressive levels of complexity in stages. Scratch returned to Logo's vision of teaching children with a graphical programming environment emphasizing playfulness rather than generality.

Technical Dimensions of Programming Systems

Some programming languages have consciously prioritized the programmer's experience of learning and using them. Ruby calls itself *a programmer's best friend* by focusing on simplicity and elegance. Elm targets the more specialized but still fairly broad domain of web applications while focusing on simplicity and programmer-friendliness. It forgoes capabilities that would lead to run-time crashes. It also tries hard to make error messages clear and actionable.

If we look beyond programming languages *per se*, we find programmable systems with better learnability. The best example is spreadsheets, which offer a specialized computing environment that is simpler and more intuitive. The visual metaphor of a grid leverages human perceptual skills. Moving all programming into declarative formulas and attributes greatly simplifies both creation and understanding. Research on Live Programming [26, 66] has sought to incorporate these benefits into general purpose programming, but with limited success to date.

HyperCard and Flash were both programming systems that found widespread adoption by non-experts. Like spreadsheets they had an organizing visual metaphor (cards and timelines respectively). They both made it easy for beginners to get started. Hypercard had layers of complexity intended to facilitate gradual mastery.

Smalltalk and Lisp machines were complex but unified. After overcoming the initial learning curve, their environments provided a complete solution for building entire application systems of arbitrary complexity without having to learn other technologies. Boxer [11] is notable for providing a general-purpose programming environment—albeit for small-scale applications—along with an organizing visual metaphor like that of spreadsheets.

4.7.2 Dimension: sociability

Over time, especially in the internet era, social issues have come to dominate programming. Much programming technology is now developed by open-source communities, and all programming technologies are now embedded in social media communities of their users. The nature of these communities often trumps purely technical and individual considerations [41]. Some of the specific concerns of sociability are:

- Easy integration into standard technology stacks, allowing incremental adoption, and also easy exit if needed. This dynamic was discussed in the classic essay *Worse is Better* [19] about how UNIX beat Lisp.
- Backing by large corporations or widespread industry investments that ensures economic sustainability.
- An open-source community of volunteers investing their time, which has proven to be as viable as financial support.
- Easy sharing of code via package repositories or open exchanges. Prior to the open-source era, commercial marketplaces were important, like VBX components for VisualBasic. Sharing is impeded when languages have competing dialects and libraries, like Scheme [68].
- Friendly and helpful user communities on social media, for example Stack Overflow. Such communities have to some extent replaced the traditional role of documentation.

The tenor of the online community around a programming system can be its most public attribute. Even before social media, Flash developed a vibrant community of amateurs sharing code and tips. The Elm language invested much effort in creating a welcoming community from the outset [9]. Attempts to reform older communities have introduced Codes of Conduct, but not without controversy.

On the other hand, a cloistered community that turns its back on the wider world can give its members strong feelings of belonging and purpose. Examples are Smalltalk, Racket, Clojure, and Haskell. These communities bear some resemblance to cults, with guru-like leaders, and fierce group cohesion.

The economic sustainability of a programming system can be even more important than strictly social and technical issues. Adopting a technology is a costly investment in terms of time, money, and foregone opportunities. Everyone feels safer investing in a technology backed by large corporations that are not going away, or in technologies that have such widespread adoption that they are guaranteed to persist. A vibrant and mature open-source community backing a technology also makes it safer.

Unfortunately, all of these issues of sociability create barriers to new programming systems targeting non-experts, and indeed the entire dimension of learnability. Large internet corporations have invested mainly in technologies relevant to their high-end needs. Open-source communities have only flourished around technologies for expert programmers “scratching their own itch”. While there has been a flow of venture funding into “no-code” and “low-code” programming systems, it is not clear how they can become economically and socially sustainable. By and large, the internet era has seen the ascendancy of expert programmers.

5 Conclusions

There is a renewed interest in developing new programming systems. Such systems go beyond the simple model of code written in a programming language using a more or less sophisticated text editor. They combine textual and visual notations, create programs through rich graphical interactions, and challenge accepted assumptions about program editing, execution and debugging. Despite the growing number of novel programming systems, it remains difficult to evaluate the design of programming systems and see how they improve over work done in the past. To address the issue, we proposed a framework of “technical dimensions” that captures essential characteristics of programming systems in a qualitative but rigorous way.

The framework of technical dimensions puts the vast variety of programming systems, past and present, on a common footing of commensurability. This is crucial to enable the strengths of each to be identified and, if possible, combined by designers of the next generation of programming systems. As more and more systems are assessed in the framework, a picture of the space of possibilities will gradually emerge. Some regions will be conspicuously empty, indicating unrealized possibilities that could be worth trying. In this way, a domain of “normal science” is created for the design of programming systems.

Technical Dimensions of Programming Systems

Acknowledgements (To be completed for publication.)

References

- [1] Antranig Basman. Building software is not (yet) a craft. In Luke Church, editor, *Proceedings of the 27th Annual Workshop of the Psychology of Programming Interest Group, PPIG 2016, Cambridge, UK, September 7-10, 2016*, page 32. Psychology of Programming Interest Group, 2016. URL: <http://ppig.org/library/paper/building-software-not-yet-craft>.
- [2] Antranig Basman. Infusion framework and components, 2021. URL: <https://fluidproject.org/infusion.html>.
- [3] Antranig Basman, L. Church, C. Klokmoose, and Colin B. D. Clark. Software and how it lives on: Embedding live programs in the world around them. In *PPIG*, 2016.
- [4] Antranig Basman, Clayton Lewis, and Colin Clark. The open authorial principle: supporting networks of authors in creating externalisable designs. In *Proceedings of the 2018 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, pages 29–43, 2018.
- [5] Joshua Bloch. How to design a good api and why it matters, 2007. URL: <http://www.cs.bc.edu/~muller/teaching/cs102/s06/lib/pdf/api-design>.
- [6] FP Brooks. Aristocracy, democracy and system design. In *The Mythical Man Month: Essays on Software Engineering*. Addison-Wesley, 1995.
- [7] Hasok Chang. *Inventing temperature: Measurement and scientific progress*. Oxford University Press, Oxford, 2004.
- [8] Michael Alan Chang, Bredan Tschäen, Theophilus Benson, and Laurent Vanbever. Chaos monkey: Increasing sdn reliability through systematic network destruction. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, pages 371–372, 2015.
- [9] Evan Czaplicki, 2018. URL: <https://www.youtube.com/watch?v=uGlzRt-FYto>.
- [10] A. A diSessa and H. Abelson. Boxer: A reconstructible computational medium. *Commun. ACM*, 29(9):859–868, September 1986. doi:10.1145/6592.6595.
- [11] Andrea A. diSessa. A principled design for an integrated computational environment. *Human-Computer Interaction*, 1(1):1–47, 1985. URL: https://doi.org/10.1207/s15327051hci0101_1, doi:10.1207/s15327051hci0101_1.
- [12] Jonathan Edwards. Subtext: Uncovering the simplicity of programming. *SIGPLAN Not.*, 40(10):505–518, October 2005. doi:10.1145/1103845.1094851.
- [13] Jonathan Edwards, Stephen Kell, Tomas Petricek, and Luke Church. Evaluating programming systems design. In *Proceedings of 30th Annual Workshop of Psychology of Programming Interest Group, PPIG 2019*, 2019.
- [14] Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, Shriram Krishnamurthi, Eli Barzilay, Jay McCarthy, and Sam Tobin-Hochstadt. A programmable programming language. *Commun. ACM*, 61(3):62–71, February 2018. doi:10.1145/3127323.

- [15] Luciano Floridi, Nir Fresco, and Giuseppe Primiero. On malfunctioning software. *Synthese*, 192(4):1199–1220, 2015.
- [16] John Foderaro. Lisp: Introduction. *Commun. ACM*, 34(9):27, September 1991. doi:10.1145/114669.114670.
- [17] Nir Fresco and Giuseppe Primiero. Miscomputation. *Philosophy & Technology*, 26(3):253–272, 2013.
- [18] Christopher Fry. Programming on an already full brain. *Commun. ACM*, 40(4):55–64, apr 1997. doi:10.1145/248448.248459.
- [19] Richard P. Gabriel. Worse is better, 1991. URL: <https://www.dreamsongs.com/WorselsBetter.html>.
- [20] Richard P. Gabriel. Designed as designer. In *Proceedings of the 23rd ACM SIGPLAN Conference on Object-Oriented Programming Systems Languages and Applications, OOPSLA '08*, page 617–632, New York, NY, USA, 2008. Association for Computing Machinery. doi:10.1145/1449764.1449813.
- [21] Richard P. Gabriel. The structure of a programming language revolution. In *Proceedings of the ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software, Onward! 2012*, page 195–214, New York, NY, USA, 2012. Association for Computing Machinery. doi:10.1145/2384592.2384611.
- [22] Erich Gamma, Richard Helm, Ralph E. Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, Reading, Mass, 1995.
- [23] John B. Goodenough. Exception handling: Issues and a proposed notation. *Commun. ACM*, 18(12):683–696, 1975. doi:10.1145/361227.361230.
- [24] T. R. G. Green and M. Petre. Usability analysis of visual programming environments: a ‘cognitive dimensions’ framework. *JOURNAL OF VISUAL LANGUAGES AND COMPUTING*, 7:131–174, 1996.
- [25] Sumit Gulwani, William R Harris, and Rishabh Singh. Spreadsheet data manipulation using examples. *Communications of the ACM*, 55(8):97–105, 2012.
- [26] C. Hancock and M. Resnick. *Real-time programming and the big ideas of computational literacy*. PhD thesis, Massachusetts Institute of Technology, 2003. URL: <https://dspace.mit.edu/handle/1721.1/61549>.
- [27] Brian Hempel, Justin Lubin, and Ravi Chugh. Sketch-n-sketch: Output-directed programming for svg. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology, UIST '19*, page 281–292, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3332165.3347925.
- [28] Robert Hirschfeld, Hidehiko Masuhara, and Kim Rose, editors. *Workshop on Self-Sustaining Systems, S3 2010, Tokyo, Japan, September 27-28, 2010*. ACM, 2010. doi:10.1145/1942793.
- [29] Robert Hirschfeld and Kim Rose, editors. *Self-Sustaining Systems, First Workshop, S3 2008, Potsdam, Germany, May 15-16, 2008, Revised Selected Papers*, volume 5146 of *Lecture Notes in Computer Science*. Springer, 2008. doi:10.1007/978-3-540-89275-5.

Technical Dimensions of Programming Systems

- [30] Daniel Ingalls. Design principles behind smalltalk, 1981. URL: <https://archive.org/details/byte-magazine-1981-08/page/n299/mode/2up>.
- [31] A. Kay and A. Goldberg. Personal dynamic media. *Computer*, 10(3):31–41, 1977. doi:10.1109/C-M.1977.217672.
- [32] Stephen Kell. The mythical matched modules: overcoming the tyranny of inflexible software construction. In *OOPSLA Companion*, 2009.
- [33] Stephen Kell. The operating system: Should there be one? In *Proceedings of the Seventh Workshop on Programming Languages and Operating Systems*, PLOS '13, New York, NY, USA, 2013. Association for Computing Machinery. doi:10.1145/2525528.2525534.
- [34] Stephen Kell. Some were meant for c: The endurance of an unmanageable language. In *Proceedings of the 2017 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, Onward! 2017, page 229–245, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3133850.3133867.
- [35] Gregor Kiczales, Erik Hilsdale, Jim Hugunin, Mik Kersten, Jeffrey Palm, and William G. Griswold. An overview of aspectj. In Jørgen Lindskov Knudsen, editor, *ECOOP 2001 - Object-Oriented Programming, 15th European Conference, Budapest, Hungary, June 18-22, 2001, Proceedings*, volume 2072 of *Lecture Notes in Computer Science*, pages 327–353. Springer, 2001. doi:10.1007/3-540-45337-7_18.
- [36] Ursula Klein. *Experiments, Models, Paper Tools: Cultures of Organic Chemistry in the Nineteenth Century*. Stanford University Press, Stanford, CA, 2003. URL: <http://www.sup.org/books/title/?id=1917>.
- [37] Clemens N. Klokmoose, James R. Eagan, Siemen Baader, Wendy Mackay, and Michel Beaudouin-Lafon. Webstrates: Shareable dynamic media. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, UIST '15, page 280–290, New York, NY, USA, 2015. Association for Computing Machinery. doi:10.1145/2807442.2807446.
- [38] Donald Ervin Knuth. Literate programming. *The computer journal*, 27(2):97–111, 1984.
- [39] H. Lieberman. *Your Wish Is My Command: Programming by Example*. Morgan Kaufmann, 2001.
- [40] John McCarthy. *LISP 1.5 Programmer's Manual*. The MIT Press, 1962.
- [41] Leo A. Meyerovich and Ariel S. Rabkin. Socio-plt: Principles for programming language adoption. In *Proceedings of the ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, Onward! 2012, page 39–54, New York, NY, USA, 2012. Association for Computing Machinery. doi:10.1145/2384592.2384597.
- [42] Stephen L. Michel. *Hypercard: The Complete Reference*. Osborne McGraw-Hill, Berkeley, 1989.
- [43] Martin Monperrus. Principles of antifragile software. In *Companion to the First International Conference on the Art, Science and Engineering of Programming*,

- Programming '17, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3079368.3079412.
- [44] T. H. Nelson. Complex information processing: A file structure for the complex, the changing and the indeterminate. In *Proceedings of the 1965 20th National Conference*, ACM '65, page 84–100, New York, NY, USA, 1965. Association for Computing Machinery. doi:10.1145/800197.806036.
- [45] James Noble and Robert Biddle. Notes on notes on postmodern programming. *SIGPLAN Not.*, 39(12):40–56, dec 2004. doi:10.1145/1052883.1052890.
- [46] Donald A. Norman. *The Design of Everyday Things*. Basic Books, Inc., USA, 2002.
- [47] Dan R. Olsen. Evaluating user interface systems research. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology*, UIST '07, page 251–258, New York, NY, USA, 2007. Association for Computing Machinery. doi:10.1145/1294211.1294256.
- [48] Cyrus Omar, Ian Voysey, Michael Hilton, Joshua Sunshine, Claire Le Goues, Jonathan Aldrich, and Matthew A. Hammer. Toward Semantic Foundations for Program Editors. In Benjamin S. Lerner, Rastislav Bodík, and Shriram Krishnamurthi, editors, *2nd Summit on Advances in Programming Languages (SNAPL 2017)*, volume 71 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 11:1–11:12, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2017/7127>, doi:10.4230/LIPIcs.SNAPL.2017.11.
- [49] David Lorge Parnas. Software aspects of strategic defense systems. 1985. URL: <http://web.stanford.edu/class/cs99r/readings/parnas1.pdf>.
- [50] Richard Pawson. *Naked Objects*. Phd thesis, Trinity College, University of Dublin, 2004.
- [51] A. J. Perlis and K. Samelson. Preliminary report: International algebraic language. *Commun. ACM*, 1(12):8–22, December 1958. doi:10.1145/377924.594925.
- [52] Tomas Petricek. Miscomputation in software: Learning to live with errors. *Art Sci. Eng. Program.*, 1(2):14, 2017. doi:10.22152/programming-journal.org/2017/1/14.
- [53] Tomas Petricek and Joel Jakobovic. Complementary science of interactive programming systems. In *History and Philosophy of Computing*, 2021.
- [54] Ian Piumarta. Accessible language-based environments of recursive theories, 2006. URL: http://www.vpri.org/pdf/rn2006001a_colaswp.pdf.
- [55] James Reason. *Human error*. Cambridge university press, 1990.
- [56] Barbara G. Ryder, Mary Lou Soffa, and Margaret Burnett. The impact of software engineering research on modern programming languages. *ACM Trans. Softw. Eng. Methodol.*, 14(4):431–477, oct 2005. doi:10.1145/1101815.1101818.
- [57] Shneiderman. Direct manipulation: A step beyond programming languages. *Computer*, 16(8):57–69, 1983. doi:10.1109/MC.1983.1654471.
- [58] Kragen Javier Sitaker. The memory models that underlie programming languages, 2016. URL: <http://canonical.org/~kragen/memory-models/>.
- [59] Yannis Smaragdakis. Next-paradigm programming languages: What will they look like and what changes will they bring? In *Proceedings of the 2019 ACM*

Technical Dimensions of Programming Systems

- SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, Onward! 2019, page 187–197, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3359591.3359739.
- [60] Brian Cantwell Smith. *Procedural Reflection in Programming Languages*. PhD thesis, Massachusetts Institute of Technology, 1982. URL: <https://dspace.mit.edu/handle/1721.1/15961>.
- [61] D. C. Smith. Pygmalion: a creative programming environment. 1975.
- [62] Guy L. Steele and Richard P. Gabriel. The evolution of lisp. In *The Second ACM SIGPLAN Conference on History of Programming Languages*, HOPL-II, page 231–270, New York, NY, USA, 1993. Association for Computing Machinery. doi: 10.1145/154766.155373.
- [63] Steven L. Tanimoto. A perspective on the evolution of live programming. In *Proceedings of the 1st International Workshop on Live Programming*, LIVE '13, page 31–34. IEEE Press, 2013.
- [64] Philip Tchernavskij. *Designing and Programming Malleable Software*. PhD thesis, Université Paris-Saclay, École doctorale n°580 Sciences et Technologies de l'Information et de la Communication (STIC), 2019.
- [65] Warren Teitelman. *PILOT: A Step Toward Man-Computer Symbiosis*. PhD thesis, 1966.
- [66] Bret Victor. Learnable programming, 2012. URL: <http://worrydream.com/#!/LearnableProgramming>.
- [67] Larry Wall. Perl, the first postmodern computer language, 1999. URL: <http://www.wall.org/~larry/pm.html>.
- [68] Rudolf Winestock. The lisp curse, 2011. URL: http://www.winestockwebdesign.com/Essays/Lisp_Curse.html.