

Why Undergraduates Should Learn the Principles of Programming Languages

ACM SIGPLAN Education Board

Stephen N. Freund (Williams College), Kim Bruce, Chair (Pomona College), Kathi Fisler (WPI),
Dan Grossman (University of Washington),
Matthew Hertz (Canisius College), Gary T. Leavens (University of Central Florida),
Andrew Myers (Cornell University), Larry Snyder (University of Washington)

January 6, 2010

Abstract

Undergraduate students obtain important knowledge and skills by studying the pragmatics of programming in multiple languages and the principles underlying programming language design and implementation. These topics strengthen students' grasp of the power of computation, help students choose the most appropriate programming model and language for a given problem, and improve their design skills. Understanding programming languages thus helps students in ways vital to many career paths and interests.

This white paper is based on contributed articles, discussions, and presentations from the 2008 SIGPLAN Programming Language Curriculum Workshop [3, 4].

Programming languages are the medium through which humans describe computations. More specifically, we use the model provided by a programming language to discuss concepts, formulate algorithms, and reason about problem solutions. Programming languages often define models tailored to thinking about and solving problems in intended application areas. For example, the C language provides a model close to a computer's underlying hardware, a spreadsheet language (such as EXCEL with Visual Basic for Applications) provides a model of cells and constraints for solving financial problems, and so on.

The languages and models used in practice change continuously. Advances in our field and the broadening uses of technology continue to drive many exciting, and sometimes dramatic, changes in how we model and express computation. The rise of the Internet and web, for example, fundamentally transformed the way many types of systems are designed, implemented, and deployed. We are in the midst of another transformation right now: the rapid development and adoption of multicore and distributed platforms is again fundamentally changing how we think about programming.

At its core, the study of programming languages examines the principles and limitations of computational (or programming) models, how to effectively design and use languages based on these models, and how to compare their relative strengths and

weaknesses in particular contexts. Undergraduate students benefit from studying this material in substantial ways, regardless of their future career paths and interests. It strengthens “a student’s understanding of computation itself, its power and its limitations” [23]. Moreover, the knowledge and skills acquired enable students to critically compare and choose the most appropriate way to describe particular programs, and to adopt and develop new models and languages reflecting advances in our field and world around us. An education providing these abilities will best prepare students to effectively innovate in all areas of computer science and engineering. We elaborate on the most salient benefits of studying programming languages below.

1. *Learning New Computational Models and Speeding Language Learning.*

The programming languages and models commonly used in practice change constantly. Witness the growth in the use of object-oriented programming in the last 20 years, as developers recognized the advantages of object-oriented techniques for building large, extensible systems, in particular for user interfaces. Similarly, most widely-used languages now manage memory via garbage collection now that processor and memory performance, as well as improved collection techniques, have made it feasible for large systems. More recently, there has been a rapid rise in light-weight scripting languages (such as Ruby) to support the new application domain of web programming, and we are in the middle of even greater changes rooted in the advent of multicore computer architectures.

Students will not use a single model or a single set of languages for their entire careers, and they will frequently need to learn new languages when they change jobs, start new projects, or begin working in new areas. The most challenging aspect of using a new language is understanding how to describe data and algorithms in a way that matches the strengths of the language’s underlying model.

To illustrate this point, consider the fundamental difficulties of exploiting multiprocessor and multicore computer architectures, which are becoming the most promising way to achieve further computer performance improvements. Significant gains will be realized only if developers can successfully leverage the inherent concurrency in these models [24, 25]. This necessity has recently begun to drive both existing and new languages to include communication and concurrency features uncommon in most prior production languages: data-parallelism with roots in languages for functional programming and high-performance computing, transactional programming with roots in database languages, and process-based and actor-based programming previously seen primarily in niche parallel programming languages.

Emerging and future languages will embody these notions of concurrency. Those who learn just the syntax of such languages but proceed to program in a style suited for older models are doomed to be ineffective. Only a thorough understanding of the concurrency models provided by new languages will enable programmers to write robust, efficient programs. For example, X10 is a new object-oriented language for concurrent and distributed programming [22]. The language has “Java-like” syntax, but if programmers write “Java-like” programs without understanding the X10 computation model, they will fail to effectively use X10’s asynchronous computation mechanism for improving performance, its notion of “places” for simplifying

the design of distributed algorithms, and its notion of atomicity and thread communication primitives for avoiding deadlocks and other subtle, but common, errors.

The best preparation for quickly learning and effectively using new languages is understanding the fundamentals underlying all programming languages and to have some prior experience with a variety of computational models. Such knowledge will endure longer than today's "hot" languages, which will undoubtedly become obsolete and give way to new languages in the future. In addition, this knowledge will enable students to quickly look beyond an unfamiliar language's surface-level details (such as syntax) and grasp the underlying computational model's design principles.

Programming languages often evolve to include successful features from other languages, and having a solid foundation in this area also enables students to readily recognize and take advantage of changes in languages they currently use. For example, functional programming techniques offer clean, robust ways to express specific types of computation, such as manipulation of XML data from web pages, or exploration of algorithms in computer graphics. These techniques have become so widely adopted that many languages (including recent revisions to C# and Java) now directly support them via features such as anonymous functions, iterators, and generic polymorphism.

2. *Choosing the Right Language.*

The availability of so many languages and models means that students will need to make educated choices about which to use for specific tasks. Even individual systems are now rarely built entirely in one language. Instead, they are the composition of various components, each written in a language chosen for its strengths in that component's particular problem domain. For example, a web application may include database queries written in SQL, server application logic written in Java, data transformers written in XSLT, and client-side code written in JavaScript.

The choice of programming language can dramatically influence how one thinks about the design and structure of computation, and while it may be *possible* to solve a problem in any reasonable language, some problems inherently lend themselves to specific ways of thinking and programming. For example, Twitter recently switched its server infrastructure from Ruby to Scala because that language better matched their needs for long running threads, high performance under heavy loads, and more robust code via compile-time type checking [26, 1]. Scala also allows one to write parts of a system using functional programming techniques, which is attractive because many data transformations performed by a server like Twitter may be most easily written in a language expressing computation as composable functions applied to streams of data.

Other companies have enjoyed similar benefits from specific language choices. The Wall Street firm Jane Street Capital attributes a major part their success to adopting the language O'Caml for their on-line trading, research, and management systems [14]. That language's module system helped them to avoid error-prone code duplication practices endemic in previous systems built with that domain's more traditional languages (eg, C++, Java, or Excel with Visual Basic), and it led to code

that was much more readable and intuitive to discuss with business people during their stringent code review procedures.

Paul Graham also notes that his company's use of the Lisp language was instrumental in the success of their online store front application, which eventually became "Yahoo Store" [10]. That language enabled them to develop and deploy new functionality more rapidly than their several dozen competitors, who were primarily using C++ and CGI scripts. Elements of the Lisp model absent in those other languages, such as meta-programming primitives enabling programs to create, modify, and execute new pieces of code, also made implementing complex features much easier.

In contrast to these examples, choosing an ill-suited model can make devising and implementing a program far more difficult, complex, and error-prone. To avoid these pitfalls, students must have the intellectual framework and skills to critically relate models to languages and determine which choices can best solve the problem at hand.

3. *Learning Widely-Applicable Design and Implementation Techniques.*

The benefits of studying programming languages extend far beyond learning new languages and making informed choices. Recently, Dean and Ghemawat recognized that the best way to support Google's need to process huge data sets on large distributed clusters was to create MapReduce, a system "inspired by the *map* and *reduce* primitives present in Lisp and many other functional languages" [6]. Leveraging these aspects of functional programming models enabled them to create a platform well-suited for implementing many algorithms that can be automatically parallelized and distributed efficiently [8, 19].

MapReduce exemplifies how language principles transfer to many other situations ranging from large-scale system architectures to API designs to configuration mechanisms, regardless of which features exist in the language ultimately chosen for implementation. As another example, consider the current state of language support for the concurrency models described above. Each model provides an abstraction well-suited for describing specific forms of parallelism or concurrency. Understanding these models provides valuable insights into how to design particular structures and components, such as a light-weight transaction mechanism or robust communication channel, irrespective of the language used for implementing them.

In addition, web browsers, printer drivers, PDF renderers, scripted robot control systems, spreadsheets, video and audio players, and many other programs all share a common structure: they take complex input data and perform symbolic computation in a way similar to how a compiler or interpreter manipulates the source code for a program [2]. Virtually every student will at some point work on a system like these, but implementing any form of symbolic computation can be quite subtle. Students must understand the underlying theory and design principles (such as recursions that follow the source language grammar, name binding and resolution rules, etc.) to produce working and maintainable artifacts. Studying programming languages prepares students with that knowledge.

Even seemingly isolated topics in programming languages provide valuable insight and principles for systems design. Web browsers, cell phones, and an increasing number of other devices execute untrusted, and potentially malicious, code. Platforms for these devices, such as the Java Virtual Machine and Microsoft's .NET framework, often enforce access control via mechanisms based on how programming language implementations manage function calls and map variable names to storage locations [9, 13]. Similar principles underly designs of other systems that involve named resources, including operating systems that must provide numerous mechanisms for naming files, processes, synchronization devices, other computers, etc., and distributed systems [18]. Also, abstract data types, objects, and modules underly the encapsulation and abstraction principles crucial to many software engineering methodologies [21].

A common issue in web applications further demonstrates the value of understanding programming language mechanics. Any such application that must warn the user not to use the "Back" button has fundamentally flawed interactions between its underlying control flow and its management of variables for session state. Programming language principles (continuations and coroutines for control flow, and stores and environments for variable management) shed light on the subtleties of these interactions in a way that both illustrates commonly-encountered problems and how to avoid them via alternative programming models.

4. *Creating New Domain Specific Languages or Virtual Machines.*

Few students will ever design a general-purpose programming language during their careers. However, many will design domain-specific APIs, languages, or virtual machines. Such systems provide a computational model for thinking about data and algorithmic structures specific to problems in one particular context. For example, the MLFi language provides a model for describing the pricing and terms of financial contracts and language primitives for computing their valuations [7, 11]. By presenting a model centered around the specific topic of contracts, the designers created a more intuitive framework for solving problems related to financial contracts than a general-purpose language, and it has been used for that purpose quite successfully.

Computing is replete with many other heavily used domain-specific languages: Mathematica [20] and MATLAB [12] for manipulating mathematical formulas, Verilog and VHDL for describing computer hardware circuit designs, Cg [15] and others for writing rendering algorithms that run directly on graphics hardware, L^AT_EX for typesetting documents, etc. These languages all exploit properties of their intended domains to facilitate writing specific types of algorithms. For example, Cg provides direct language support for graphics concepts, such as vertices and textures, as well as operations that can execute efficiently on the highly data-parallel processing units present on graphics cards.

Designers of domain-specific languages must always address the same basic issues: How expressive must the language be? What abstract model does it provide? How will it support user-defined naming and abstraction? How will programs communicate with the rest of the computing environment? Will any specific features interact

in undesirable ways? Similar issues affect the design of API layers in systems work, such as a virtual machine to encapsulate a hardware interface and enhance portability.

Lack of knowledge of programming language fundamentals can lead to languages that are difficult to understand and use or that require later repair. For example, dynamically scoped function texts (as opposed to lexically-scoped closures for functions) make higher-order abstractions unusable in many cases, and leads to problems in type checking and optimization; this problem has had to be fixed in LISP, and Smalltalk. A solid programming languages foundation enables students to effectively recognize both when designing a new language is appropriate and how to avoid these problems.

Understanding programming language principles and models often provides the insights leading to new innovations as well. Features of MLFi, for example, were inspired by Haskell and other languages that, while not currently pervasive, are often examined in programming languages courses. Type checking and event handling models provided key insights into the development of Hancock [5], a language used successfully by AT&T to write statistical analyses for identifying patterns in huge streams of call records, such as patterns indicating fraud. Sawzall [17], Dryad [27], and Pig [16] leverage the same language principles as MapReduce to model computations for distribution across large networks of computers at high levels of abstraction.

5. *Developing Fundamental Concepts and Problem Solving Skills.*

More broadly, programming languages embody many concepts central to all of computer science, including abstraction, generalization and automation, computability, and resource management. Studying programming languages enables one to examine these topics in a precisely defined, accessible domain, and the lessons learned from programming languages thus provide immediate insight into all aspects of our discipline.

By teaching multiple ways to express computation, students also learn the importance of critically comparing and balancing multiple perspectives, and that changing perspectives can have a huge impact on one's experience and success. In fact, recognizing and incorporating diverse perspectives and thought processes can dramatically improve communication and problem solving skills for situations well outside the realm of just computer science.

Summary. The study of programming languages is an invaluable component of an undergraduate education. The knowledge and skills imparted by the study of languages and their underlying models help prepare students for successful and productive careers in computing. It teaches students to effectively use current and future languages and programming techniques, to apply sound design principles when building new languages or systems, and to think critically about many central themes in computer science. Moreover, exposing students to more than just a single programming language mono-culture equips students with skills for approaching problems from multiple, different perspectives. This invariably leads to simpler and more elegant solutions.

2008 SIGPLAN Programming Language Curriculum Workshop Contributors. *Steering Committee Co-Chairs:* Kathleen Fisher (AT&T Research), Chandra Krintz (UC Santa Barbara) *Steering Committee Members:* Eric Allen (Sun Microsystems), Ras Bodik (UC Berkeley), Kim Bruce (Pomona College), Matthias Felleisen (Northeastern Univ.), Stephen Freund (Williams College), Robert Harper (CMU), Michael Hind (IBM Research), Jim Larus (Microsoft Research), Doug Lea (SUNY Oswego), Greg Morrisett (Harvard Univ.), Lori Pollock (Univ. of Delaware), Stuart Reges (Univ. of Washington), Martin Rinard (MIT), Olin Shivers (Northeastern Univ.). *Participants:* Mark Bailey (Hamilton College), William Cook (UT Austin), Kathi Fisler (WPI), Daniel Friedman (Indiana University), John Hughes (Chalmers), Shriram Krishnamurthi (Brown), Gary T. Leavens (University of Central Florida), John Reynolds (CMU), Peter Sestoft (ITU), Lynn Andrea Stein (Olin College of Engineering), Mark Sheldon (Wellesley College), Larry Snyder (University of Washington), Franklyn Turbak (Wellesley College), Mitchell Wand (Northeastern University).

References

- [1] April 2009. <http://lambda-the-ultimate.org/node/3261>.
- [2] Eric Allen. Some things that computer science majors should know. *SIGPLAN Not.*, 43(11):32–35, 2008.
- [3] Eric Allen, Mark W. Bailey, Rastislav Bodík, Kim B. Bruce, Kathleen Fisher, Stephen N. Freund, Robert Harper, Chandra Krintz, Shriram Krishnamurthi, James R. Larus, Doug Lea, Gary T. Leavens, Lori L. Pollock, Stuart Reges, Martin C. Rinard, Mark Sheldon, Franklyn A. Turbak, and Mitchell Wand. SIGPLAN programming language curriculum workshop: Discussion summaries and recommendations. *SIGPLAN Notices*, 43(11):6–29, November 2008.
- [4] Mark Bailey, Kim Bruce, Kathleen Fisher, Robert Harper, and Stuart Reges. Report of the 2008 SIGPLAN Programming Languages Curriculum Workshop: Preliminary Report. In *Proceedings of the ACM Technical Symposium on Computer Science Education*, pages 132–133, 2009.
- [5] Corinna Cortes, Kathleen Fisher, Daryl Pregibon, Anne Rogers, and Frederick Smith. Hancock: A language for analyzing transactional data streams. *ACM Trans. Program. Lang. Syst.*, 26(2):301–338, 2004.
- [6] Jeffrey Dean and Sanjay Ghemawat. MapReduce: simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, 2008.
- [7] Jean-Marc Eber. The financial crisis, a lack of contract specification tools: What can finance learn from programming language design? In *European Symposium on Programming*, pages 205–206, 2009.
- [8] The Apache Software Foundation. Hadoop: Open source implementation of MapReduce. <http://hadoop.apache.org>, 2009.

- [9] Li Gong, Gary Ellison, and Mary Dageforde. *Inside Java 2 Platform Security: Architecture, API Design, and Implementation (2nd Edition)*. Prentice Hall, 2003.
- [10] Paul Graham. *Beating the Averages*, pages 169–180. O’Reilly, Cambridge, 2003.
- [11] Simon L. Peyton Jones, Jean-Marc Eber, and Julian Seward. Composing contracts: an adventure in financial engineering, functional pearl. In *International Conference on Functional Programming*, pages 280–292, 2000.
- [12] The Mathworks. Matlab. <http://www.mathworks.com/>, 2009.
- [13] Microsoft. Microsoft .NET framework. <http://www.microsoft.com/.NET/>, 2009.
- [14] Yaron Minsky and Stephen Weeks. Caml trading - experiences with functional programming on wall street. *J. Funct. Program.*, 18(4):553–564, 2008.
- [15] NVIDIA. Cg — the language for high-performance realtime graphics. http://developer.nvidia.com/page/cg_main.html, 2009.
- [16] Christopher Olston, Benjamin Reed, Utkarsh Srivastava, Ravi Kumar, and Andrew Tomkins. Pig Latin: a not-so-foreign language for data processing. In *SIGMOD International Conference on Management of Data*, pages 1099–1110, 2008.
- [17] Rob Pike, Sean Dorward, Robert Griesemer, and Sean Quinlan. Interpreting the data: Parallel analysis with Sawzall. *Scientific Programming*, 13(4):277–298, 2005.
- [18] Rob Pike, David L. Presotto, Sean Dorward, Bob Flandrena, Ken Thompson, Howard Trickey, and Phil Winterbottom. Plan 9 from Bell Labs. *Computing Systems*, 8(2):221–254, 1995.
- [19] Colby Ranger, Ramanan Raghuraman, Arun Penmetsa, Gary R. Bradski, and Christos Kozyrakis. Evaluating mapreduce for multi-core and multiprocessor systems. In *International Conference on High-Performance Computer Architecture*, pages 13–24, 2007.
- [20] Wolfram Research. Mathematica. <http://www.wolfram.com/>, 2009.
- [21] 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, 2005.
- [22] Vijay Saraswat and Nathaniel Nystrom. Report on the experimental language X10. <http://dist.codehaus.org/x10/>, 2009.
- [23] Olin Shivers. Why teach programming languages. *SIGPLAN Notices*, 43(11):130–132, 2008.

- [24] H. Sutter. The free lunch is over: a fundamental turn toward concurrency in software. *Dr. Dobbs's Journal*, 30(3), 2005.
- [25] Herb Sutter and James Larus. Software and the concurrency revolution. *Queue*, 3(7):54–62, 2005.
- [26] Bill Venners. Twitter on Scala: A conversation with Steve Jenson, Alex Payne, and Robey Pointer. http://www.artima.com/scalazine/articles/twitter_on_scala.html, 2009.
- [27] Yuan Yu, Michael Isard, Dennis Fetterly, Mihai Budiu, Úlfar Erlingsson, Pradeep Kumar Gunda, and Jon Currey. DryadLINQ: A system for general-purpose distributed data-parallel computing using a high-level language. In *USENIX Symposium on Operating Systems Design and Implementation*, pages 1–14, 2008.