Modelling and Extracting the Build-Time Software Architecture View

Qiang Tu, Michael W. Godfrey, and Xinyi Dong
Software Architecture Group (SWAG)
School of Computer Science, University of Waterloo
email: {qtu,migod,xdong}@swag.uwaterloo.ca

Abstract

Large software systems often have complex subparts that engage in subtle relationships with the underlying technologies from which they are built. Consequently, many such systems have complex and interesting architectural properties that can only be understood in the context of the various phases of system construction. We consider that modelling and extracting build-time architectural views of such systems are key to the software comprehension process. In this paper, we explore the notion of build-time architectural views and their relationship to well-known taxonomies of software architecture views. We present a simple notation for modelling build-time views, we present three detailed examples of systems with interesting build-time view, and we introduce a related architectural style — the "code robot". Finally, we discuss some of our recent work in the design of a toolkit for extracting and modelling build-time architectural views.

1   Introduction

The grand goal of software architecture modelling is to extract and usefully represent the most important high-level design information about a software system. Previous research and practice within the field of software architecture has affirmed the need to consider software systems from several distinct points of view. The well known architectural taxonomies of Kruchten [13] and Hofmeister *et al.* [9] suggests that four or five points of view may be sufficient: the logical view (*i.e.,* the domain object model), the (static) code view, the process/concurrency view, the deployment/execution view, plus scenarios and use-cases. However, we have found that some classes of software systems exhibit interesting and complex *build-time* properties that are not explicitly addressed by previous models.

In this paper, we explore the idea of build-time architectural views. We explain what they are, and how they fit into traditional models of software architecture, and outline how they can be extracted and modelled. We also present three detailed examples of software systems with interesting build-time architectural views.

The structure of the rest of this paper is as follows: Section 2 explores the need for the explicit modelling of build-time architectural properties: it summarizes previous definitions of software architecture and taxonomies of architectural views; it introduces the idea of build-time architectural views; it outlines some of the common activities performed at build-time and the reasons behind them; and it discusses the relationship between build-time views and configuration and build management. Section 3 introduces a visual notation for modelling build-time views, and presents three detailed examples of systems that have interesting build-time architectural views: the GCC compiler suite, the Perl language interpreter/run-time system, and a simple use of the Java Native Interface (JNI). Section 3 also introduces the "code robot" architectural style, which is common to systems that have interesting build-time views. Section 4 describes some preliminary work on a toolkit for extracting and modelling build-time views. Finally, Section 5 summarizes the work presented here, and suggests some future research directions.

2   Background: Why the build-time view is important

In this section, we explore the motivation for studying build-time architectural views. We first describe the kinds of software comprehension problems that explicitly modelling build-time properties can help to ameliorate. We then discuss software architecture and architectural views, and discuss how build-time views fit in.

2.1   The build/comprehend pipeline

Large software systems often have complex subparts whose realization involves complex and often subtle relationships with the various technologies employed in designing and implementing the system [10]. For example, a linker has a defined search order for resolving missing symbols; a programmer may take advantage of this knowledge to manage the configuration of software built for different deployment environments, or to select the correct routine from a collection of

libraries in the case of name collisions. A more complex example occurs during the building of the GCC compiler suite, when the build process first probes the target environment, and then generates core algorithms and data structures that are specialized for that environment.

Sometimes, these phenomena are relatively easy to understand, and knowledge about them can simply be recorded as annotations on appropriate software artifacts, as in the case of the linker example. More complex phenomena, such as the generation of specialized source code, are much harder to document or model without a great deal of intellectual effort; in such cases, it is common for detailed information about these properties to exist only in the minds of the system experts.

As long as an understanding of these relationships and their effects remains in the minds of the developers, it is likely that these software systems can be successfully maintained. However, it is also possible (and likely, judging by our anecdotal experiences) that such a design decision will not be documented anywhere in the source code or other development artifacts; indeed, in our own work uncovering the build-time architecture of GCC, we found very useful little information in the various official documents and web sites. Furthermore, personnel turnover, the passage of time, "off-site" system building and development (such occurs as in open source systems) can all contribute to a loss of knowledge about these aspects of the architecture of these systems. In turn, the systems become brittle and increasingly difficult and expensive to maintain.

So, therefore, it is our argument that if a software system with complex build-time properties is to be successfully maintained, then the software comprehension process must provide an intellectual means for modelling this information, as well as semi-automated tools for extracting it from the source artifacts and the build processes themselves. We now explore these ideas in more detail.

2.2 Software architecture and architectural views

Although the idea of considering software at the "architectural" level has been practised since the very first large software systems were built, as an area of formal study software architecture is still in its early stages. Several definitions of software architecture have been proposed [16, 18, 2], various architectural styles have been documented [18], and different taxonomies of architectural views have been suggested [13, 9]. Bass *et al.* define software architecture as "the design and implementation of the high-level structure or structures of the system; it comprises software components, the externally visible properties of these components, and relationships among them". Perry and Wolf presented their definition of software architecture as a mathematical model:

Software architecture = {elements, form, rationale}

Others have continued to refine and adapt the definition. For example, Boehm uses the term "connections" instead of form, and supplements constraints to the rationale. Shaw and Garlan define software architecture as components, connectors, and configurations [18]. For our purposes, we will take a hybrid view; we consider that the *software architecture* of a software system consists of a set of inter-related high-level design (structural) views; each view consists of a set of components and connectors, plus annotations. We now discuss the idea of software architecture views in more detail.

2.2.1 Architectural views

Large, long-lived software systems typically accrue design complexity as they age; it is generally agreed that to be comprehended, such systems need to be considered from different points of view. Different stakeholders (*e.g.,* users, business analysts, programmers, testers, and maintainers) typically care about only specific aspects of the software system, and have widely varying mental models based on how they experience the system. To separate the different areas of concerns, and to reflect the dynamic nature of software architecture, the notion of "views" of software architecture is introduced.[1]

The first and best known taxonomy of software architecture views was suggested by Kruchten; he proposed a "4+1" view model to describe the software architecture of a system from multiple perspectives, to separate concerns from various stakeholders of the architecture, and to differentiate the functional and non-functional requirements [13]:

- The *logical view* is the analyst's abstract object model, which captures the abstraction, encapsulation, and inheritance in the problem domain.

- The *process view* models concurrency and synchronization aspects of the software system.

- The *physical view* maps the software elements onto the hardware execution environment; this view also capture the distributed aspects of some systems.

- The *development view* models the the static organization of the software in its development environment. It mainly focuses on the source code structure.

- *Scenarios* and use-cases show how all four views work together to satisfy user requirements.

Hofmeister *et al.* have defined a similar taxonomy, which we shall refer to as the "four views" model [9]:

- The *conceptual architecture view* is analogous to the "4+1" logical view.

---

[1]Buschmann *et al.* define a view as "a partial aspect of a software architecture that shows specific properties of a software system". [4].

- The *module architecture view* shows the high-level design of the system. Its purpose is twofold: it describes how to map the conceptual view onto the high-level design artifacts, and it describes the layering structure of subsystems and specifies the interfaces between subsystem layers and low-level design artifacts (*e.g.,* source code modules).

- The *code architecture view* is the low-level design view that describes the interfaces and interdependencies between source code, files, directories, libraries, *etc.*

- The *execution architecture view* describes the structure of a system in terms of its run-time platform elements, such as tasks, processes, threads, and address spaces. This view captures the distributed and concurrent aspects of the system, as well as other non-functional requirements such as resource sharing, scheduling policies and load balancing.

Roughly speaking, we can consider that the above models concern three kinds of views: requirements views (*logical*, *process*, *scenarios*, and *conceptual*), development views (*conceptual*, *module*, and *code*), and deployment views (*physical*, *execution*). However, we have noticed that many software systems exhibit interesting structural and behavioural properties that are apparent only at system configuration and build-time, and that these properties are not explicitly considered by either the "4+1" or "four views" models. Therefore, we now examine the idea of *build-time architectural views* in more detail.

2.3 Build-time architectural views

Once a software system has been designed and implemented, it must be configured, compiled, and linked for a particular environment before it can be deployed. For a small software system written for a single platform, the `make` utility and a single `Makefile` is often sufficient to manage system building. However, for systems that are large and complex, that run on a variety of platforms, and that support multiple functional configurations, build management is non-trivial. For example, the Perl build process has become so complicated that a separate development effort has been created just to write the configuration scripts and `Makefile` for each release [20]. Furthermore, the GCC system (GNU Compiler Collection) exhibits even more dynamic and interesting build-time behaviours, including as multiple passes of compilation and a significant amount of automatic source code generation.

For such large and complex systems, build management is typically performed by dedicated employees ("build engineers") using a specialized build management system. The main responsibilities of a build management system are to configure and manage the build scripts, to provide the ability to define repeatable system build procedures, to maintain
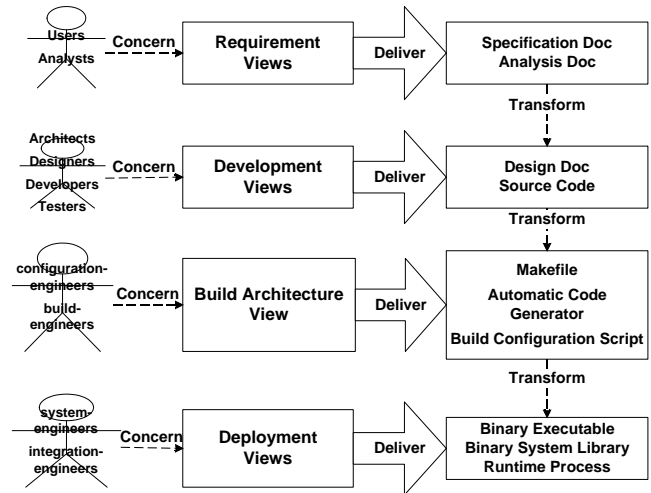

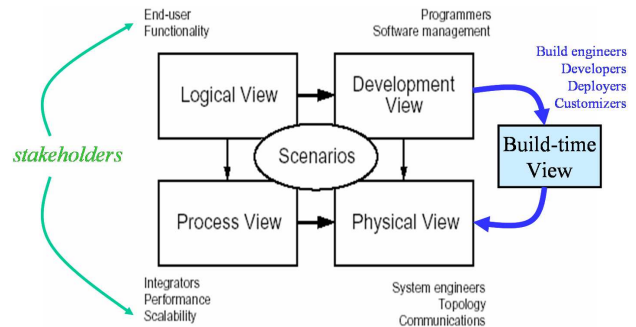Figure 1: Transformation between architecture views.

Figure 2: Kruchten's "4+1" model enhanced with the build-time architectural view.

consistency in system builds, and to control the various build tools that are involved in the process.

Build-time architectural views should capture configuration and built-time properties that are extractable from build management artifacts, such as build scripts, the source and object files, and configuration choices. At the very least, they should model the compilation dependencies among compilation units, time-sequence configuration of the compilation procedure, and source code generation at build-time.

As mentioned above, neither the "4+1" or "four views" models explicitly addresses the idea of build-time architectural views. Figure 1 shows the relationships between build-time views and the other kinds of architectural views, and Fig. 2 shows how the "4+1" model may be extended to include build-time views.

As shown in Fig. 1, a software system undergoes a series of fundamental transformations as the result of development and deployment efforts, and different architectural views are important to different stakeholders throughout this

process. For example, the software designer studies the logical view of the software system, and applies various design techniques to layout the development view. When the development phase is complete, the build engineers write the `Makefile` and actually build the system from a collection of source code files (development view or code architecture view) into a suite of executables and libraries (execution architecture view or physical view) that work together closely to deliver the functional and non-functional requirements demanded by the customer.[2]

2.4 Build-time views and configuration/build management

Build-time architectural views can serve as high-level and visual documentation for software configuration and build management. As the build architectures of software systems become more dynamic and complex, the importance of explicitly modelling build-time artifacts and their related activities increases.

One common reason to have a complex build-time architecture is to aid in the cross-platform construction problem, when software systems are required to build and run on a variety of hardware platforms and operating systems using the same source code distribution. The usual solution involves running a configuration tool, such as the Unix tool `configure`; such a tool queries the operating system for hardware and operating system details. This information is then used together with a set of provided `Makefile` templates to create customized build scripts that are tailored for the target environment. Often, this information is straightforward in nature, and entails the setting of various source code macros and compiler flags. However, in the case of GCC the data structures of the Register Transfer Language (RTL) and algorithms to manipulate them are fundamentally different for each CPU architecture. The task is far beyond the capability of `imake` and conditional compiling techniques. The solution employed by GCC is to develop a special program to behave as an automatic code generator. It takes architecture description files as input and emits appropriate algorithms and data structures as output. As a result, some fundamental parts of the source code for GCC are created during build-time. This unique design approach makes the build-time architecture of GCC very dynamic, and it is not well represented by static architectural views. Many interesting dynamic build architectures may be found within the domain of programming language compiler/interpreter systems. This is because the data structures and algorithms (especially algorithms for code generation and optimization) are highly dependent on the target CPU architecture and operating system.

---

[2]We note that this diagram is somewhat idealized, and shows the logical progression of the development of a new software system. Experience suggests that industrial practice is much more conceptually messy.
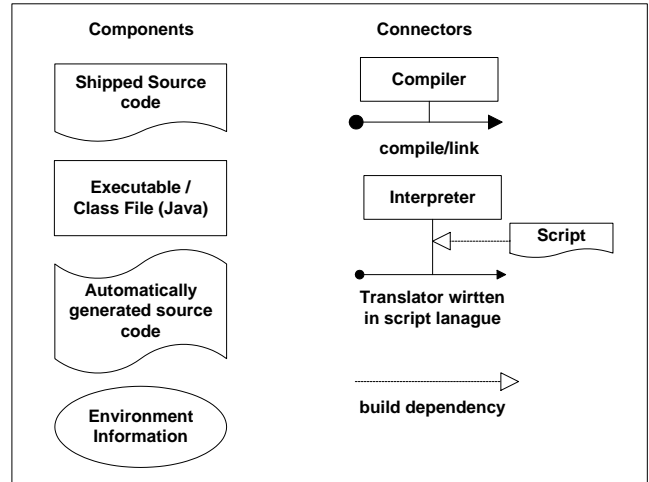

Figure 3: Notation of build-time architectural view.

It is important to point out that there are many software/computing systems that use various dynamic configuration techniques at run-time, such as Microsoft's .NET/COM/DCOM, CORBA, Enterprise Java Beans, and the Java reflection API. For example, CORBA allow distributed applications to call a remote service without knowing its exact location or even the programming language in which it is implemented. However, these techniques are part of the execution architecture, where the distributed components have already been build and deployed. Their dynamic run-time architectures can be modelled by the "physical view" in the "4+1" model and the "execution architecture view" in the "four views" model.

3 Modelling build-time views

We now give an informal definition of our notation for visualizing build-time architectural views; this notation will be used to represent build-time views of the three example systems in the next sections.

As shown in Fig. 3, we define four types of components and three types of connectors in our build-time architecture view:

Components: A box with a straight top edge and a curved bottom edge represents source code that has been written before the build and is provided within the original source distribution. A box with curved top and bottom edges represents source code that is dynamically generated at build-time. A box with straight edges represents a program executable or Java class file. An ellipse denotes environmental information, such as CPU and operating system details.

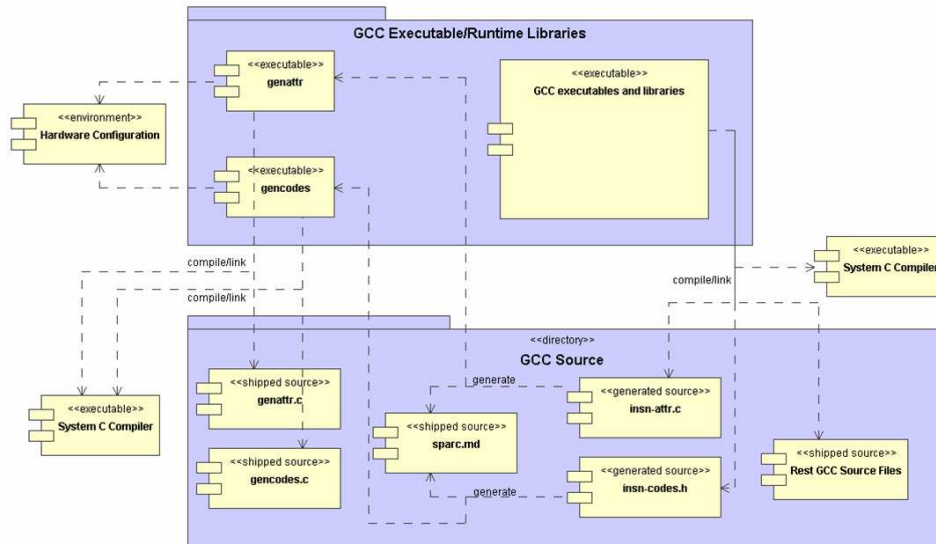**Connectors:** For a compilation link, an arrow is drawn from

Figure 4: A built-time view of source file dependence modelled using a UML component diagram.

the source code to the compilation target; the compiler used may be indicated as an association to the relation arrow. A common generalization of this connector is the interpreter/translator arrow, where a transformation is performed based on a provided script; in this case, the script is denoted by the addition of a hollow-headed arrow. The third type of link is a general build dependency, which is represented with a hollow-headed arrow.

We note that one could also use the UML notation to model build-time views. As we have stated elsewhere [23], "static" dependencies between components (possibly involving components generated at build time) can be modelled using a UML component diagram (as shown in Fig. 4), and dynamic build-time behaviours can be modelled using UML sequence diagrams. However, it is not our intent here to discuss the additional customizations that would be required (for example, UML stereotypes); instead, we merely mention that our notation shown in Fig. 3 is but one possible approach to modelling build-time properties of software systems.

We now examine as case studies three software systems that have interesting build-time views: GCC, Perl, the use of the Java Native Interface (JNI).

3.1 Example: Build-time views of GCC

The GNU Compiler Collection (GCC) is a suite of compilers that supports several programming languages (C, C++, Object-C, *etc.*) and runs on a variety of hardware architectures and operating systems [11]. GCC is an open source software system: the source code is freely available, and it is

therefore possible for users to compile and install their own customized versions. In fact, portability and customizability are major design goals for GCC, and this has had a significant influence on its software architecture.

In examining the GNU C compiler (gcc), a subcomponent of GCC, we have found that it exhibits two particularly interesting phenomena at build time: it uses a multiple-pass compilation process (also known as "bootstrapping"), and it automatically generates significant portions of the system source code.[3] In this section, we will try to capture these phenomena with a formalized build-time architectural views, which will clearly define the components, form, and rationale. The GCC version used in this case study is 2.7.2.3.

3.1.1 Bootstrapping in gcc

Bootstrapping is a common technique employed in compiler design to "use the facilities offered by a language to compile itself" [1]. A typical bootstrapping process consists of several steps [17]:

1. Compile a reduced version of the compiler in a different environment.

2. The first reduced compiler is used to compile itself for the target machine.

3. The compiled compiler then compiles itself on the target computer.

---

[3]Note that we use the term 'GCC' when referring to the full compiler suite, and we use the term 'gcc' when referring only to the C compiler component of GCC. Although confusing, this convention is the standard one.

4. The language is enhanced to its full capability and the enhanced version recompiled.

5. If code optimization was left out in previous step, then subsequent recompilation will improve the compiler's own performance.

A full-scale bootstrapping is essential for developing the first compiler of a new programming language, or to support a new hardware architecture. More commonly, a simplified bootstrapping procedure is used to develop a new version of the same compiler system. Bootstrapping also provides an excellent opportunity to look for bugs in the compiler, since the compiler source code is usually large and complex.

In many compiler systems, the only artifact that documents the bootstrapping procedure and configuration is a `Makefile` (or hierarchically organized family of `Makefiles`). However, the `Makefile` itself is usually inscrutable to humans due to the size and complexity of the system, and because it is usually highly templated and parameterized to account for cross-platform issues. The `Makefile` is expanded at build-time after configuration tools have collected information about the hardware and operating system then tailored the `Makefile` for this specific environment. As a result, the only simple and practical way to infer the behaviour that occurs during bootstrapping is to log the execution of `make` at build time, and then to review the log file.

The build-time behaviour of GCC during bootstrapping is shown in Fig. 5. During the bootstrapping process, three different GCC compilers are built. The first one is built by the default system C compiler and linker, and the remaining two are built by GCC itself. In all three builds, the same source files are compiled.[4] Three copies of the GCC compiler executables `gcc` are created at different time and each but the last is immediately used to compile the next version.

First, the existing C compiler on the build platform is used to compile the GCC driver, the C language compiler `gcc`, and the GCC libraries (not the C system library, but the supporting routines for other parts of GCC) from the GCC source code. When the first build is completed, we have an intermediate GCC C compiler that is fully functional. We call it "stage 1" GCC.

For the second build, we run the product from last build, the "stage 1" GCC, on the same GCC source code again. This time, we compile not only the GCC driver, the C compiler, and the GCC libraries, but also the C++ compiler, the Object-C compiler and their supporting libraries. At this moment, all the functional components of GCC 2.7.2.3 (C, C++ and Object-C compilers) are built and integrated with a uni-
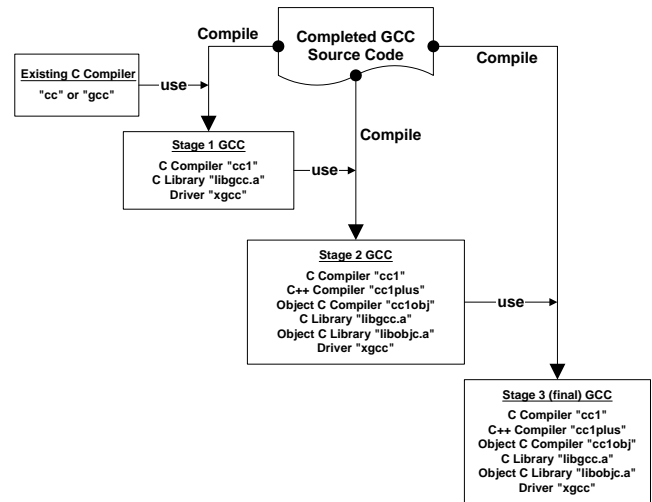
[4]The `gcc` compiler supports various extensions to the C language that are not part of the ANSI C standard, such as comments that begin with "//". Many systems, such as the Linux kernel, take advantage of these language extensions. However, since GCC is usually first compiled by a non-GCC C compiler, the GCC system itself is written entirely in ANSI C.

Figure 5: Build-time view of GCC bootstrapping.

fied compiler driver `xgcc`. We call this intermediate compiler suite "stage 2" GCC.

"Stage 2" GCC is still not the final product. We run it over the GCC source for the third and the last time. The product of this round of build should be identical to the "stage 2" GCC, because they are complied from the same source code and both by the GCC 2.7.2.3 C compiler. The purpose of this build is for self-testing only: if "stage 2" GCC and "stage 3" GCC are different, this implies that something has gone wrong with the compilation and that there are still some bugs to be worked out.

The build architecture view of this aspect of GCC is shown in Fig. 5; the details were manually reverse engineered from the log file of an actual build of GCC 2.7.2.3 on a Sun Solaris 2.6 server with the bootstrapping build option turned on.

This build architecture view shows the dynamic behaviour and configurations of the build process of GCC. The constructs of the view, such as boxes and arrows, have different semantics from other architectural views we have been familiar with. For example, in the traditional module view, the software system is represented as a static structure of a group of programming modules, while the links between components (modules) represent the semantic dependencies and run-time interactions. In the process view, the system is represented as a snapshot of a running organism with different processes collaborating with each other. However, the build architecture view is much more similar to a time-space diagram. The components inside the view do not co-exist at the same moment. The sequence of events happen in the order of from top to the bottom, and from the left to the right. The components, which should be interpreted as either static data or execution units, came into existence in sequence along with the events that modelled as arrows in

the view. The arrows between components also represent the build-time (dynamic) dependencies between compilation entities.

3.1.2 Automatic code generation during build-time

In this section, we will discuss another interesting aspect of the build architecture of GCC, and its role in the architectural transformation of system views.

The build process of GCC consists of two major activities: `configure` and `make`. The "`configure`" process first checks whether the build environment is supported by GCC. Next, `configure` probes the build environment for information such as the CPU architecture, operating system, available system libraries, *etc.* Based on the probed information, it generates the `Makefiles` using the templates and system configuration information that `configure` just collected. In the "`make`" phase, the compilation of the source, linking of object file and installation of final binaries are performed.

For many software systems, including the Linux kernel, Mozilla, and VIM whose architectures we have analyzed before [3, 8, 21], the build process is fairly straightforward: virtually all of the source code that will be used to compile the executable system has been created during the development phase and shipped with the source distribution, except a few system-dependent header files that are automatically generated by `configure`; most of these files contain simple macros and system-dependent constants. However, GCC has a much more complex and dynamic building process, where the software architecture undergoes a significant reshaping — some of the most important source code files in the GCC system are automatically generated, then compiled and linked to create the final compiler system.

One of the main design goals of GCC is to support multiple programming languages, hardware architectures, and operating systems. Not only does GCC support a wide range of languages and platforms "out of the box", the compiler suite has also been designed to allow easy expansion and customization by third-party developers. To achieve these goals, the architecture of GCC must be both flexible and configurable. To that end, the individual GCC compilers share many of the front-end and back-end components, such as parts of the compiler driver, the preprocessor, the RTL (an intermediate code format), the object code generator, and the optimizer. Support for multiple programming languages is achieved entirely within the downloaded source code of GCC; however, support for multiple CPU architectures is implemented largely by generating part of the source code at build-time, based on information gleaned about the target environment by the configuration process.

In GCC, the Register Transfer Language (RTL) is an intermediate representation used to represent the target system's code after parsing, similar to Java byte code. However, unlike Java bye code, RTL is hardware dependent. The spec-
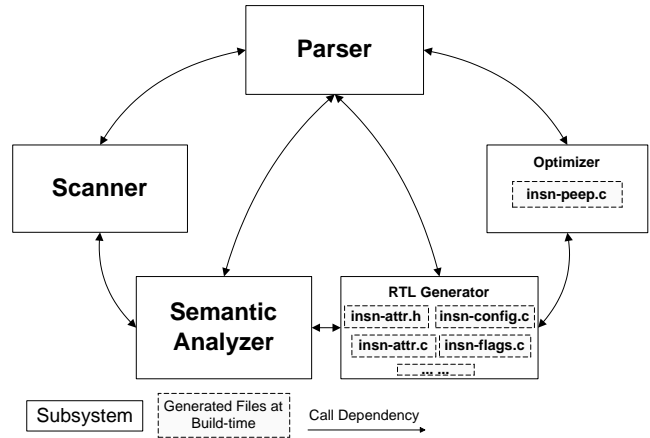

Figure 6: Code architecture view of the compiler-core subsystem in GCC.

ification of RTL and the portion of source code that operates on RTL are generated at build time, using machine description information and collected system parameters from `configure`. The main benefit of having a target-dependent RTL representation is that we can immediately generate the target machine language (assuming an infinite number of registers), but in a way that the compiler can understand and manipulate the emitted instructions. Hardware-dependent optimizations also operate on this intermediate format, and only valid instruction for the machine is generated as results of all passes of transformation, for RTL has built-in knowledge about target CPU architecture [15].

Figure 6 shows a portion of the code architecture view of GCC 2.7.2.3 with "holes" (dashed boxes) that represent the "missing" source code files. The GCC system consists of five major subsystems: the driver, the preprocessor, the compiler-core, the code-generator, and the utility libraries. Both the compiler-core and code-generator subsystems contain the RTL manipulation code that is missing from the distribution. The internal code architecture of compiler-core subsystem is illustrated in Fig. 6.

The "missing" files in the compiler-core subsystem are generated at build-time from code templates by source code generators.[5] The procedure is explained here and illustrated in Fig. 7 with a build architecture view diagram. In addition, Fig. 7 shows the relationship between build architecture view and code architecture view from above, and execution view from below.

1. First, the (build-time) source code generators are compiled; the source code for these generators are contained within files whose names begin with "`gen`". The result is a set of executable programs.

---

[5]The source code generators shipped with the GCC source should not be confused with the object code generator subsystem of GCC, which is a standard component of any compiler.

Filename	Role
sparc.h	Contains C macros that define the general attributes of the Sun SPARC architecture
sparc.c	Contains machine description supporting functions and macro expansion
sparc.md	Contains RTL expressions that define the instruction set (template). This file provides the input to programs that produce insn*.h and insn*.c files.

Table 1: GCC machine description files for Sun SPARC.

Filename	Created by	Description
insn-attr.h	genattr	Definition for any defined attributes and delay-definitions
insn-attr.c	gen-attrtab	Functions to access the attributes
insn-codes.h	gencodes	Definitions for named pattern

Table 2: Some GCC source files generated at build-time.

2. Next, these code generators are executed in sequence. They take machine description files for the target machine as input. The machine description files are picked by configure. Table 1 lists the machine description files for Sun SPARC architecture. The output is a collection of C source files, such as those listed in Table 2; these generated files have names that begin with "insn". These C files are used to fill the "holes" in the code view of compiler-core subsystem.

3. Finally, the source files to build a working GCC are all available. We now compile the code from the source distribution together with build-time generated code, and link them together to create the GCC compiler system.

Thus, the build-time architecture shows how the GCC system "fills in the gaps" of the code view of the shipped source code that were (intentionally) left by the GCC developers.

3.2 Example: Build-time views of Perl

Perl version 5.6 [24, 12] is an interpreted high-level programming language in very wide use, originally developed by Larry Wall. Like GCC, Perl is an open source system, which means that portability of the source is a paramount design goal. Also like GCC, Perl exhibits complex and interesting properties in its build-time architecture, which we now discuss.

By comparing the content of Perl source directory before and after build, we found that there are 22 C source code files
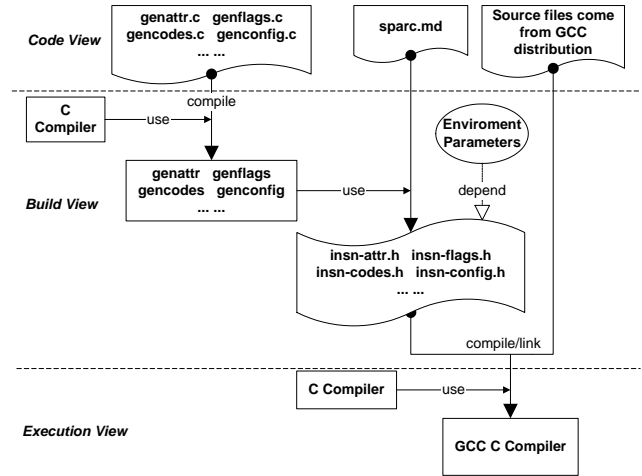

Figure 7: Build-time view showing GCC source code generation.

Filename	Description
B.c	Allow a Perl program to delve into its own innards. It is used to implement the "backends" of the Perl compiler.
ByteLoader.c	Used to load byte compiled Perl code. It uses the source filter mechanism to read the byte code and insert it into the compiled code at the appropriate point.
DB_File.c	Allows Perl programs to make use of the facilities provided by Berkeley DB.

Table 3: Some C files generated at build-time in Perl.

newly generated during the "make". The log file of make reveals that these C files originate from templates written in a special language called XS. The XS templates are then processed by an intermediate Perl interpreter miniperl with a translation script xsubpp.pl. The results of the processing are those new C source files mentioned above. Finally, the C files are compiled and linked by C compiler gcc as a part of Perl 5.6 run-time libraries.

Table 3 lists the filenames and descriptions of some of the C files that are generated at build-time; these particular files enable the Perl system to call Unix system libraries, Berkeley DB and Sys V IPC. It also enable Perl programs to directly manipulate the backend Perl run-time system. These C files are part of the Perl source code but they are generated at build-time. The templates for the C files are written in a special language called XS and they come with the Perl source distribution. We now explain the XS language and its relationship to Perl [24].

Assume you are a programmer who wants to implement a function in C or call a C system library from Perl. To do so, you need to write a special C library that can be either

8

dynamically loaded by the Perl run-time system or statically linked into the Perl executable. This library acts as "glue" that links the C code with the Perl system. For example, when a C library function is called from the Perl program, the glue library first pulls arguments of the call from Perl's argument stack and converts the Perl values to the formats expected by the particular C function. Then, it calls the C function and finally transfers the return values of the C function back to Perl. To pass the result back to Perl, the glue library either puts them on the Perl stack or modifies the arguments supplied by Perl.

Of course, a programmer can craft a "glue" library directly in C, but this is very tedious and requires knowledge about the working mechanisms of Perl stack. Therefore, the Perl distribution provides a standard interface and mechanism to create glue libraries that is both flexible and easy to use. The interface is written in XS language that allows a programmer to describe the behaviours of the glue. The XS language describes the mapping between a Perl function and a C function. It can also define a wrapper Perl function that wrap around a C function. The XS compiler `xsubpp`, which comes with the Perl source distribution, compiles the XS file and generates the glue automatically. The output of `xsubpp` is the C source file that implements the glue library, as we can see in Table 3.

Now that we understand the procedure to create a Perl-C extension, we can investigate and create models for the build-time architecture of of Perl 5.6. The Perl source distribution includes 22 Perl-C extension templates written in XS language. These extensions are either part of the Perl run-time environment, or interface to important Unix/C libraries. The translation process to generate the Perl-C extension from the XS templates is operating system dependent. Therefore, these extensions are created at build-time, since the operating system environment parameters are only available after running `configure` on the build platform.

Similar to GCC, Perl also bootstraps itself. An intermediate Perl interpreter called `miniperl` is built first, then `miniperl` and `gcc` are used together to build the complete Perl interpreter and run-time system. The translation script that transform XS file to C file is also executed by this `miniperl`.

The build architecture view of Perl is shown in Fig. 8, where the build-time code generating and bootstrapping are clearly illustrated. At the top of Fig. 8, the code architecture view corresponds to the directory structure and organization of Perl source code. Components of the code architecture view shown here include the source files to build `miniperl`, the translation script from XS to C `xsubpp`, Perl-C extension templates `.xs` files and other C code to build the rest of Perl interpreter and run-time.

The first step of build creates an intermediate Perl interpreter `miniperl`, which provides a functionality subset of the full-scale Perl interpreter. Although `miniperl` is not as
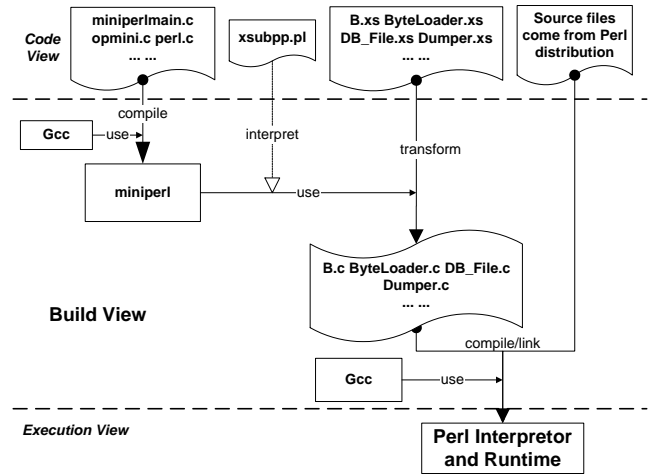

Figure 8: The build-time view of compiling Perl 5.6.

efficient and optimized as the final Perl interpreter `perl`, it is sufficient to accomplish the required tasks.

The second step is to create another important component of the build architecture: the C files that implement the "glue" libraries for system Perl-C extensions. The translation is carried out by the Perl interpreter `miniperl` and translator utility script `xsubpp`. The translator script takes `.xs` files as input, and outputs the C files, such as those listed in Table 3.

The final step is to compiler all of the C source files and linked them together to create the Perl interpreter and run-time libraries.

3.3 Example: Build-time views of the JNI

The Java Native Interface (JNI) is a programming interface for Java. Using the JNI allows Java programs running within a Java virtual machine (JVM) to interact at run-time with applications and libraries written in other languages, such as C, C++ and assembly language. Similar to the Perl-C extension provided in Perl, JNI enables Java programs to take advantage pre-existing native applications, to exploit platform-dependent utilities, or to improve the performance of system bottlenecks [5].

Similar in spirit to the Perl XS interface, the JNI also provides a standard framework with a collection of Java APIs and tools. They work together to allow the programmer to "glue" Java application with native methods (typically, C or C++ library functions).

A typical use of JNI involves several steps:

1. First at the development stage, we write Java code that declares only the signature of the native method. The native method can be called from any other part of the Java program as if it were a "normal" Java method. We also implement the native method in C or C++.
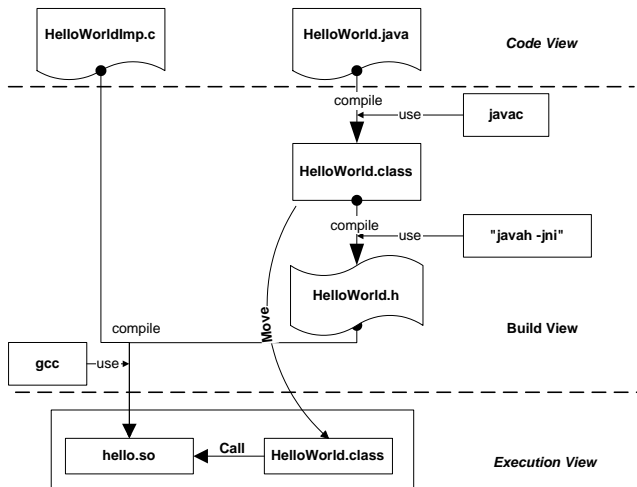
9

Figure 9: The build-time view of a simple use of the JNI.

2. At build-time, we compile the Java class that declares the native method and all other Java classes including those who call the native method. Next, we use the JDK tool `javah` to generate the header file for the native method. Now we have the native signature for our C/C++ implementation. Then, we compile both the header file and the implementation C/C++ file into a shared library file.

3. Finally, we have the running application, where the Java code from the JVM can call the functions defined in the native library.

A build-time architectural view of using Java JNI is shown in Fig. 9 using our notation. This simple example shows a Java program that calls a C program that prints "Hello world" [5]. We can compare this view to the the build-time view of Perl shown in Fig. 8; we note that unlike the previous two example systems, in this case the development of the applications and their (partial) construction are intertwined, rather than being neatly separated into distinct phases.

3.4 The "code robot" architectural style

In their book on software architecture, Shaw and Garlan discussed several *architectural styles* that model recurring abstract patterns of high-level structure within software systems, such as "pipeline", "layered", "client-server" and "interpreter" [18]. In Kruchten's "4+1" paper, each of his architectural views indicates some representative architecture styles. For example, the "object-oriented" style is used in logical view, while the "pipe and filter" and "client-server" styles are applied in the process view [13].

We consider that the dynamic build-time behaviour of systems such as GCC and Perl defines a new architectural style
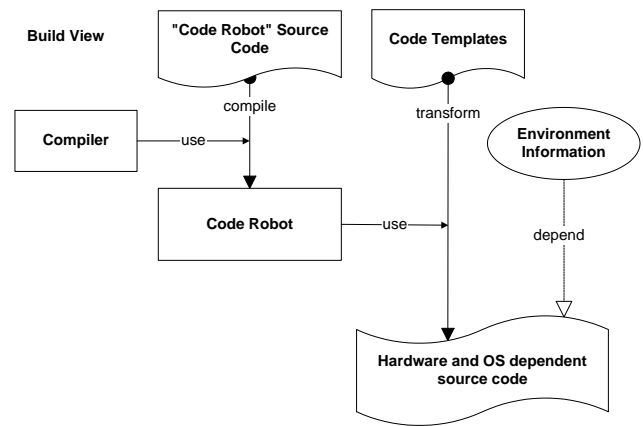
that applies to build-time architectures; we call it the "code robot" architecture style. The idea is, if the behaviour of the software system depends heavily on the hardware architecture or operating system, the software designer must devise an effective and sophisticated strategy for customization of the system source code at build-time. For example, the strategy taken by the developers of GCC and Perl is to write a code generator, a "code robot", such as the `gen*.c` in GCC and `xsubpp` in Perl. In Perl, the system-dependent code is specified by a template written in XS language. In GCC, the rules of how to create them from hardware architecture description files are embedded in the code robot itself. Given a description of hardware architecture, the code robot knows how to generate corresponding system-dependent code. Figure 10 shows the "code robot" architecture style.

Figure 10: The "code robot" architecture style.

3.4.1 Code robots and open source software

Automatic source code generation is not new; it is a well known technique that is in wide use in both industry and by open source projects. For example, the code wizard in Microsoft's VisualStudio can create generic MFC skeleton code for a Win32 application, which includes standard windows, menus, dialog boxes, and shortcuts. Others examples include `lex`, `yacc`, and their relatives which generate scanners and parsers for compilers.

Automatic source code generation *during system building* is much less common. The reason seems to be that most commercial software systems are targetted for a relatively small number of possible hardware architectures and operating systems. Most companies do not ship the source code for their systems to clients; consequently, system building is performed prior to shipping, and the software company can choose its preferred platforms, compilers, subtools, and libraries. This means that system building can be managed in a straightforward manner, often by creating a product line for each target environment.
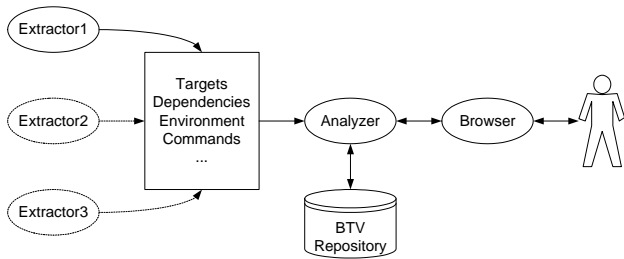
Figure 11: Architecture of the BTV Toolkit.

However, open source software systems are, in general, designed to be as portable as possible. Since the source code is available, the system is typically built by the user; this means that the source code distribution must be designed to buildable on a large number of possible platforms, using different compilers and subtools, and using different system libraries. Rather than create a separate source distribution for each set of alternatives, open source systems usually abstract the commonalities into a single distribution and rely on configuration tools to aid in building. This is why the code robot style is most commonly found within open source software systems.

4 Extracting Build-Time Views: The *BTV Toolkit*

When we did the work to examine the build-time views of GCC and Perl, the build architectures were manually reverse engineered based on execution logs of make; we found this process to be inefficient and error-prone. We have been therefore begun developing a toolkit for supporting the extraction and visualization of build-time views from software systems called the *BTV Toolkit* [6]. Figure 11 shows the current proposed architecture for the tool.[6] There are four major components:

1. a set of *fact extractors* that operate on source artifacts (usually source code, configuration scripts, and Makefiles),

2. an *analysis engine* that can reconcile "facts" from different extractors to produce a unified set of abstracted build views, and also respond to user queries from the browser to generate new views,

3. a *repository* for build-time view facts, and

4. a *browsing tool* to support visualization, structural navigation, and querying of build-time views.

---

[6]We note that this architecture is very similar to that of most reverse engineering tools, such as SWAGkit [19] or Rigi [14]. This is not surprising, since the *BTV Toolkit* is effectively a tool for reverse engineering build-time architectural models of software systems.

Our current implementation[7] consists of

- a fact extractor that is effectively an instrumented version of gmake, the GNU version of make,

- an analysis engine,

- a plain-text file-based repository mechanism, and

- a browsing/visualization tool that is based on the graphviz engine [7].

Figure 12 shows an example (partial) view of the construction of the Apache httpd program, as extracted and modelled by our current prototype. In this diagram,

- a box represents a source code directory in which the system build process has been active,

- a double ellipse represents a top-level goal of an invocation of make,

- a single ellipse represents a *target* of the make (source and derived files are eventual targets of make),

- a solid arrow represents a build dependency defined explicitly or implicitly in the Makefile, and

- a dotted arrow represents the instantiation relationship between a (parent) make process, and the sub-make processes it creates.

If a box contains another box, this indicates that the corresponding directories are nested in the obvious way. If a box contains an ellipse, this indicates that the corresponding target "lives" in the corresponding directory; that is, either the "phony" target was executed in that directory, or the (source or derived) file-target exists there.

In this particular figure, some relationships have been elided: some arrows are not drawn, and all but one of the boxes (*i.e.,* srclib) have their contents hidden. An arrow going into a box (*e.g.,* the os directory on the right) indicates that something inside the box/directory is the object of a build relationship.

We have used the current implementation of the *BTV Toolkit* to create models of the construction of several large software systems, including GCC, Apache, Perl, PostgreSQL, and Vim. The applicability of the current implementation is limited to those systems that can be constructed using gmake. Also, the information that is currently extracted is limited to compilation dependencies and actions as observed during the build.

We note that the visual notation shown in Figure 12 is not consistent with that introduced in Section 3. This is because it is at a lower level of abstraction than the relatively high-level phenomena described in the previous examples. We

---

[7]The current version of the *BTV Toolkit* can be downloaded from http://www.swag.uwaterloo.ca/~xdong/btv
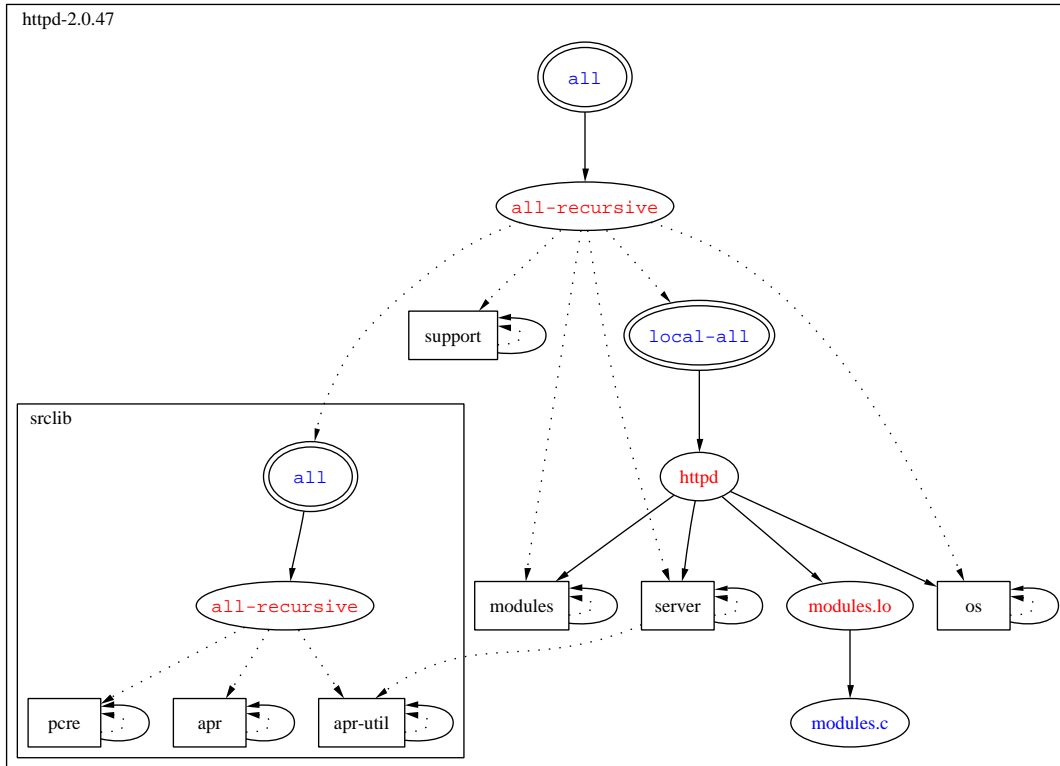
11

Figure 12: Build-time view of Apache `httpd`.

have elected to show "raw" build relationships here between various build targets (*i.e.,* source and generated files, as well as "phony" targets) embedded in the context of the directory hierarchy of the source code. We intend to provide support for creating build-time view models similar in abstraction level to those diagrams shown in Section 3, based on the kind of information shown in this diagram.

5 Summary

In this paper, we have explored the idea of build-time software architectural views. We motivated the need for explicitly extracting and modelling build-time views for software systems that have complex and subtle dependencies on the the underlying technologies from which they are constructed, and we showed how the build-time architectural views relate to well-known software architectural taxonomies. We presented a notation for modelling build-time views, and we explored the characteristics and the significance of software build-time architectures in three case studies: GCC, Perl, and the use of Java's JNI. Through our case studies, we discussed how explicitly modelling build-time behaviours with these architectural views can aid developers in gaining a better understanding the software system itself, as well as in managing the build process in the

software development/deployment cycle. We introduced the "code robot" architectural style to aid in modelling build-time views of systems, and discussed why it is most common in open source software systems. Finally, we sketched the design of a toolkit that is under development for the semi-automated extracting and modelling of build-time architectural views.

6 Acknowledgments

We thank our colleagues Ric Holt and Andrew Malton for their ideas, comments, and feedback. Also, a preliminary version of some of these results were described in a conference paper presented at ICSM 2001 [22].

References

[1] Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman. *Compilers: Priciples, Techniques and Tools.* Addison Wesley, 1986.

[2] L. Bass, P. Clements, and R. Kazman. *Software Architecture In Practice.* Addison Wesley, 1998.

[3] Ivan T. Bowman, Richard C. Holt, and Neil V. Brewster. Linux as a case study: Its extracted software architecture. In

Proc. of the 21^{st} *Intl. Conference on Software Engineering (ICSE-21)*, Los Angeles, CA, May 1999.

[4] F. Buschmann, R. Meunier, H. Rohnert, P. Sommerlad, and M. Stal. *Pattern-Oriented Software Architecture*. John Wiley and Sons, 1996.

[5] M. Campione, K. Walrath, and A. Huml. The Java tutorial: A practical guide for programmers. Website. http://java.sun.com/docs/books/tutorial/java/.

[6] Xinyi Dong. BTV: The Built-Time View toolkit. Website. http://www.swag.uwaterloo.ca/~xdong/btv.

[7] John Ellson *et al*. Graphviz: Open source graph drawing software. Website. http://www.research.att.com/software/tools/graphviz/.

[8] Michael W. Godfrey and Eric H. S. Lee. Secrets from the monster: Extracting Mozilla's software architecture. In *Proc. of 2000 Intl. Symposium on Constructing software engineering tools (CoSET-00)*, Limerick, Ireland, June 2000.

[9] C. Hofmeister, R. Nord, and D. Soni. *Applied Software Architecture*. Addison Wesley, 2000.

[10] Richard C. Holt, Michael W. Godfrey, and Andrew Malton. The build / comprehend pipeline. In *Proc. of 2003 ASERC Workshop on Software Architecture*, Banff, Alberta, February 2003.

[11] http://www.gnu.org/software/gcc/. GCC homepage. Website.

[12] http://www.perl.com. The Perl homepage. Website.

[13] P. Kruchten. The 4+1 view model of architecture. *IEEE Software*, 12(5), November 1995.

[14] Hausi A. Müller and Karl Klashinsy. Rigi: A system for programming-in-the-large. In *Proc. of the* 10^{th} *Intl. Conference on Software Engineering (ICSE-10)*, Singapore, April 1988.

[15] Hans-Peter Nilsson. Porting GCC for dummies. Website. ftp://ftp.axis.se/pub/users/hp/pgccfd/.

[16] Dewayne E. Perry and Alex Wolf. Foundations for the study of software architecture. *ACM SIGSOFT Software Engineering Notes*, 17(4), October 1992.

[17] Thomas Pittman and James Peters. *The Art of Compiler Design, Theory and Practice*. Prentice-Hall, 1992.

[18] Mary Shaw and David Garlan. *Software Architecture: Perspectives on an Emerging Discipline*. Prentice Hall, 1996.

[19] SWAGkit homepage. Website. http://www.swag.uwaterloo.ca/swagkit/.

[20] The Perl mailing lists. Website. http://www.perl.org/support/mailing_lists.html.

[21] John B. Tran, Michael W. Godfrey, Eric H. S. Lee, and Richard C. Holt. Architecture analysis and repair of open source software. In *Proc. of 2000 Intl. Workshop on Program Comprehension (IWPC-00)*, Limerick, Ireland, June 2000.

[22] Qiang Tu and Michael W. Godfrey. The build-time software architectural view. In *Proc. of 2001 Intl. Conference on Software Maintenance (ICSM-01)*, Florence, Italy, October 2001.

[23] Qiang Tu and Michael W. Godfrey. Representing build-time software architecture views with UML. In *Proc. of 2001 ICSE Workshop on Describing Software Architecture with UML*, Toronto, Ontario, May 2001.

[24] Larry Wall, Tom Christiansen, and Jon Orwant. *Programming Perl*. O'Reilly & Associates, 2000.