**Is it getting any easier
to understand
other people's code?**

Code Spelunking Redux

George V. Neville-Neil, Consultant

It has been five years since I first wrote about code spelunking,⁸ and though systems continue to grow in size and scope, the tools we use to understand those systems are not growing at the same rate. In fact, I believe we are steadily losing ground. So why should we go over the same ground again? Is this subject important enough to warrant two articles in five years? I believe it is.

The oft-quoted Moore's law about the increasing power of computers actually works against the code spelunker. The more powerful computers become, the more we demand that they do, which increases the complexity of the software that runs on them. Processor speeds increase, and that

means more lines of code can now be run in the same amount of time. Available memory gets larger, so we can now keep more state or code in memory. Disks get larger and require less power (in the case of flash), and suddenly we're able to carry around what were once considered huge amounts of data in our pockets. What was termed the "software crisis" in the 1970s has never really abated, because each time software engineers came up with a new way of working that reduced complexity, the industry moved forward and demanded more.

Complexity increases in many directions, including lines of code, numbers of modules, and numbers of systems and subsystems. More complex systems require more lines of code to implement. As they grow their systems, software teams often integrate more code from outside resources, which leads to complex interactions between systems that may not have been designed with massive integration in mind.

These numbers should not be surprising to any software engineer, but they are a cause for concern. Although it was unlikely that the numbers would shrink, all but one of them has grown by more than 50 percent, and although the number of lines may have grown linearly, the interactions among the new components that these numbers represent have not grown in a linear fashion. If we assume that all modules in a system can interact freely with all other modules, then we have a system in which the potential number of interactions is expressed as $n(n-1)/2$, an equation that should be familiar to those who work in networking as it represents a fully connected network. If a system grows from 100 modules to 200 modules, a 100 percent

growth rate, then the number of potential connections grows from 4,950 to 19,900, a 302 percent growth rate. One reliable measure of the number of interfaces into a system is the number of system calls provided to user programs by an operating-system kernel. Since the publication of my first article on code spelunking, the Linux kernel has grown from just shy of 200 system calls to 313, an increase of more than 50 percent (see table 1).⁶

TWO NEW TOOLS

My first article on code spelunking covered several tools, including global,³ Cscope,¹ gprof,⁴ ktrace,⁷ and truss.¹¹ I continue to use these tools on a daily basis, but in the past five years two new tools have come to my attention: Doxygen² and DTrace.⁹ Although they may not have been specifically designed with code spelunking in mind, both make significant contributions to the field. Here, I discuss each tool and how it can help us understand large code bases.

Doxygen. Right at the top of the Doxygen Web page² is the following statement: "Doxygen is a documentation system for C++, C, Java, Objective-C, Python, IDL (Corba and Microsoft flavors), Fortran, VHDL, PHP, C#, and to some extent D." As the blurb says, Doxygen was designed with documenting source code in mind—and it is quite a good system for documenting source code so that the output is usable as man pages and manuals—but it has a few features that make it applicable to code spelunking, too.

What Doxygen does is read in all, or part, of a source tree, looking for documentation tags that it can extract

TABLE 1
Growth in Program Size

Program	Version	Files	Lines	%Chg Lines
Apache Web Server	1.3	471	158,332	--
	2.2.8	1108	374,993	136%
Emacs	21	2586	1,317,915	--
	22	2598	1,771,282	34%
FreeBSD Kernel	5.1	4758	2,140,517	--
	7.0	6723	3,556,087	66%
Linux Kernel	2.4.20-8	12417	5,223,290	--
	2.6.25-3	19483	8,098,992	55%
Python	2.2.3	1158	356,314	--
	2.5.2	2379	910,573	155%

and turn into nicely formatted output suitable for documenting a program. It can produce Unix man pages, LaTeX, HTML, RTF, PostScript, and PDF.

What is most interesting for the code spelunker is Doxygen's ability to extract information from any source code by running preprocessors over the code in question. Doxygen is a static analysis tool in that it analyzes the source code of a program but does not look into the program state while it is running. The great thing about a static analysis tool is that it can be run at any time and does not require that the software be executing. In analyzing something like an operating system, this is extremely helpful. The features that make Doxygen most relevant to our work are those related to how data is extracted from the source code. When you start out with the intention of documenting your own code with Doxygen, you are already working with the system and very little extra needs to be done. If you're code spelunking an unknown code base, then you will need to be more aggressive and manually turn on certain features in the Doxyfile, which is Doxygen's configuration file. These features are listed here:

Feature	Meaning
EXTRACT_ALL	Extract everything you can from the source code
SOURCE_BROWSER	Create a full cross reference of the source code
CLASS_DIAGRAMS	Create class diagrams and inheritance graphs
HAVE_DOT	Create useful code spelunking graphs
CALL_GRAPH	Make a call graph following all function calls
CALLER_GRAPH	Output a graph of the caller dependencies

The option `HAVE_DOT` is the most important one because it's what allows Doxygen to generate the most useful output for the code spelunker, including class, collaboration, call, and caller graphs. We'll take a brief look at two of these types of graphs: *call* and *caller*. The code that we're analyzing in this article is the TCP/IP stack of the FreeBSD operating system. The BSD TCP/IP stack has been studied in the past¹² and continues to be studied by researchers working on the TCP/IP suite.

For our examples we look at a single function, `ip_output()`, which is called in various parts of the network stack in order to send an IP datagram to the network. The `ip_output()` function is quite important to the stack because all normal packet transmissions flow through it. If a bug were found in this function, or if the API needed to be changed for some reason, it would be important to trace back all of the current users (callers) of the function. In figure 1 we see the caller graph produced by Doxygen for `ip_output()`.

In figure 1 no fewer than 16 separate routines, in nearly as many modules, depend on the `ip_output()` function. To effect a fix or update the API, all of these routines need to be investigated.

The opposite of a caller graph is a call graph. A call graph is familiar to users of the tools mentioned previously, such as Cscope and global, which allow the user to move interactively through the call graph of a function, jumping into and out of underlying functions while browsing the source code. Doxygen gives us a different way of interacting with the call graph. Figure 2 shows the call graph for the `ip_output()` function.

The call graph, like the caller graph, provides a good visual overview of how the function fits into the overall system. Both of these figures function as maps from which we can derive clues as to how the software is structured. One clue that is relatively easy to see is another hot spot in the packet output code, namely `tcp_output()`, which is called from seven different routines.

The kind of information that Doxygen can show comes at a price. Generating the graphs shown here, which required analyzing 136 files consisting of 125,000 lines of code, took 45 minutes on a dual-core 2.5-GHz MacBook Pro laptop. Most of the time was taken up by generating the call and caller graphs, which are by far the most useful pieces of information to a code spelunker.⁵

DTrace. One of the most talked-about system tools in the past few years is DTrace, a project from Sun Microsystems released under the CDDL (Common Development and Distribution License) that has been ported to the FreeBSD and Mac OS X operating systems. Regardless of whether the designers of DTrace were specifically targeting code spelunking when they wrote their tool, it is clearly applicable.

DTrace has several components: a command-line program, a language, and a set of probes that give information about various events that occur throughout the system. The system was designed such that it could be run against an application for which the user had no source code.

DTrace is the next logical step in the line of program-tracing programs that came before it, such as ktrace and truss. What DTrace brings to code spelunking is a much richer set of primitives, both in terms of its set of probes and the D language, which makes it easier for code spelunkers to answer the questions they have. A program such as ktrace shows only the system calls that the program executes while it's running, which are all of the application's interactions with the operating system.

Code Spelunking Redux

On a typical operating system, these number in the low hundreds, and although they can give clues to what a complex piece of software is doing, they are not the whole story. Ktrace cannot trace the operating system itself, which can now be accomplished using DTrace.

When people discuss DTrace they often point out the large number of probes available, which on Mac OS X is more than 23,000. This is somewhat misleading. Not all of the probes are immediately usable, and in reality, having such an embarrassment of riches makes picking the most useful probes for a particular job difficult. A probe is some piece of code in an application, library, or the operating system that can be instrumented to record information on behalf of the user. The probes are broken down into several categories based on what they record. Each probe is delineated by its Provider, Module, Function, and Name. Providers are named after systems such as io, lockstat, proc, profile, syscall, vminfo, and dtrace itself. Several distinct providers are available in Mac OS X,

FIGURE 1 Caller Graph for `ip_output()`

FIGURE 2

Call Graph for ip_output()

although naively printing them all will show you that several exist on a per-process basis. The per-process probes show information on core data within the process, as well as on locks. Some of the providers available on Mac OS X are shown here:

Provider	Purpose
dtrace	Probes related to dtrace itself
fbt	Entry and exit points for functions
io	I/O probes
lockstat	Probes related to locking
plockstat	pthread lock-related probes
proc	Process-specific information
profile	Profiling and performance data
syscall	Information on system calls
vminfo	Virtual memory probes

Probes are associated not only with providers but also with modules, which are those you want to instrument, as well as with functions, which are also subject to observation. The name of a trace point specifies a single probe. All of these categories are put together in a DTrace script or command line to form a sort of address that specifies what the engineer is trying to observe. The canonical form is provider:module:function:name, with an empty section indicating a wildcard. Although the two manuals from Sun are excellent references on DTrace,^{9, 10} a quick example will demonstrate how it can be used for code spelunking.

When presented with a new and unknown system to spelunk, one of the first things to find out is which services the program uses. Programs such as ktrace and truss can show this type of information but DTrace greatly extends this ability. We will now find out which services the ls program requires to execute, as well as which ones are used most often.

```
1: pid$target:::entry
2: {
3:   @[probefunc] = count();
4: }
```

The script here is written in the D language and should be relatively easy to decipher for anyone familiar with C. The script contains a single function, which counts the entry into any call that the ls program makes. Where a C programmer might find a function name and argument list, we instead see what is called a *predicate*. The predicate is a way of selecting the probes for which DTrace will record data. The predicate on line 1 selects the entry into any call for the associated process. When the calls.d script is executed with dtrace, its pid\$ variable is replaced with the process ID of the program that is given after the -c command-line argument:

```
> sudo dtrace -s calls.d -c ls
dtrace: script 'calls.d' matched 5906 probes
```

[output of ls command removed for brevity]

```
dtrace: pid 7008 has exited
```

strcoll	148
strcoll_l	148
__error	326
free	353
strcmp	381
wcwidth	614
_none_mbrtowc	662
mbrtowc	662
putchar	705
pthread_getspecific	1424

```
>
```

DTrace also allows the tracing of live processes by replacing -c with -p and the program name with a live process ID. Here we show the abbreviated output from the execution of ls under DTrace. Only the last several lines, those with high entry counts, are shown. From this snapshot we can see that ls does a lot of work with the string functions strcoll and strcmp, and if we were trying to optimize the program we might look first at where these functions were called.

With thousands of predefined probe points, and the ability to create probes dynamically for user processes, it's obvious that DTrace is the most powerful code-splunking tool developed in the past decade.

CONTINUING CHALLENGES

In reviewing the tools mentioned here—as well as those that are not—a few challenges remain apparent. The first

challenge is the move by some developers away from a tool-based approach to an all-in-one approach.

A tool-based approach can best be understood by looking at the programs available on any Unix-like system. The use of several programs, mixed and matched, to complete a task has several obvious benefits that are well documented by others. When working with large code bases, the downfalls of an all-in-one approach, such as an IDE, become a bit clearer. A system such as the FreeBSD kernel is already several hundred megabytes of text. Processing that code base with tools such as Cscope and global in order to make it more easily navigable generates a further 175 MB of data. Although 175 MB of data may be small in comparison with the memory of the average desktop or laptop, which routinely come with 2 GB to 4 GB of RAM, storing all that state in memory while processing leads to lower performance in whatever tool is being used. The pipeline processing of data, which keeps in-memory data small, improves the responsiveness of the tools involved. Loading the FreeBSD kernel into Eclipse took quite a long time and then took up several hundred megabytes of RAM. I have seen similar results with other IDEs on other large code bases.

An even larger challenge looms for those who work on not only large but also heterogeneous code bases. Most Web sites today are a mélange of PHP or Python with C or C++ extensions, using MySQL or PostgreSQL as a database back end, all on top of an operating system written in C. It is often the case that tracking down particularly difficult problems requires crossing language barriers several times—from PHP into C++ and then into SQL, then perhaps back to C or C++. Thus far, I have seen no evidence of tools that understand how to analyze these cross-language interactions.

The area that deserves the most attention is visualization. Of all the tools reviewed, only Doxygen generates interesting and usable visual output. The other tools have a very narrow, code-based focus in which the user is usually looking at only a small part of the system being investigated.

Working in this way is a bit like trying to understand the United States by staring at a street sign in New York City. The ability to look at a high-level representation of the underlying system without the fine details would be perhaps the best tool for the code spelunker. Being able to think of software as a map that can be navigated in different ways—for example, by class relations and call graphs—would make code spelunkers far more productive.

One last area that has not been covered is the network. Network spelunking, the ability to understand an application based on its network traffic, is still in a very nascent state, with tools such as Wireshark being the state of the art. Many applications are already running online, and being able to understand and work with them at the network level is very important. **Q**

REFERENCES

1. Cscope man page; http://cscope.sourceforge.net/cscope_man_page.html.
2. Doxygen Web site; <http://www.stack.nl/~dimitri/doxygen/>.
3. GNU global source-code tag system (Apr. 21, 2008). Tama Communications Corporation; <http://tamacom.com>.
4. gprof; <http://www.gnu.org/manual/gprof-2.9.1/gprof.html>.
5. Graphviz Web site; <http://www.graphviz.org/>.
6. Jones, T., Tauferner, A., Inglett T. 2007. HPC system call usage trends. Linux Clusters Institute; http://www.linuxclustersinstitute.org/conferences/archive/2007/PDF/jones_21421.pdf.
7. ktrace: standard tool on open-source operating systems.
8. Neville-Neil, G.V. 2003. Code spelunking: Exploring cavernous code bases. *ACM Queue* 1(6): 42-48. <http://doi.acm.org/10.1145/945131.945136>.
9. Sun Microsystems. 2005. How to Use DTrace; <http://www.sun.com/software/solaris/howtoguides/dtrace-howto.jsp>.
10. Sun Microsystems. 2005. Solaris Dynamic Tracing Guide; <http://docs.sun.com/app/docs/doc/817-6223>.
11. Truss is available on Solaris.
12. Wright, G.R., Stevens, W.R. 1995. *TCP/IP Illustrated, Vol. 2: The Implementation*. Boston, MA: Addison-Wesley Professional.

LOVE IT, HATE IT? LET US KNOW

feedback@acmqueue.com or www.acmqueue.com/forums

GEORGE V. NEVILLE-NEIL (gnn@acm.org) is a columnist for *Communications of the ACM* and *ACM Queue*, as well as a member of the *Queue* Editorial Board. He works on networking and operating system code and teaches courses on various subjects related to programming.

© 2008 ACM 1542-7730/08/1100 \$5.00

This article appeared in print in the October 2008 issue of Communications of the ACM.