Dynamic Analysis for Locating Product Features in Ada Code

Laura White[1], Norman Wilde
Department of Computer Science
University of West Florida

Introduction

One of the most difficult tasks facing any Software Engineer is that of updating an unfamiliar program. Despite years of development of Software Engineering techniques, the job still somehow comes down to long hours of paging through code in an editor or debugger.

Practicing Software Engineers tell us that reliably locating *product features* in unfamiliar programs can be one of their biggest headaches. Our research in feature location was motivated several years ago by a visit to a site that maintained the software of a telephone PBX switch. To make changes correctly, maintainers needed to find all the code fragments involved in end-user features such as "call forwarding" or "call waiting". In a large and frequently-modified system such as this one, the code for a feature is often not contiguous or located in obvious places.

In response to this need, we developed a simple but elegant dynamic analysis method called *Software Reconnaissance* to aid in locating such product features. The method has been tried out and refined with support from the Software Engineering Research Center (SERC)[2]. Case studies at several sites indicate that the method seems to be useful [WILD.96]. For C software, students at the University of West Florida have developed a public-domain Reconnaissance tool called Recon2 [RECON] and Telcordia, a SERC affiliate, has included a Reconnaissance facility in their χSUDS testing toolkit for C and C++ [TELCO].

Several SERC industrial affiliates also expressed an interest in applying Reconnaissance to Ada code, and particularly to Ada embedded systems. However such code has characteristics that can make dynamic analysis difficult. This paper describes the progress of our current efforts to address the issues involved in extending Software Reconnaissance to Ada and to develop a public-domain Ada Software Reconnaissance tool.

The Software Reconnaissance Method

Software Reconnaissance uses test cases as probes to locate code for a particular product feature. The program is first instrumented in much the same way as programs are instrumented to determine test coverage. Then it is run with a few test cases that exhibit the desired feature and with a few others that do not. The executions of the instrumented code produces trace files showing which code components (usually branches or basic blocks) were used in each test.

The Software Reconnaissance method is described formally in [WILD.95]. Any one of several different heuristics described there may then be used to locate the components for a particular feature. The simplest, and often the best, is simply to take the set of *marker* components for the feature, defined as those components that are executed in at least one test case that exhibits the feature and not executed in any test case without the feature (Figure 1).

---

[1] Address correspondence to: Laura White, Department of Computer Science, University of West Florida, 11000 University Parkway, Pensacola, FL 32514, tel. 850-474-2542, fax. 850-857-6056, email: lwhite@uwf.edu.
[2] SERC is an NSF Industry/University Cooperative Research Center, with sites at Purdue University, the University of Florida, the University of Oregon, and the University of West Virginia. SERC, and this research, has been supported by NSF grant EEC-9418762 and others, and by SERC industrial affiliates.

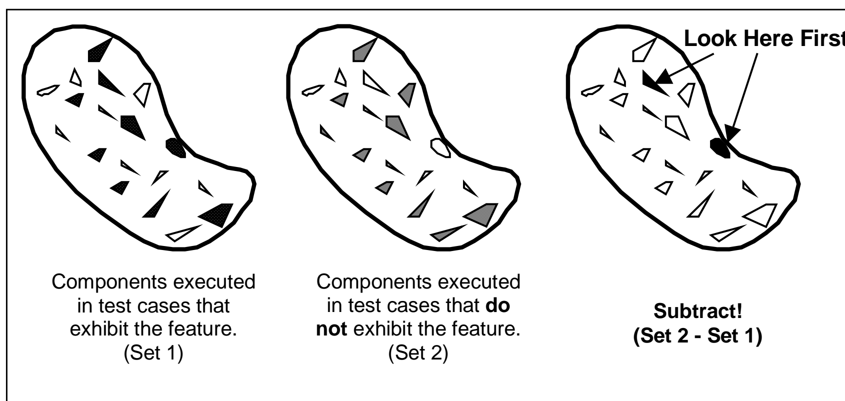

Figure 1
Locating "Marker" Components for a Feature

The techniques in the literature that most closely resemble Software Reconnaissance are methods for debugging. A common witticism in Software Engineering is that a program fault may be viewed of as simply an undesired "feature". In desperation, experienced debuggers sometimes resort to a technique called "dump and diff" that involves inserting print statements and running a file difference utility to compare different runs. The first published suggestion that trace information could be used systematically for fault localization seems to have been by Collofello and Cousin [COLO.87] who suggested tracing "decision-to-decision paths" and comparing runs that exhibit a failure with runs that do not.

Experience from C Trials of Software Reconnaissance

The Software Reconnaissance method of locating features has been tried in a series of case studies, most of them involving C code. The first trials used University or public domain code, as a "sanity check" to see if the method made sense. One of these trials was a protocol study in which experienced C programmers were asked to "think aloud" as they used the method to solve a problem. Such trials are useful to evaluate the usability of a technique by its prospective users [WILD.95].

Additional trials used small software systems (from about 10 KLOC to 50 KLOC, raw line counts) from SERC affiliates and other companies [WILD.96]. In these trials a Software Engineer familiar with the system was asked to locate features that he had recently needed to study, or that could be important for a new maintainer. These trials gave us increased confidence in the value of the method, as well as insight into the kinds of test cases that are most effective.

The results of these trials have generally been favorable. Software Reconnaissance does not locate all the code a maintainer needs to study to understand a feature, but it usually finds very good starting points for code exploration in an unfamiliar system. It focuses attention on a relatively small part of the code and can often provide insights that surprise programmers, even if they thought they already understood the program. "I didn't know it was doing that!" is a comment we have heard several times in our studies.

Software Reconnaissance for Ada

These experience with C would seem to show that an Ada Software Reconnaissance tool would be desirable. However Ada systems have characteristics that make the production of traces somewhat problematic:

- Ada systems are often embedded, which makes it more difficult to produce and store trace files.

- Ada systems often use multi-tasking, so naive methods of tracing, such as direct writes to the file system, may fail due to task contention.

- Ada systems often have real-time constraints, which may limit the amount of tracing which can be performed.

An Ada tool for Software Reconnaissance will thus require careful design. Our approach to developing such a tool has included a survey to identify current instrumentation practices for embedded systems, design of an architecture to accommodate multi-tasking as well as the results from the survey, and timing trials to estimate the performance impact of instrumentation.

Instrumentation Practices for Embedded Systems

There is a considerable literature on the problems of tracing or monitoring embedded, real-time and distributed systems. An extensive survey of the debugging problems is given by McDowell and Helmbold [MCDO.89]. Schutz also gives an extensive review of the problems, emphasizes the problem of observability, and comments on the difficulties of using conventional debuggers [SCHU.94]. He also describes several monitoring systems.

Many other authors also discuss the theory of monitoring methods and/or describe specific monitoring or debugging systems, for example [MARI.90, TSAI.90, SCHM.94, SIDE.94]. Yan, Sarukkai and Mehra describe some experience with a very extensive set of monitoring and visualization tools from NASA's Ames Research Center [YAN.95]. Hollingsworth, Miller, Goncalves, Naim, Xu and Zheng describe a technique they call "dynamic program instrumentation" in which the running program is modified periodically to collect information about its execution [HOLL.97]. Waheed, Rover and Hollingsworth give a detailed analysis of some of the design alternatives in monitoring distributed systems, such as the trade-off between sending event data immediately or batching it for efficiency [WAHE.98]. Heath and Etheridge describe methods for graphically displaying the results of monitoring parallel systems, a topic that is very important for effective program comprehension [HEAT.91].

Our group took a different tack and surveyed practicing Software Engineers, mostly from SERC affiliate companies, to see what instrumentation methods were actually being used. The survey covered 10 embedded systems ranging from 2 KLOC to 3 MLOC in size, and from mid 70's to late 90's in development epoch. Only two of the systems were written in Ada, but they were all embedded and had real-time constraints. A more complete report of the survey is given in [WILD.99].

In general, we found that instrumentation was often ad-hoc, and was only added after severe problems were encountered in integration or testing. Few engineers were aware of or had available to them instrumentation systems of the sophistication described in the studies mentioned previously.

The survey clearly showed that extracting trace data from an embedded system is difficult; there can be no general solution because the hardware differs so greatly from system to system. Software engineers exercised great ingenuity to extract trace information. Methods described to us used everything from light emitting diodes to logic analyzers to conventional flat files depending on the available output devices.

We also found that our interviewees wanted to go considerably beyond the needs of Software Reconnaissance. Reconnaissance only requires a trace that shows if a particular component (e.g. subprogram or basic block) was executed or not. However for understanding embedded systems it is often essential to know the number of times the component was executed, the actual sequence of execution, or possibly even the time at which execution occurred.

Based on these interviews, we defined three primary circumstances in which a Software Engineer could make use of Ada Software Reconnaissance:

- Testing in a development environment where timing constraints are not stringent. Reconnaissance could help acquire an understanding of the features of a legacy software system.

- Testing using a monitor connected to the actual target system. Traces would be sent to the monitor by a variety of methods (socket, communications port, logic analyzer, etc.) depending on the target system hardware available. Reconnaissance could be used to analyze and evaluate run-time behavior of the target system.

- Remote monitoring of a deployed system for troubleshooting purposes, making use of instrumentation previously installed. Reconnaissance could be used to highlight changes in behavior as the system executes [LUKO.00].

There are thus three decisions to be made for each tracing session:

- What events are instrumented (e. g. subprogram entries/returns, basic block executions, task activation and rendezvous?)

- What information is captured at each event (e.g. just whether the event occurred, a count of occurrences, the sequences of occurrences, or the times of occurrences)

- How trace monitoring is performed (e. g. via a socket, a logic analyzer, a communications port, etc.)

The Ada Tracing Architecture

It is clear that no one solution will meet all these needs. Instead, we have defined a set of interchangeable components that can be used to instrument an Ada program. The general architecture is shown in Figure 2.

The instrumented program actually contains the `trace_manager_interface` task embedded in the declarative part of the main program unit. Trace data is stored within the `trace_array_pkg` which is made visible using a `with` context clause. Additional instrumentation which calls the `trace_it` procedure is used to record each trace event in the `trace_array`, which is declared using the `Atomic_Components` pragma to prevent contention.

The `trace_manager_interface` task runs concurrently with the users program units. At the `trace_now` entry call, `copy_traces` copies the `trace_array` to `tmi_array` during the rendezvous, after which `send_traces` and the users tasks continue concurrently.

The `send_traces` procedure would have to be written specially for each environment, since the mechanism for trace monitoring is hardware-dependent. The trace manager may be as simple as a program to write trace data to a file, or it could interface to a more sophisticated analysis or display tool.

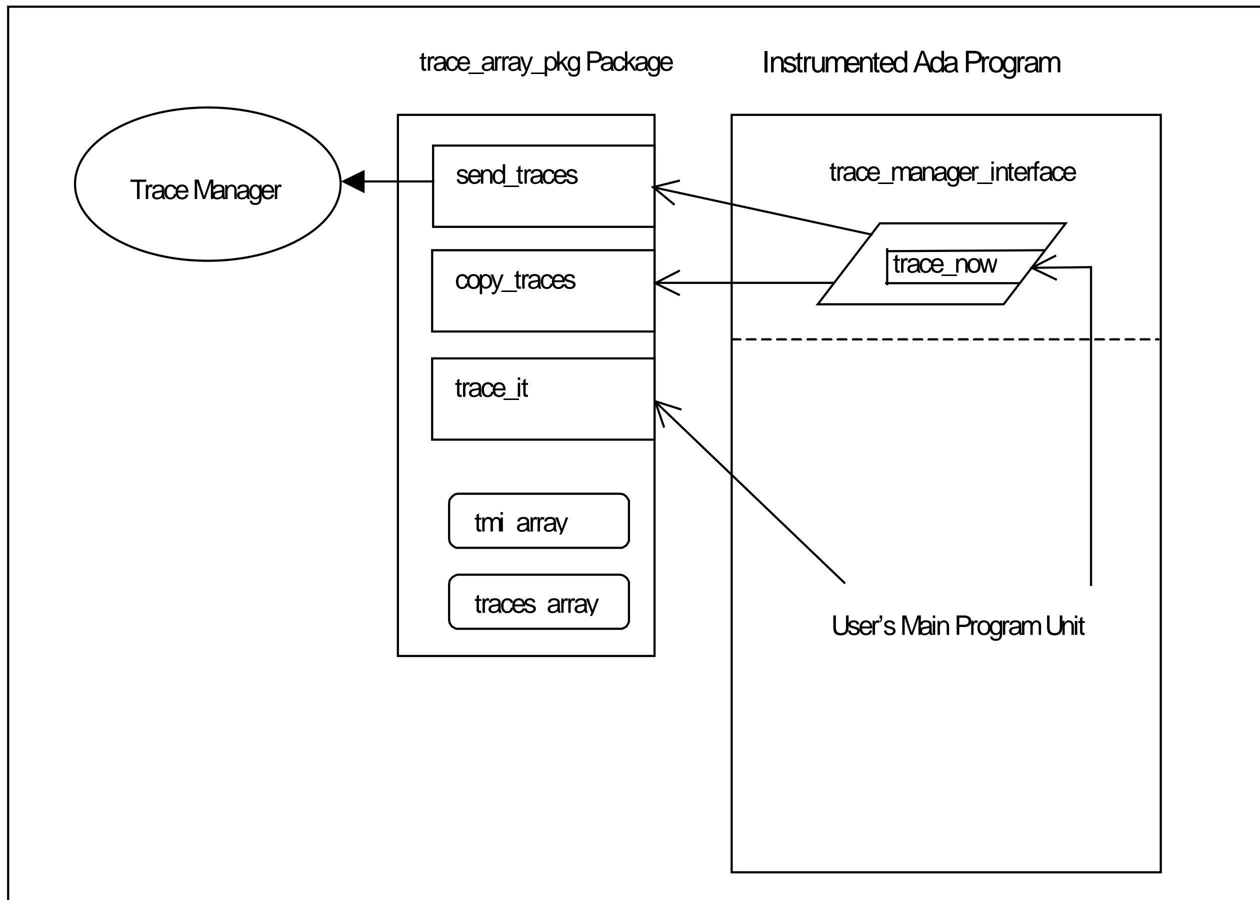

Figure 2
General Architecture for Ada Tracing

Performance Trials for Ada Tracing

All methods of producing traces may be to some extent *intrusive*, that is, they may change the behavior of the system being studied. For real-time software, intrusion is a key issue, since the behavior of the instrumented system may not be the same as that of the original, and the conclusions of the analysis may thus become invalid.

Intrusion is less important in Software Reconnaissance than in other kinds of analysis since the main goal is simply to locate code that implements a feature; the analysis of the feature is performed by a human studying the identified code. Unless system behavior is so radically changed by the instrumentation that the feature is not performed, we should still locate it. If necessary, instrumentation may be disabled for some of the tightest loops.

However it would certainly be useful to have a rough idea of the performance impact of the different kinds of trace instrumentation, at least to give Software Engineers a reference point in deciding what strategies they might try. Accordingly we are carrying out a performance study using several small, hand-instrumented Ada programs that are publicly available.

The problems with basing decisions on any kind of time benchmarking are well known. For example [CLAP.86] describes the difficulties in developing a set of benchmarks for different features of Ada. The programs and the environment used for the benchmark may not be representative of your system so "your mileage may vary".

(The full paper will contain a more detailed description of the performance trials and their results)

Conclusions

We have described the Software Reconnaissance method for locating product features in software systems as well as our experience so far with this method as applied to C programs. We believe that Reconnaissance can also be very useful to developers and maintainers of Ada systems.

To create an Ada Reconnaissance tool we have surveyed practicing Software Engineers to see how they use instrumentation with embedded software today. Based on these results, and on our earlier experience with C, we have attempted to define an Ada tracing architecture that will be flexible enough to be usable in the kinds of situations that arise in practice.

We think we are well on the way to providing a public domain Ada Reconnaissance tool that will be made available to the Ada community.

Acknowledgements

We would like to thank very much all the students who have participated in different facets of this project over the last few years and particularly Gus Lorberg and Andre Wacaster, for their assistance with the architecture and Barry Coker and Carlos Trani for their work on the performance trials.

References

[CLAP.86] Clapp, R.; Duchesneau, L.; Volz, R.; Mudge, T., and Schultze, T., "Toward Real-Time Performance Benchmarks For Ada", Communications of the ACM, Volume 29, Number 8, August 1986, pp. 760-778.

[COLO.87] Collofello, James, and Cousin, Larry, "Towards automatic software fault location through decision-to-decision path analysis", *Proceedings National Computer Conference 1987*, pp. 539-544.

[HEAT.91] M. Heath and J. Etheridge, "Visualizing the Performance of Parallel Programs", *IEEE Software*, Vol. 8, No. 5, (September 1991), pp. 29 - 39.

[HOLL.97] J. K. Hollingsworth, B. P. Miller, M. J. R. Goncalves, O. Naim, Z. Xu, L. Zheng, "MDL: A Language and Compiler for Dynamic Program Instrumentation", *PACT'97 - 1997 International Conference on Parallel Architectures and Compilation Techniques*, November 11-15, 1997, San Francisco

[LUKO.00] K. Lukoit, N. Wilde, S. Stowell, T. Hennessey, "TraceGraph: Immediate Visual Location of Software Features", *Proc. International Conference on Software Maintenance - ICSM 2000*, IEEE Computer Society, Los Alamitos, CA, October 2000, pp. 33 - 39.

[MARI.90] D. Marinescu, J. Lumpp, T. Casavant, and H. J. Siegel, "Models for Monitoring and Debugging Tools for Parallel and Distributed Software", *Journal of Parallel and Distributed Computing*, Vol. 9, No. 2, (June 1990), pp. 171 - 183.

[MCDO.89] C. McDowell and D. Helmbold, "Debugging Concurrent Programs", *ACM Computing Surveys*, Vol. 21, No. 4, (December 1989), pp. 593 - 622.

[RECON] "Recon Tool for C Programmers", http://www.cs.uwf.edu/~recon/.

[SCHM.94] U. Schmid, "Monitoring Distributed Real-Time Systems", *Real-Time Systems*, Vol. 7, No. 1, (July 1994), pp. 33 - 56.

[SCHU.94] W. Schutz, "Fundamental Issues in Testing Distributed Real-Time Systems", *Real-Time Systems*, Vol. 7, No. 2, (September 1994), pp. 129 - 157.

[SIDE.94] R. S. Side and G. C. Shoja, "A Debugger for Distributed Programs", *Software - Practice and Experience*, Vol. 24, No. 5, (May 1994), pp. 507 - 525.

[TELCO] "Telcordia Software Visualization and Analysis Toolsuite", http://xsuds.argreenhouse.com/.

[TSAI.90] J. Tsai, K. Fang, H. Chen, and Y. Bi, "A Noninterference Monitoring and Replay Mechanism for Real-Time Software Testing and Debugging", *IEEE Trans. on Software Engineering*, Vol. 16, No. 8, (August 1990), pp. 897 - 916.

[WAHE.98] A. Waheed, D. Rover, J. Hollingsworth, "Modeling and Evaluating Design Alternatives for an On-Line Instrumentation System: A Case Study", *IEEE Trans. on Software Engineering*, Vol. 24, No. 6, (June 1998), pp. 451 - 470.

[WILD.95] N. Wilde and M. Scully, "Software Reconnaissance: Mapping Program Features to Code", *Journal of Software Maintenance: Research and Practice*, Vol. 7, (1995) pp. 49-62.

[WILD.96] N. Wilde and C. Casey, "Early Field Experience with the Software Reconnaissance Technique for Program Comprehension", *Proc. International Conference on Software Maintenance - 1996*, IEEE Computer Society, Los Alamitos, CA, November 1996, pp. 312 - 318.

[WILD.99] Norman Wilde and Dean Knudson, "Understanding Embedded Software Through Instrumentation: Preliminary Results from a Survey of Techniques", Report SERC-TR-85-F, Software Engineering Research Center, Purdue University, 1398 Dept. of Computer Science, West Lafayette, IN 47906, February, 1999. Available at http://www.cs.uwf.edu/~wilde/publications/TecRpt85F_ExSum.html.

[YAN.95] J. Yan, S. Sarukkai, P. Mehra, "Performance Measurement, Visualization and Modeling of Parallel and Distributed Programs using the AIMS Toolkit", *Software - Practice and Experience*, Vol. 25, No. 4, (April 1995), pp. 429 - 461.