

# Statement of Purpose

Philip Guo (pgbovine@mit.edu)

I want to pursue a Ph.D. in Computer Science, and my career aspiration is to become a professor. My research interests include program analysis, software engineering, operating systems, and human-computer interaction (HCI).

My motivation for pursuing research related to program analysis stems from learning about the challenges of software development during 3 summers spent as a software engineering intern (1 at Codehost, a start-up, and 2 at Teradyne, a large corporation). Although I found the actual development work to be educational, what really sparked my interest in research was experiencing first-hand the challenges that software engineers face daily.

I aspire to create tools that can increase programmer productivity and improve software quality and security. As a first step towards this goal, for the past 2 years I have been working with Prof. Michael Ernst in the MIT Program Analysis Group. The research problem that my work addresses is the difficulty of building dynamic analysis tools that work reliably on large C and C++ programs, due to the presence of complex constructs and the lack of memory and type safety. I have developed a framework named Fjalar that enables the construction of scalable and robust dynamic analyses for C and C++ programs. It operates on unmodified x86 binaries of programs as large as 1 million lines of code, and it provides rich language- and machine-level information at run time. I have used Fjalar to build two dynamic analysis tools. The first, Kvasir, records a detailed trace of data structure contents during execution. The second, DynComp, performs dynamic type inference to find sets of related variables whose values interact at run time. My Master's thesis is based on Fjalar and the tools built upon it. I also have experience with static analysis. I am currently working on type inference of unit types (e.g., physics units) for C programs.

My research has been directly applied to several projects and produced visible results. Prof. Ernst has used Kvasir to extend the Daikon dynamic invariant detector to analyze large C and C++ programs, and he has used DynComp to improve Daikon's performance and the quality of inferred invariants. We collaborated with Prof. Martin Rinard to use Daikon together with Kvasir and DynComp to generate constraints for performing run time data structure repair. This research was funded by DARPA, and they recently evaluated its effectiveness by hiring an external company to try to attack and crash a specific program: the Freeciv multiplayer game server. Our repair system successfully protected the server from a high percentage of their attacks. In another assessment, we were able to use this repair system to mitigate the effects of several real security flaws in the BIND DNS server. It is very encouraging to see my research applied to real-world problems.

Taking an Operating Systems Engineering course has inspired me to explore OS-related research. In particular, I am interested in applying program analysis to make low-level software such as operating systems more reliable and secure. In addition, I have worked on 2 research projects that have stimulated my interest in user interfaces and HCI: I created a suite of 5 educational games for Palm OS devices for the Teacher Education Program (2002-2003) and an image editing tool for the Computer Graphics Group (Fall 2003).

Although I am open to a variety of research, there are several professors at Stanford whose projects are especially appealing to me: Profs. Aiken, Engler, and Lam (program analysis); Profs. Fox and Mazieres (OS and distributed systems); Prof. Klemmer (HCI). Studying some of these professors' papers for classes and for my research and reading project websites have given me a sense that Stanford's Ph.D. program is a great match for my interests.

Fjalar project: [pag.csail.mit.edu/fjalar/](http://pag.csail.mit.edu/fjalar/)

See Additional Information page for publications and additional details