

Analyzing Distributed Java Programs

Mariana Barca
Ohio State University

Distributed applications

Distributed applications are used extensively to solve problems in a variety of disciplines due to the significant benefits they provide. These benefits include software performance, reliability, interoperability and extensibility.

A distributed application is made of components that communicate with each other. Web browsers, seen today on every desktop, are a familiar example of distributed application because they connect to remote servers to retrieve web pages.

The design of a distributed application is inherently complex and raises many problems related to scalability, security, and maintenance. There has been extensive research in the past decades to find solutions for these problems. Various approaches to distributed application design have been explored, however the approach I will focus on in my presentation is program analysis.

What is Program Analysis?

Program analysis is an area that is concerned with developing techniques for analyzing programs, either source code or compiled code. The results of these analyses are designed to help developers in their daily tasks. They are incorporated in tools used for program compiling, understanding, debugging, testing, reverse engineering and software visualization.

My work

My work explores the relatively new area of creating analysis techniques for distributed applications. This area adapts the existing techniques of program analysis to the challenges of distribution. Useful information about distributed applications is found by analyzing their code. My research focuses on distributed Java programs built with the help of Remote Method Invocation (RMI) mechanism. I consider the particular technique of points-to analysis for such applications. Points-to analysis determines the set of objects pointed to by a variable or a object field in a program. My presentation will show how the points-to information resulted from this analysis may reveal unexpected data correlations between different parts of a distributed application.