

Name: Sang Boem Lim

Advisor: Geoffrey Fox

PhD Title: Platforms for HPJava: Efficient Runtime Support for Scalable Programming with Java

Expected Defense Date: Summer 2002

Remained courses:

- CAP 5600 Artificial Intelligence
- CIS 5920r Colloquium

Portfolio Status: Jun. 2001

Title of Survey Paper: Environments for Parallel Computing with Current Object-Oriented Languages

Expected date of Completion of Survey Paper: Jun. 2001

Expected Date of Oral exam: Aug. 2001

Environments for Parallel Computing with Current Object-Oriented Languages

By Sang Lim
slim@csit.fsu.edu

Abstract

The report surveys contemporary systems and environments that support parallel computing using modern object-oriented languages. It concentrates on Java as characteristic example of this kind of language, although we also briefly discuss related languages like C++ and C#. The emphasis is on platform independent software for scalable computation, founded on modern Internet technologies. We will be especially interested in technologies developed in recent years by participants of bodies like the Java Grande Forum and the Global Grid Forum.

The paper will broadly review ongoing activities in this area, and then discuss in detail particular environment under development at FSU called *HPJava*. This survey is intended as foundation for further research and development on the *HPJava* environment. In particular the proposed work will concentrate on issues related to the development of efficient *run time support software* for parallel languages extending an underlying object-oriented language. So the survey emphasizes aspects most relevant to this topic — aspects like high performance interprocessor communication and dynamic discovery of available compute resources.

The first part of the report reviews related researches and systems. We start with discussion of communication approaches relevant to parallel computing that includes Java RMI and Message Passing libraries for Java. PVM and MPI are discussed as examples of message passing libraries. We discuss *MPJ*, an effort by members of the Java Grande Forum to define a consensus MPI-like interface for Java. We review general parallel computing environments for Java including JavaParty and Javelin. We also briefly review more contemporary Java-specific distributed computing technology like JINI, JMS, EJB, and JXTA and discuss how they relate to our high performance computing infrastructures.

The report continues with discussion of ongoing work within our research group on the *HPJava* system. The *HPJava* language design and its features are discussed first. Associated runtime requirements, including distributed data descriptors, high level communications libraries for distributed arrays, and dynamic processor discovery are considered. A C++ library that provides some of these features is described. Issues about providing these functionalities in pure Java based programming environment are discussed. As part of the *HPJava* environment, *mpiJava*, an object-oriented Java interface to MPI, has been developed; design and implementation issues discussed.

The HPJava system is compared and contrasted with related systems. Java as language support parallel and scientific computing is compared with C++ and C#.

Provisional table of Contents

Chapter 1 Introduction

Chapter 2 Survey Java for Parallel Programming

2.1 Communication in Parallel Programming

2.1.1 Java RMI

2.1.2 PVM and MPI

2.1.3 MPJ

2.2 Java-specific Distributed Computing

2.2.1 JINI, JMS, EJB, and JXTA

Chapter 3 Review of the HPJava System

3.1 Introduction to HPJava

3.2 Runtime Support

3.2.1 Data Descriptors

3.2.2 Communication Libraries for Distributed Arrays

3.2.3 Processor Discovery: JINI and beyond?

3.3 Message Passing for HPJava

Chapter 4 Comparison with Related Systems

4.1 Related Systems

4.2 Related Languages

4.2.1 Java vs C++

4.2.2 Java vs C#