Process Management in a Distributed Operating System

Sape J. Mullender
Centre for Mathematics & Computer Science
Amsterdam
and
Computer Laboratory
Cambridge University

As part of designing and building the Amoeba distributed operating system, we have come up with a simple set of mechanisms for process management that allows downloading, process migration, checkpointing, remote debugging and emulation of alien operating system interfaces.

The basic process management facilities are realized by the Amoeba Kernel and can be augmented by user-space services: Debug Service, Load-Balancing Service, Unix-Emulation Service, Checkpoint Service, etc.

The Amoeba Kernel can produce a representation of the state of a process which can be given to another Kernel where it is accepted for continued execution. This state consists of the memory contents in the form of a collection of segments, and a *Process Descriptor* which contains the additional state, program counters, stack pointers, system call state, etc.

Careful separation of mechanism and policy has resulted in a compact set of Kernel operations for process creation and management. A collection of user-space services provides process management policies and a simple interface for application programs.

In this paper we shall describe the mechanisms as they are being implemented in the Amoeba Distributed System at the Centre for Mathematics and Computer Science in Amsterdam. We believe that the mechanisms described here can also apply to other distributed systems.

CR Categories: D.4, C.2.4, D.2.5.
Key Words & Phrases: Distributed operating system, process management, migration, Amoeba.

1. Introduction

Our goal in designing the process management primitives described in this paper was to provide mechanisms that can do what process management primitives in existing general-purpose operating systems can do and much more. The added functionality has to do with the properties of the kinds of distributed systems we are interested in: personal workstations, shared server machines and *guest systems*, connected by a fast local-area network.