

High Speed Communication for Simulation of Large VHDL Models *

*David Charley Debra Hensgen Timothy J. McBrayer Philip A. Wilsey
Manish Ankola*

Department of Electrical and Computer Engineering
University of Cincinnati, Cincinnati, OH. 45221-0030
{dcharley, dhensgen, tmcbraye, paw, mankola} @thor.ece.uc.edu

October 21, 1992

Large VHDL models can be verified using simulation techniques. Simulations of this sort require extensive computing power and can run for long periods of time. Few machines are available which can execute large simulations in a reasonable time period and they are quite expensive.

A worthwhile environment for simulating extremely large VHDL models is to use an existing network of workstations. The Quest project at the University of Cincinnati was originally designed to run large simulations on a specialized hardware box and is currently being ported to run on a network of UNIX workstations.

One of the limiting factors in a distributed network environment is communication between simulation objects (which represent VHDL processes in this case). Simulation objects communicate by transferring messages. Objects which reside on a common machine can take advantage of high speed shared memory, while objects residing on different machines rely on network communications.

A communication system has been developed to run on a network of SGI or SUN workstations at Wright Patterson Air Force base. Workstations contain between one and four processors, and are connected with Ethernet. The communication system has been designed to take advantage of both shared memory (within a workstation) and Ethernet (between workstations) to facilitate the transfer of information between simulation objects. The system has also been designed to take advantage of SCRAMNet (a ring of memory expansion cards, connected with fiber optic links to emulate shared memory between machines) networks which may exist between selected workstations to speed up inter-machine communications.

The simulator does not require ordered delivery of messages and to date our system does not address this issue (messages are not guaranteed to be ordered). This allows the communication system to be fast. We plan to study the impact of ordering messages based on simulation time (we expect the simulation time to improve, even though the message passing system will be slower).

Another interesting aspect of the design is that migration of simulation objects will be implemented. Thus as computing power is needed, long running simulations can temporarily move off selected workstations. As computing power becomes available, a simulation can take advantage of the free resources (by migrating simulation objects on to different workstations) and reduce the simulation execution time.

*This work is partially supported by the Air Force under order number 054458 and also by the Defense Advanced Research Projects Agency under order number 7056, monitored by the Federal Bureau of Investigation under contract number J-FBI-89-094.

**High Speed Communication
for Simulation of Large VHDL Models**

October 21, 1992

David Charley
University of Cincinnati
Department of Electrical and Computer Engineering
Cincinnati, OH 45221-0030
(513) 556-2499 or dcharley.@thor.uc.ece.edu

Additional authors

Debra A. Hensgen
Timothy J. McBrayer
Philip A. Wilsey
Manish Ankola

D. Charley

Agenda

Background

Benefits

Details

Optimizations

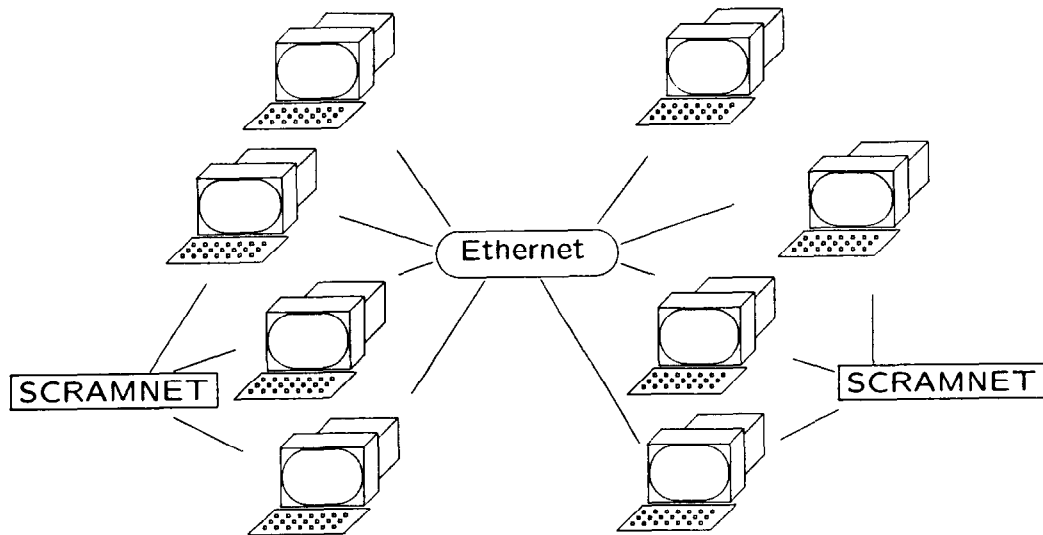
Summary

D. Charley

Background

U.C. Simulator for VHDL simulations

Hardware



D. Charley

Benefits

Communication System

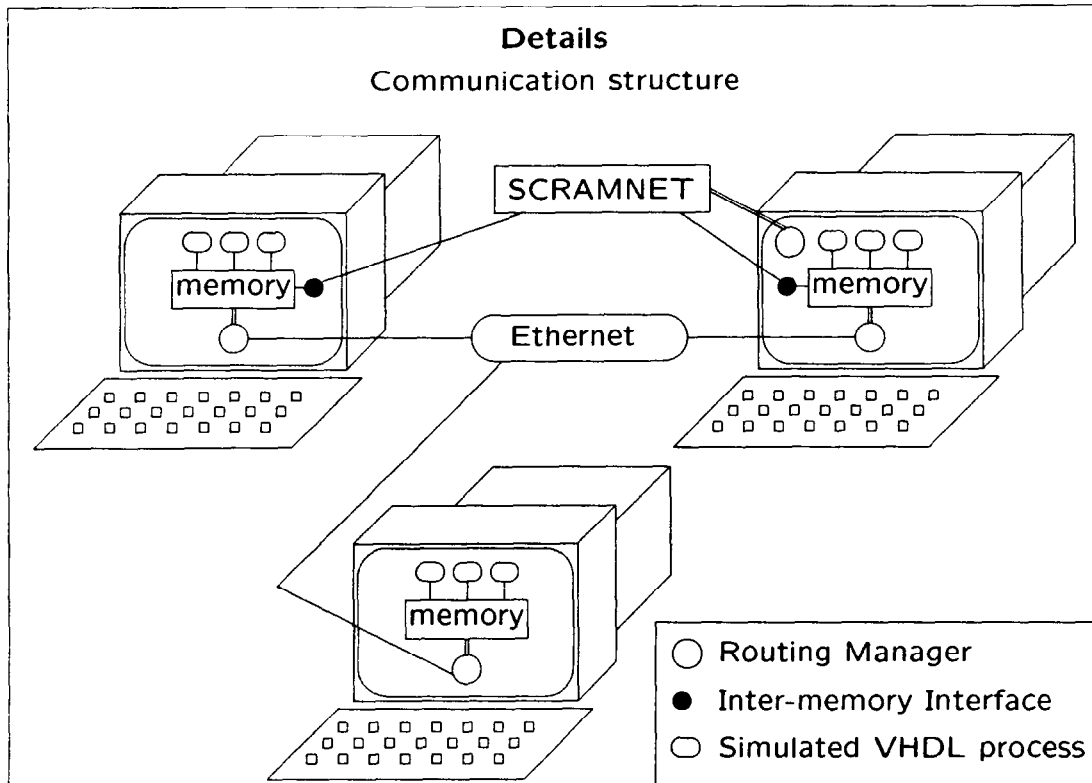
Use existing resources

Hide architecture details

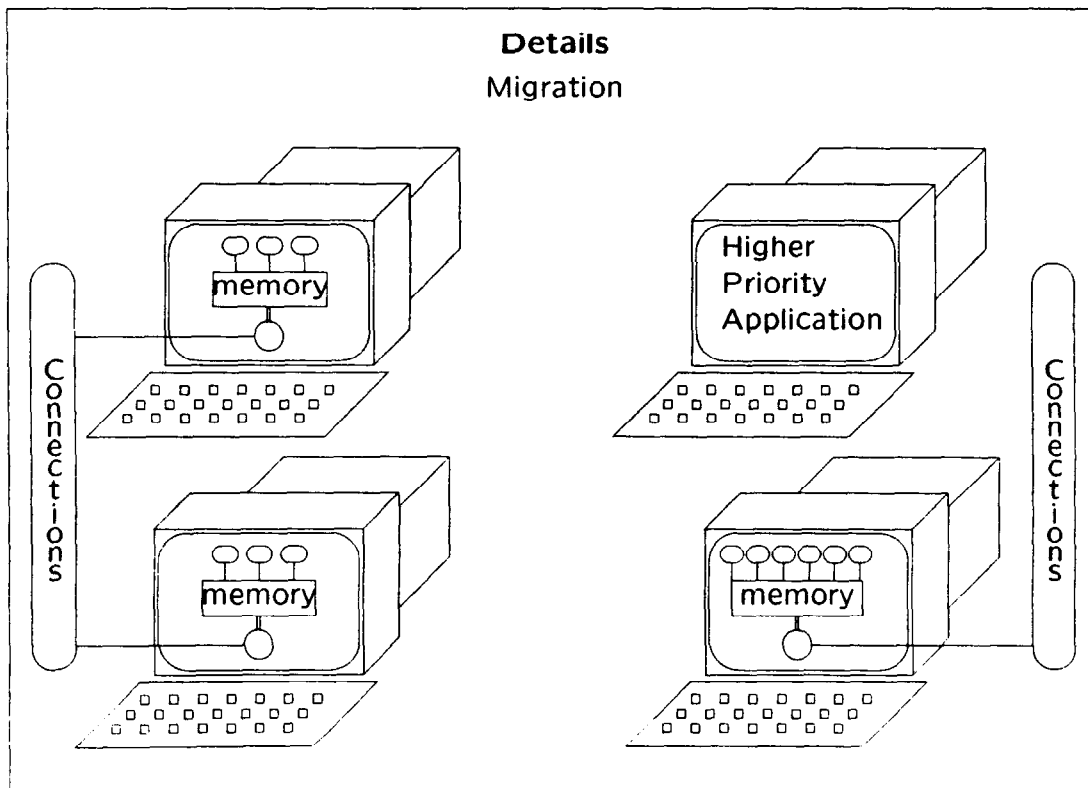
Migration Utility

Priority usage of resources

D. Charley



D. Charley



D. Charley

Optimizations

Ordered Linked lists

Route selection (Scramnet vs. Ethernet)

Fault Tolerance

Load balancing

D. Charley

Summary

Take advantage of idle workstations

Distribute VHDL simulations across many workstations

D. Charley