

Ateneo de Manila University

Building the AGILA HPCS



Ateneo High Performance Computing Group

1st Semester 2001-2002

<http://www.math.admu.edu.ph/ahpc/>

william.s.yu@ieee.org

Section I

Introduction

Is It Primetime for Parallel Computing?

- ★ increase in performance of mass-market commodity microprocessors
- ★ declining costs of high speed network hardware
- ★ scientific computing applications that can benefit from clustering:
 - ★ Artificial Intelligence Models
 - ★ Weather and Climate Modelling
 - ★ Computational Fluid Dynamics
- ★ commercial applications that can benefit from clustering:
 - ★ Email Systems
 - ★ Proxy Systems
 - ★ Web Services

What Makes an Application Parallelizable?

- ★ Is the problem inherently parallel?
 - ★ Cellular Automata
 - ★ Neural Networks
 - ★ Genetic Algorithms
- ★ Is the Problem Data-Parallel?
- ★ What is the granularity of the problem?
 - ★ Fine Grain - no way!
 - ★ Medium Grain
 - ★ Coarse Grain

Amdahl's Law

$$\frac{R(P)}{R(1)} = \frac{(T_s + T_p)}{(T_s + (P * T_{is}) + (\frac{T_p}{P}) + T_{ip})} \quad (1)$$

- ★ Amdahl's Law strictly limits the amount of speedup
- ★ serial components
- ★ communications overhead
- ★ optimistic estimation

Nature of the Problem: Some Questions

- ★ Is it compute intensive or communications intensive?
- ★ How much memory does the application need?
- ★ Does it benefit from the presence of a local disk?
- ★ How big should the cluster be?

Section II


Design

Hardware Components for the AGILA HPCS

ITEM	Number
Node System: CPUs(Athlon 650MHz w/ 128 MB SDRAM, 6.4 GB HD, Intel EEPro100 NIC)	7
Server System: (Athlon 650MHz w/ 256MB SDRAM, 10.4 GB HD, (2) Intel EEPro100 NIC Monitor, Trident 3D AGP, 1.44 Floppy, Keyboard	1
Intel 410T Switch	1 x 24 port
APC Smart UPS 3000VA	1 (20 CPUs)
Cabling and others	1 box and 1 pack

Table 1: Hardware Components of the AGILA HPCS

AGILA HPCS



AGILA HPCS Design

AGILA HPCS Software Setup

	Software Used
Linux Distribution	Redhat Linux 6.2 (Zoot)
Linux Kernel	Linux Kernel v. 2.2.17
Compilers	GCC
Message Passing Libraries	PVM 3.4.3, MPI-MPICH 1.2.1, MPI-LAM 6.3.3b48
Scientific Libraries	SCALAPACK, FFTW, PETSc

Table 2: AGILA HPCS Software Setup

Proposed Additional 8-nodes

ITEM	COST(pesos)
Node System: CPUs(Athlon 750MHz w/ 128 MB SDRAM, 10.4 GB HD, Intel EPro100 NIC)	240,000
Total	240,000

Table 3: Proposed Additional 8-nodes

Conclusion

- ★ not all problems are parallelizable
- ★ clusters are built to solve problems
- ★ Beowulfery is not simply about putting computers together
- ★ a well planned beowulf is a good thing (TM)



Copyright © 2000-2001 by William Emmanuel S. Yu. This material may be distributed only subject to the terms and conditions set forth in the Open Content License, v1.0 or later (the latest version is presently available at <http://opencontent.org/opl.shtml>).