CS550

- Distributed Operating Systems (Advanced Operating Systems)
- Instructor: Xian-He Sun
  - Email: sun@iit.edu, Phone: (312) 567-5260
  - Office hours: 2:10pm-3:10pm Tuesday, 3:30pm-4:30pm Thursday at SB229C, or by appointment
- TA: TBA
  - Email: tba@iit.edu
  - Office hours: TBA
- Blackboard*:*
  - http://courseinfo.iit.edu
- Class Web site
  - http://www.cs.iit.edu/~sun/cs550.html

Outline

- Course information
- Key issues of distributed operating systems
- Hardware concepts
  - Multiprocessors
  - Multicomputers
  - Distributed systems
- Software concepts
  - Uniprocessor OS
  - Distributed OS
  - Network OS
  - Middleware

What This Course is About

- Understanding the *fundamental concepts* of distributed operating system, and distributed systems in general
- Learning *distributed programming* techniques
  - Multithreading, RPC, RMI, Sockets, etc.
- Understanding the *general principles* of distributed paradigms
  - MPI, JINI, NFS, Web Service, Grid, etc.

Prerequisite

- CS450 “Operating Systems”
- Familiar with
  - Programming in C/C++ or Java
  - UNIX tools and development environment
    - Command
    - Editors (vi, emacs), compilers (gcc), makefiles (GNU make)
  - Networking programming
    - Sockets
    - Multithreaded
    - RPC, Java RMI
  - Basic concepts of computer architecture

Course Materials

- Required:
  - "Distributed Systems: Principles and Paradigms" by Tannenbaum and Van Steen, Prentice Hall 2002
- Recommended:
  - "Distributed Operating Systems & Algorithms" by Randy Chow and Theodore Johnson, Addison-Wesley, 1997
- Supplemental readings

Misc. Course Details

- You are expected to attend all of the lectures and presentations
- Grading
  - written and programming assignments (35%): individual work
  - One exam (35%)
  - Final project (30%): individual or group with 2-3 students
- Use the course blackboard
  - Announcements
  - Lecture notes
  - Assignments
  - Discussion
  - …

Policies

- Collaboration
  - Encouraged for high level concepts and understanding the courses materials
  - but …..
- Cheating
  - Copying all or part of another student's homework
  - Allowing another student to copy all or part of your homework
  - Copying all or part of code found in a book, magazine, the Internet, or other resource

Any Questions?

Personal Introduction

- Research interests
  - Middleware
  - Performance Analysis and Modeling
  - Pervasive Computing
  - Scientific High Performance Computing
- Research group:
  - Scalable Computing Software Laboratory (SCS)
  - http://www.cs.iit.edu/~scs/
  - Weekly Research seminar

Distributed Computing at SCS

Many workstations are made available for graduate students

Scalable Computing Software (SCS) Lab.

Parallel Computers at SCS

Distributed Optical Testbed (Grid)

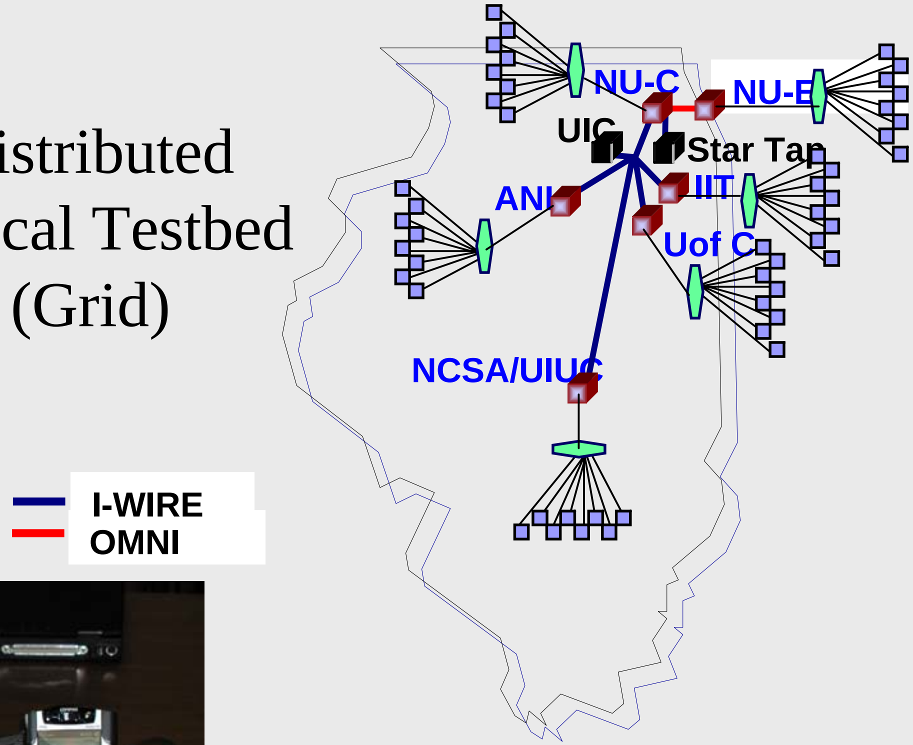

Pervasive Computing Environments at SCS

Evolution of Computing

Bigger becomes even bigger

Smaller becomes ever smaller, & connected

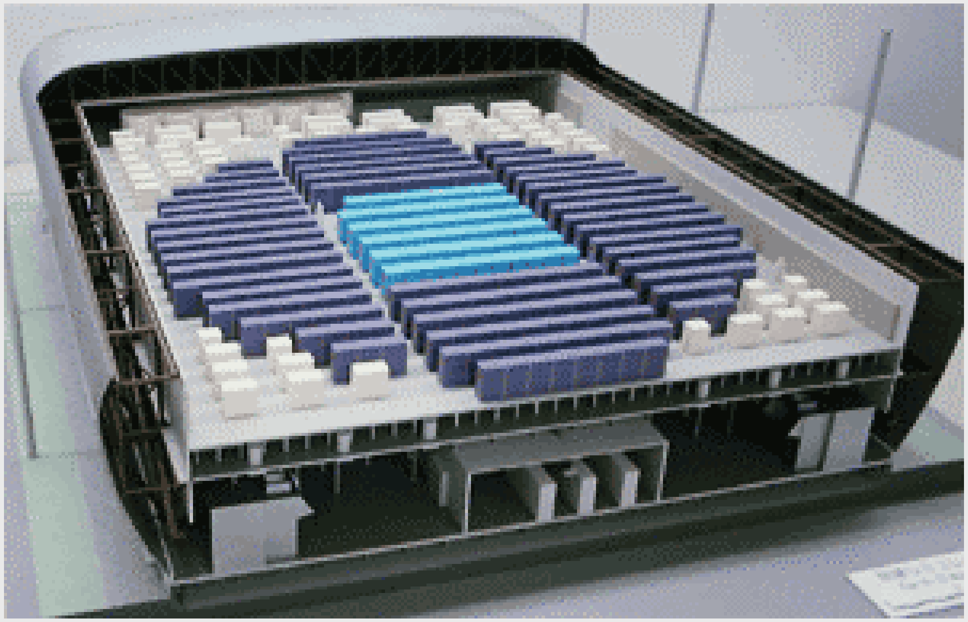

Japan’s Earth Simulator

- 640 processor nodes (PNs)
- Each PN is a system with 8 vector-type arithmetic processors (APs)
- Peak performance 40Tflops

approx; 50m x 65m x 17m

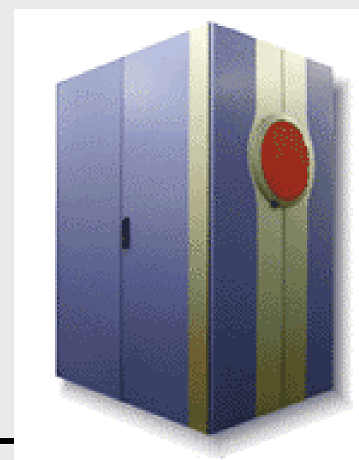

1.4m x1m x 2m

Embedded Systems: What is the new

- Devices become smaller and more powerful
- Devices are coordinated via network
- From “autonomous computing” to coordinated “human-center computing”

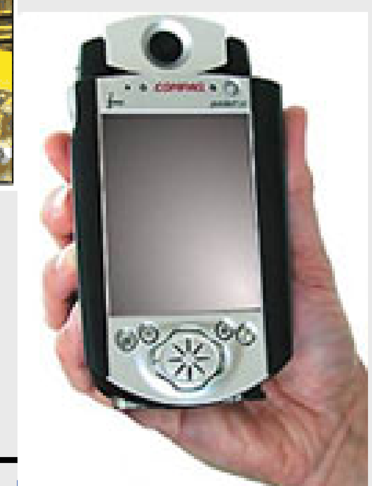

Pervasive Computing

MIT's view of pervasive computing

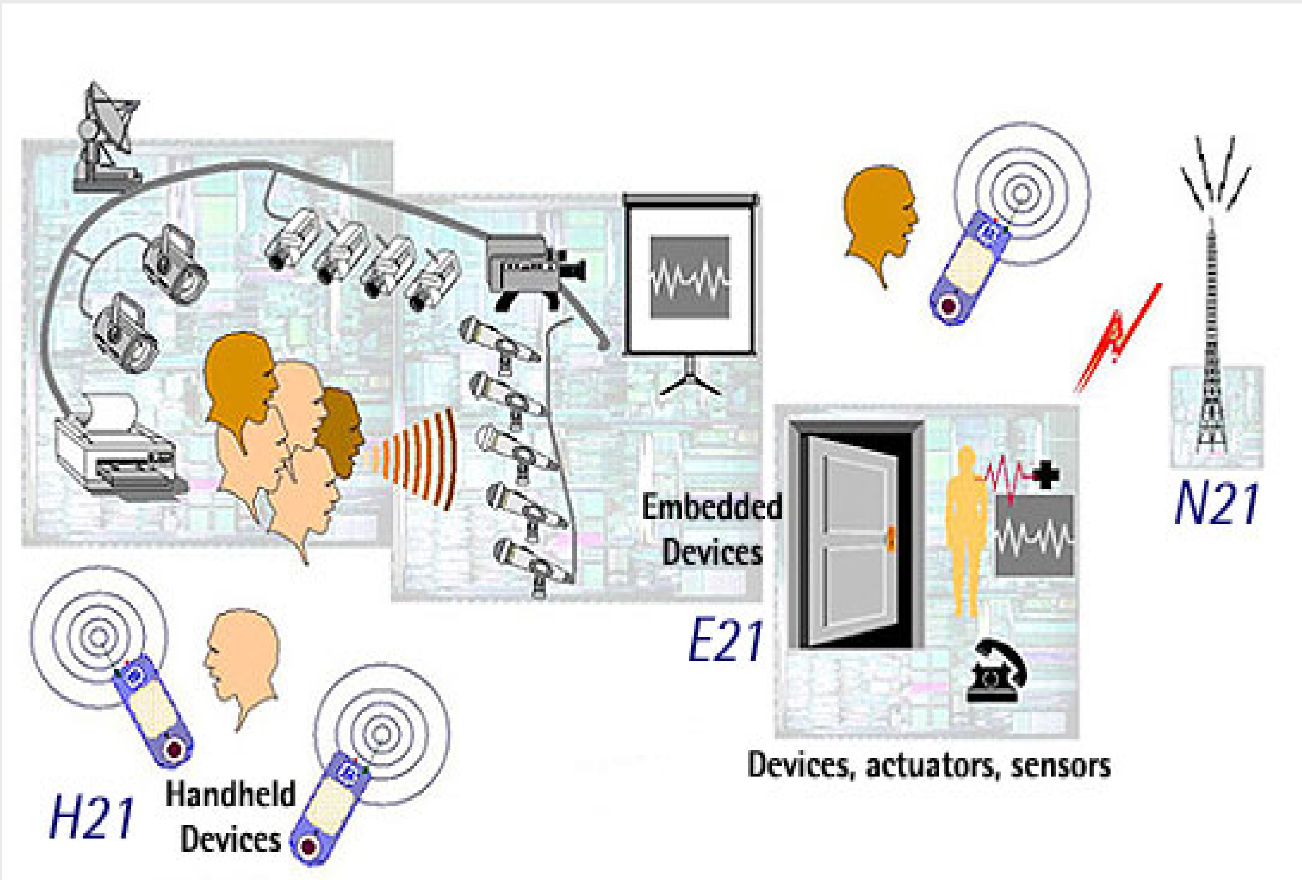

Pervasive Computing

continue

LAST WEEK'S MARKETS

DOW INDUSTRIAL AVERAGE 21.92 8604.60 Up 0.3 percent | NASDAQ COMPOSITE 17.27 1520.15 Up 1.2 percent | STANDARD & POOR'S 500 3.33 933.41 Up 0.4 percent

BUSINESS

CHICAGO SUN-TIMES MONDAY MAY 12, 2003

EDITOR: Dan Miller | TO REACH US: (312) 321-2841, stbusiness@suntimes.com

PAGE 51

'Plain old telephone service' could get a lot more exciting

BY HOWARD WOLINSKY
Business Reporter

Four years ago, telecom experts were about ready to write off the 100-year-old wired phone system, known in the industry as POTS (Plain Old Telephone Service), predicting it would soon be eclipsed by the Internet and wireless phones.

But researchers from the Illinois Institute of Technology and Lucent Technologies' Naperville campus have developed new software that turns POTS into PANS (an industry term for Pretty Amazing New Stuff), moving the old phone system into the Internet loop.

Over the last two years, Vijay Gurbani, 37, a doctoral candidate in computer science at IIT and a Lucent researcher, and his colleagues have added some Internet digital magic to old-fashioned analog phones. This work, which has earned one patent thus far for Gurbani, is expected to result in Lucent commercializing the services within a few years. So the technical breakthrough could be coming soon to phones in the home and office.

"In the Internet era, wired phones were being written off," Gurbani said. "Cellular companies were spending billions of dollars on new infrastructure to make 3G [third-generation] services available for cell phones. I wanted to see if it was possible to make 3G-type services available on landlines."

It was possible.

In fact, IIT and Lucent researchers have developed software that recycles POTS so it can handle buddy lists and instant messaging, the same kinds of services that have been wildly successful over the Internet for America Online and MSN, and also are available over some Internet appliances, such as 2½G and 3G cell phones and personal digital assistants.

"The phone network has been used for its voice capabilities and the Internet for its data capabilities. The two networks have virtually never talked to each other—until now," said Xian-He Sun, Gurbani's computer science professor.

Researchers developed software that connects landline phones to the Internet. Ironically, in the past, the old system mainly was used to connect computers to the Net.

In a demonstration at Lucent, a landline phone was lifted and replaced on the cradle, and it sent a signal that made the network aware that the user was home, popping up an icon on a buddy list on a computer screen.

Buddy lists have been done pre-

See PHONE, Page 56

Xian-He Sun (from left), Vijay Gurbani and Nehal Mehta have developed a way for typical landline phones to tap into buddy lists to aid in communication and even notify a person when a friend is near. -RICH HEIN/SUN-TIMES

PHONE

Continued from Page 51

'Reliable and ubiquitous' wired phones here to stay

viously with computers and cell phones, but never before with signals from landlines, which in the past were considered "dumb terminals," Gurbani said.

Sun said the wired phone network could send an instant message to an Internet user, notifying him of missed calls. This could be a sort of mobile caller-ID to check on traditional phones at home or work.

He said companies could use these buddy lists to set up large-scale phone conferences. In addition, caller ID screens on landline phones could be used to receive text messages.

Nehal Mehta, a doctoral student at IIT, and Lucent researchers Byron Williams and Sudha Gouthma also have worked on the project.

Mehta said the new technology has great commercial potential because wired phones are not disappearing in North America, "where the system is so reliable and ubiquitous."

This opens market possibilities for local phone companies, such as SBC, BellSouth and Verizon, all of which are Lucent customers.

Doug Varney, technical manager of Lucent's network services architecture group, said, "Traditional phone companies are looking for new services like this because they are losing customers to long-distance companies and wireless carriers. They could leverage their embedded investments, and bring in new revenues by offering new services. I think [buddy lists and instant messaging] are just the tip of the iceberg."

Gurbani said the services next will be tested on older-model cell phones, and other services will be developed. For instance, he said the framework could make possible "proximity notification," a service that triggers an e-mail from a friend's cell phone when the friend is nearby. He said the new technology could be used by law-enforcement officials to track suspects.

What about telemarketers or other unwanted parties tapping into the system to find out if someone is home, in effect hacking into buddy lists?

Gurbani said these are real concerns: "Security and privacy issues will have to be worked out."

He sees more advantages than disadvantages to upgrading the old system: "This technology creates even more communication freedom for busy people on the go."

- Cross network(PSTN phone, internet) Service
- Scarlet framework for context aware computing
- Mobility

The IIT HawkTour

Pushing the Boundaries of Pervasive Computing

- Provides tour information based on location and orientation
- Location-aware with WiFi positioning and GPS
- Intelligent content delivery
- Scarlet is used to provide context awareness

HawkTour cont

- Location Awareness
  - Tracking and Web Services
- Application Design
  - User Interface
- Content
  - Information relevant to the HawkTour

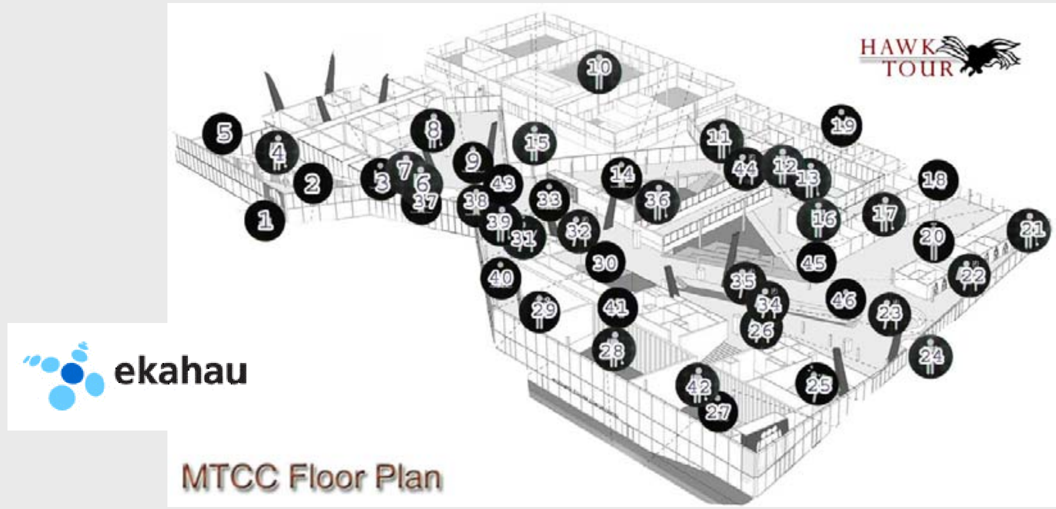

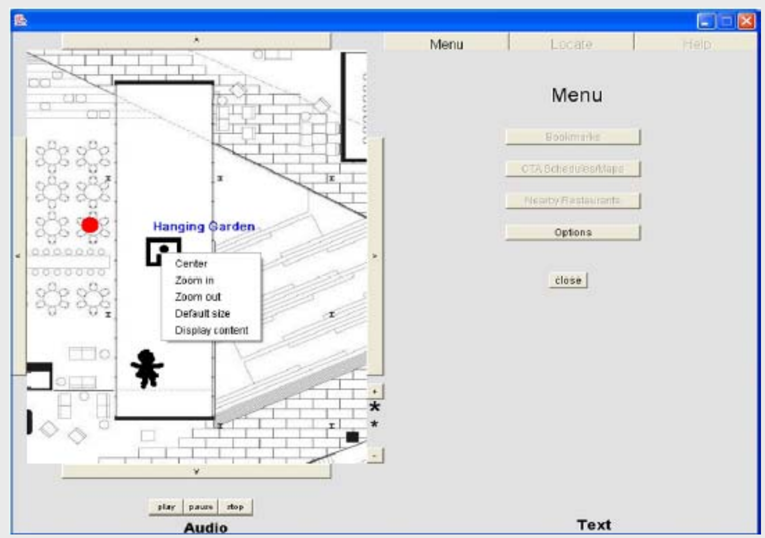

Distributed Computing: What is the new

- Supercomputers become ever powerful
- Communities of "Virtual organizations" are formed
- No VO possesses all required skills and resources
- From "community sharing" to "information grid"

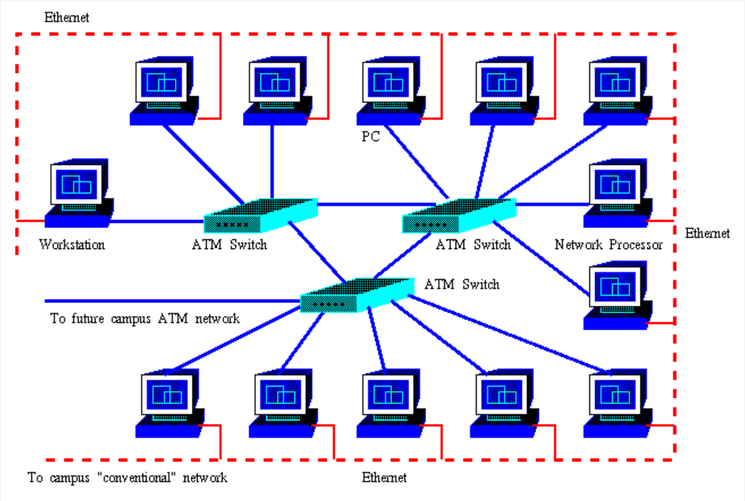

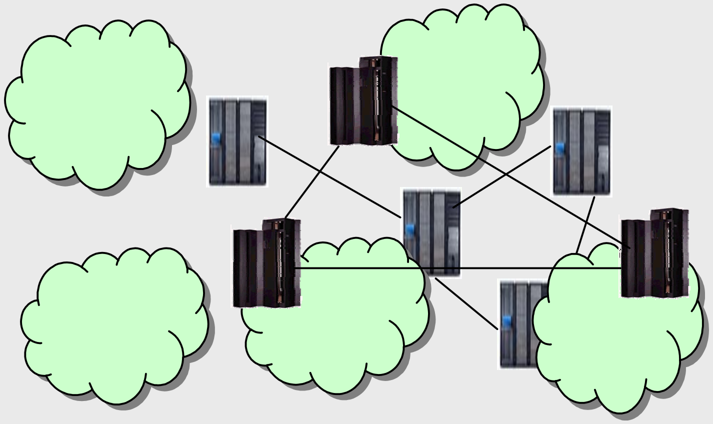

Integrated VOs: the Grid

Mimic the electrical power grid

The Challenge of Grid Computing

Virtualization and Resource Management

Many sources of data, services, computation

Registries organize services of interest to a community

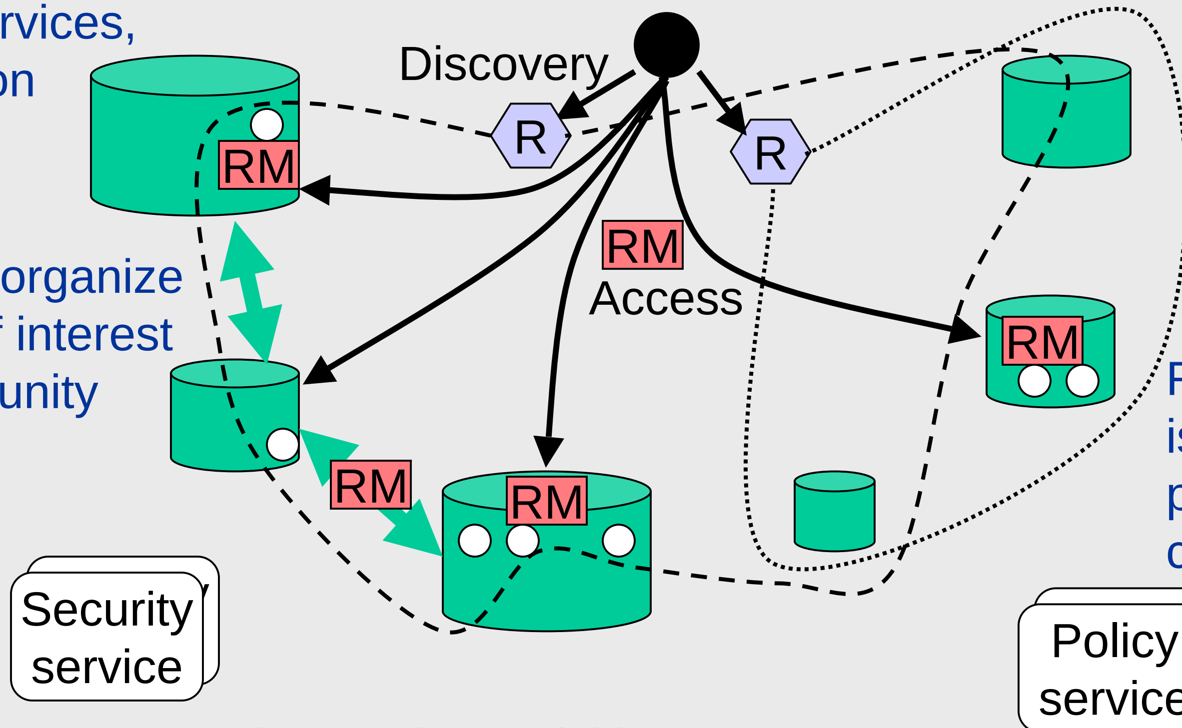

Security & policy must underlie access & management decisions

Resource management is needed to ensure progress & arbitrate competing demands

Data integration activities may require access to, & exploration/analysis of, data at many locations

Exploration & analysis may involve complex, multi-step workflows

Virtual Private Grid (VP-Grid)

- A hosting platform where each user can create and operate in a private grid(s), based on the same shared Grid infrastructure, achieving:
  - Virtualization
  - Isolation and Protection
  - Privacy
  - Accountability and QoS
  - On-demand creation and provisioning

Overview of VP-Grid

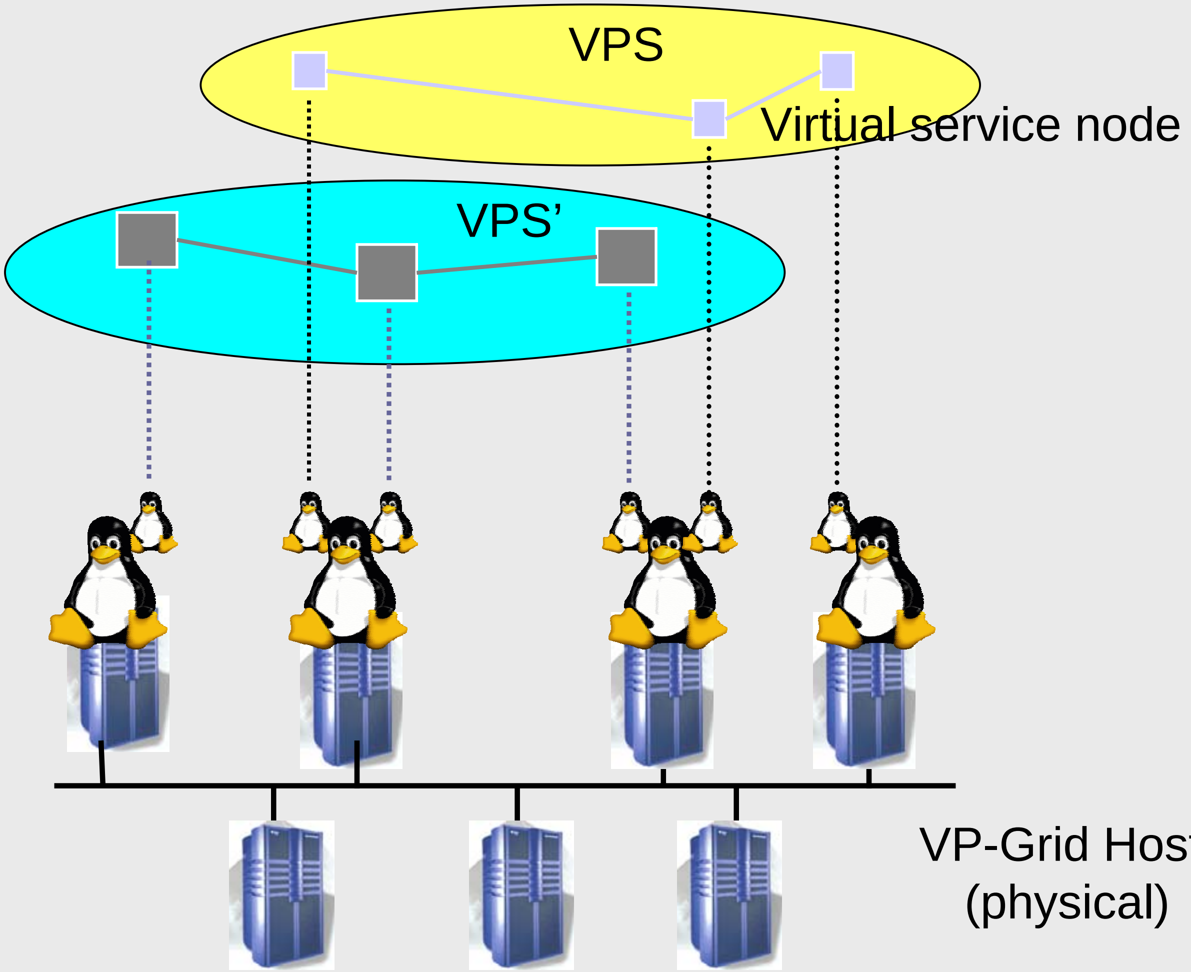

Virtualization: Key Technique

- Two-level OS structure
  - Host OS
  - Guest OS
- Strong isolation
  - Administration isolation
  - Installation isolation
  - Fault / attack Isolation
  - Recovery, migration, and reconfiguration
- Virtual service node
  - VP-Grid Service (VPS)
  - Guest OS
  - Internetworking enabled

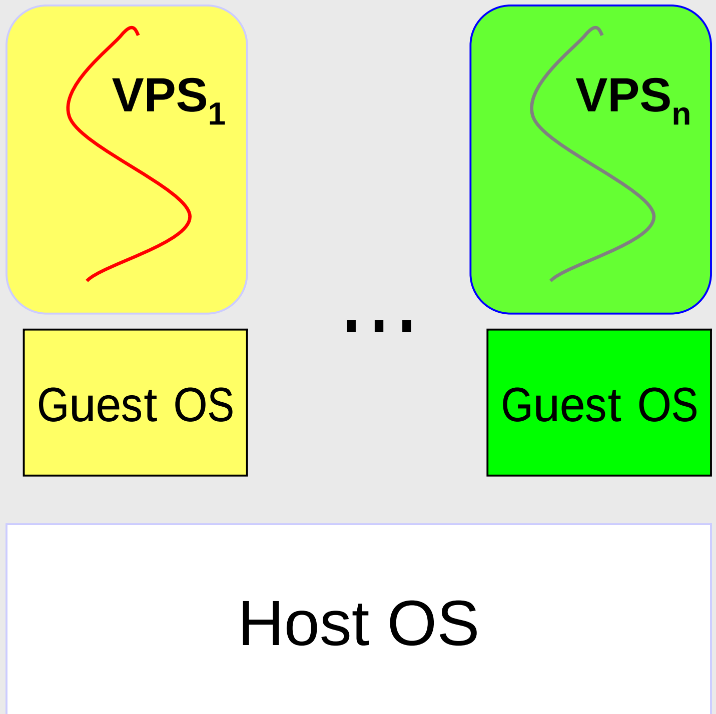

One VP-Grid host

Resource Management & Task Scheduling

- VP-Grid provider selection:
  - Among a set of VP-Grid providers, which one should be chosen to host an VPS?
- VPS selection:
  - Among a set of potential tenants (VPSes), which ones to host? (for QoS, resource utilization, security…)
- The Grid Harvest Service (GHS) System
  - A long-term application-level performance prediction and task scheduling system for non-dedicated distributed (Grid) environments
  - Reservation-based versus shared resources

Rescheduling Algorithm

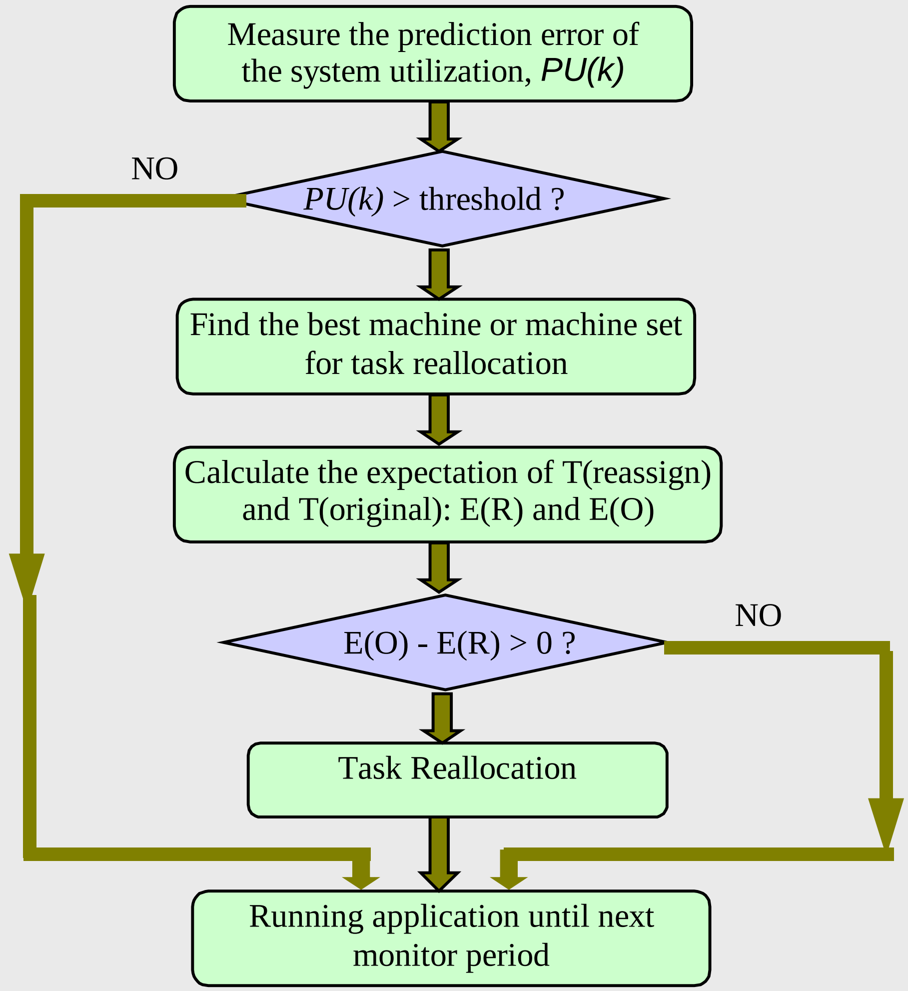

Mobility of VP-Grid

- Mobility is needed for dynamic scheduling of tasks and reconfiguration of the VPS
- Current successes of mobile computing are based on safe-languages such as Java, which is slow and cannot apply to legacy codes
- VP-Grid supports mobility at two-levels, at the virtual machine level and at the code level
- VP-Grid supports mobility of legacy codes written in traditional languages such as Fortran, C, C++

Mobility of Legacy Code

- We have developed novel methodologies and a prototype system, HPCM, to migrate codes written in traditional languages such as Fortran, C, C++

  – Two level mobility: migrate native codes under Java virtual machine

  – General methods: migrate between different computing systems and different virtual organizations.

  – Leading technology, strong mobility

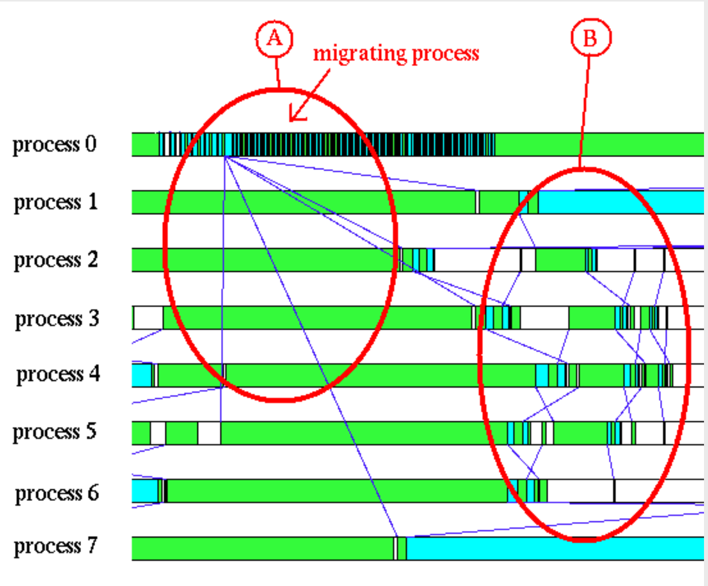

Any Questions?