

# COMP 7970 Storage Systems

## I/O-Aware Load Balancing

**Dr. Xiao Qin**

*Department of Computer Science and  
Software Engineering*

*Auburn University*

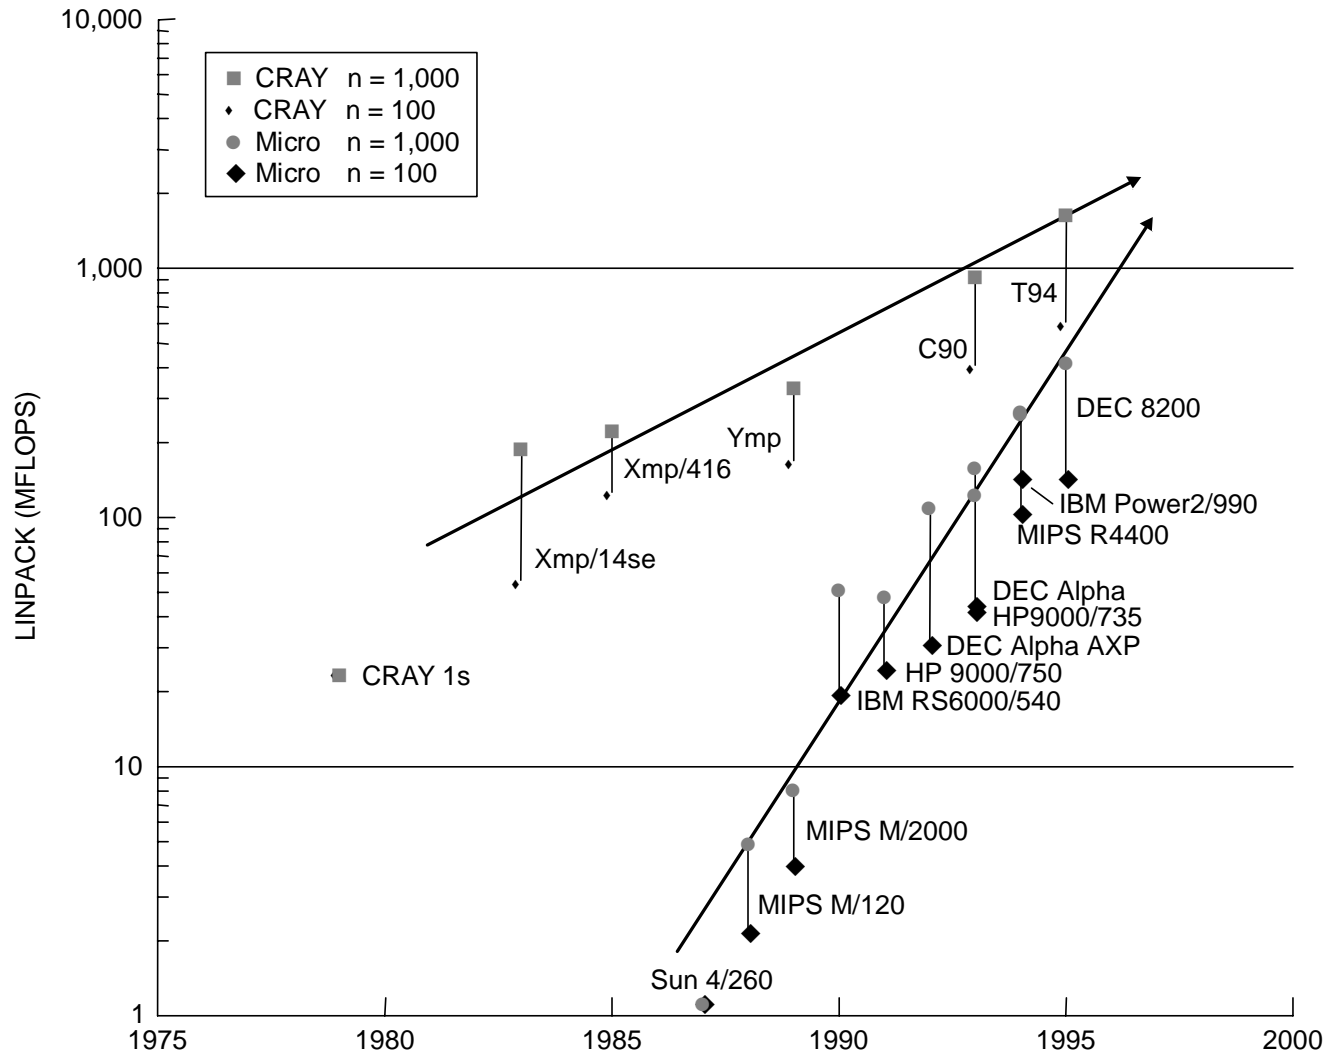
*<http://www.eng.auburn.edu/~xqin>*

*[xqin@auburn.edu](mailto:xqin@auburn.edu)*

# How to Choose a Research Topic?

- Select a topic that **interests** you
- **Job Market**
- Read through **background** information
- Start making a list of **key words**
- Write out your topic as a **statement** and select the **main concepts**
- Start making a list of **words** to describe your **new ideas**


# Raw Uniprocessor Performance: LINPACK



Observations?

Why?

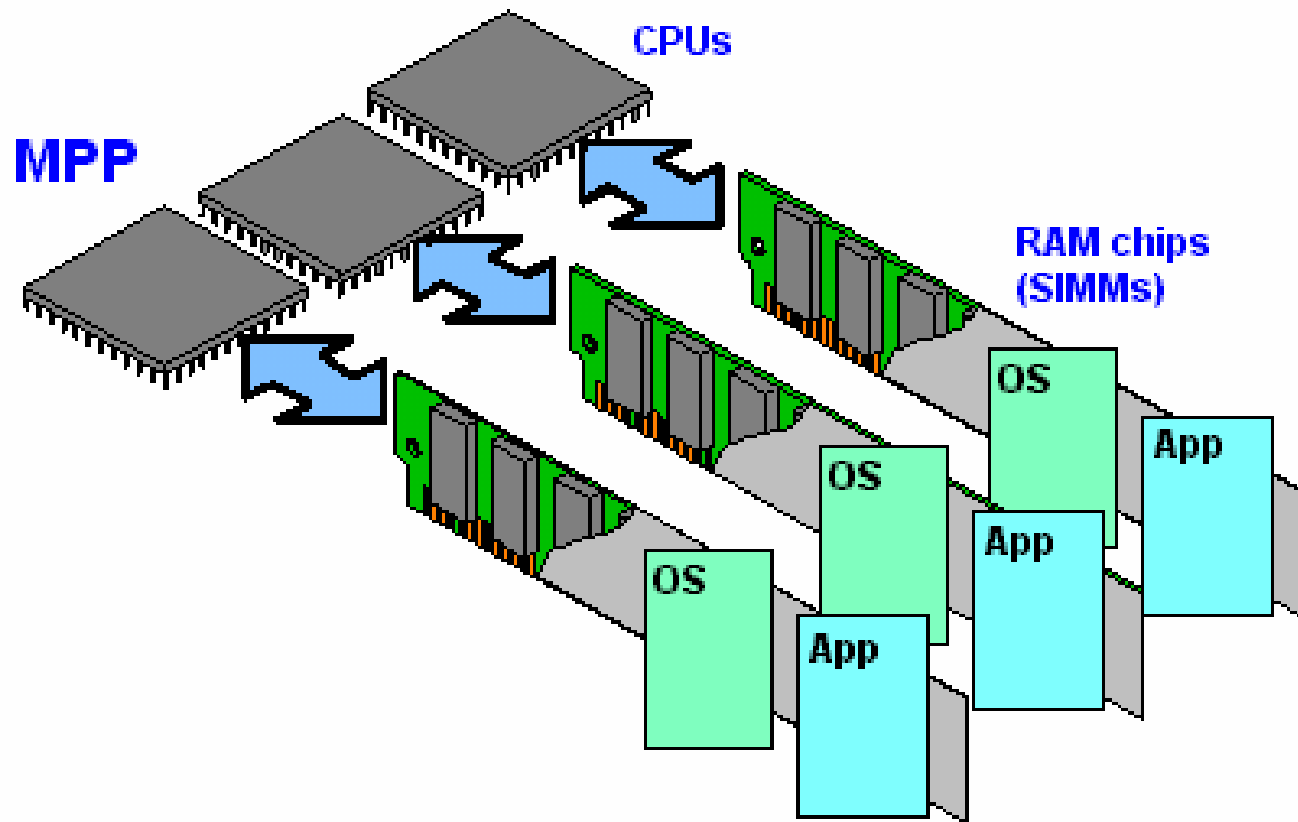
# Raw Parallel Performance: LINPACK



- Even vector **Crays became parallel**
  - X-MP (2-4) Y-MP (8), C-90 (16), T94 (32)
- Since 1993, Cray produces MPPs too (T3D, T3E)

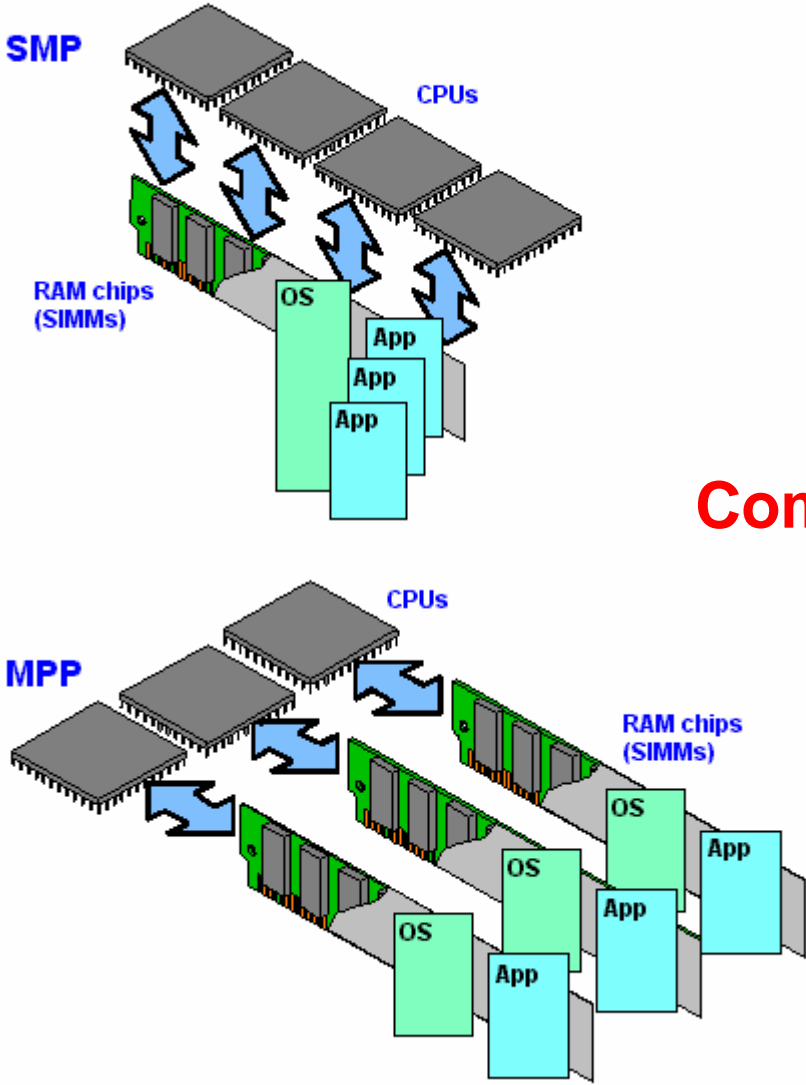
# MPP - Massive Parallel Processing

From Computer Desktop Encyclopedia  
© 1998 The Computer Language Co. Inc.



# SMP - Symmetric Multiprocessing

From Computer Desktop Encyclopedia  
© 1998 The Computer Language Co. Inc.

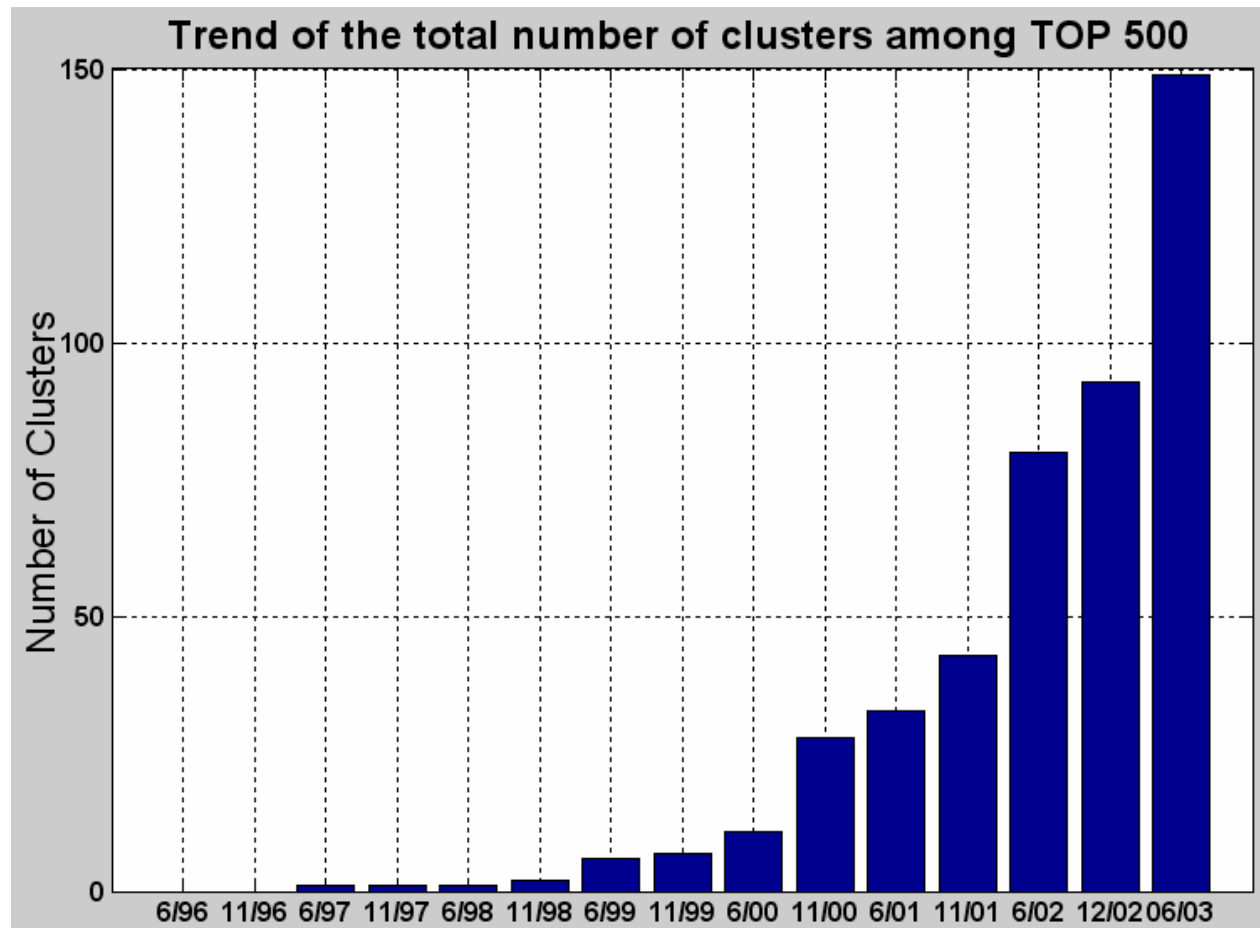


**Comparisons?**

# The PrairieFire Cluster at the University of Nebraska-Lincoln



# Clusters – fastest growing platform



Cluster development trend in the past 7 years.


# I/O-intensive Applications



long running simulations




remote-sensing database systems



biological sequence analysis

# Motivation



- I/O-intensive Applications require input and output of **large amounts of data**.
- I/O performance can be a potential **bottleneck**.