Shared Memory Computing on Networks of Workstations

Willy Zwaenepoel
Department of Computer Science
Rice University

Distributed Shared Memory (DSM)

Software provides shared memory image

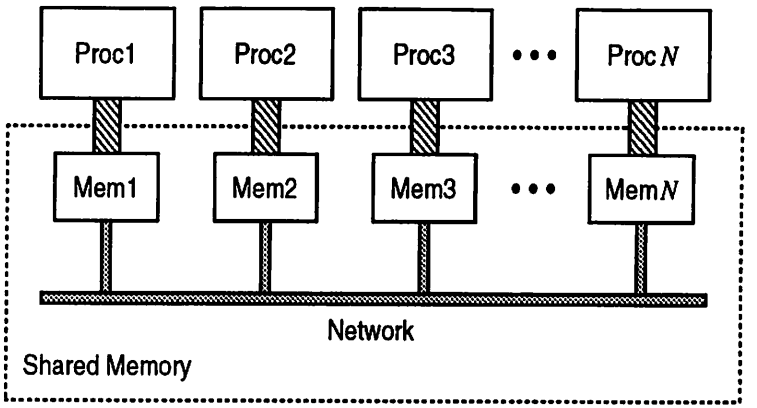

Motivation:

Easier to program than message passing

DSM Programming Model

Shared Memory Programs:

- Threads
- Synchronization
  - Locks
  - Barriers
  - Flags

Key Point

Distributed shared memory:

- support for parallel processing
- on networks of workstations
- for real problems
- with reasonable efficiency
- with reasonable programmer effort

Implementation Issues

Conventional implementations:

- Sequential consistency
- False sharing

Goal:

- Reduce communication
- Keep shared memory model

Solutions:

- Lazy release consistency [Keleher 92]
- Multiple writer protocol [Carter 91]

TreadMarks

Standard kernel and compilers

User-level library for C and Fortran

Implemented on

- SPARC
- DECStation, DEC Alpha
- IBM RS-6000/SP-1
- SGI Indy

[Keleher et al. 94]

Two Applications

Mixed Integer Programming

Genetic Linkage

Mixed Integer Programming

Mixed Integer Programming =

Linear Programming +

Some of the variables are integers

A 2-dimensional example:

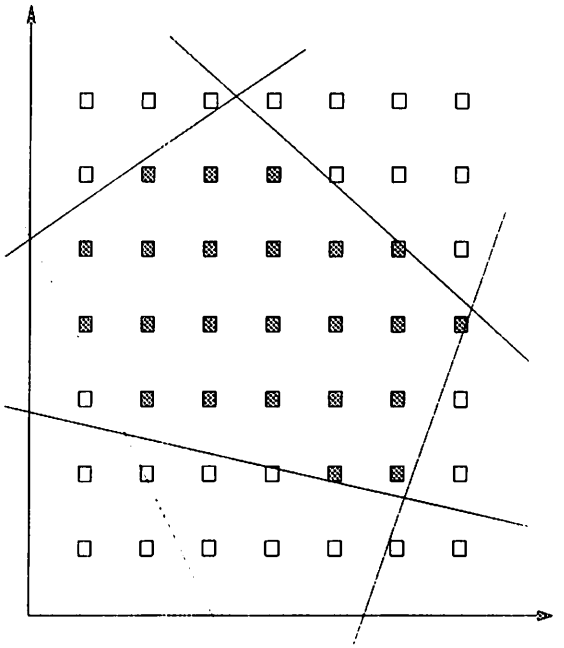

Complexity

Hard in a theoretical sense

Hard in a practical sense:

real instances run for a long time

Branch-and-Bound

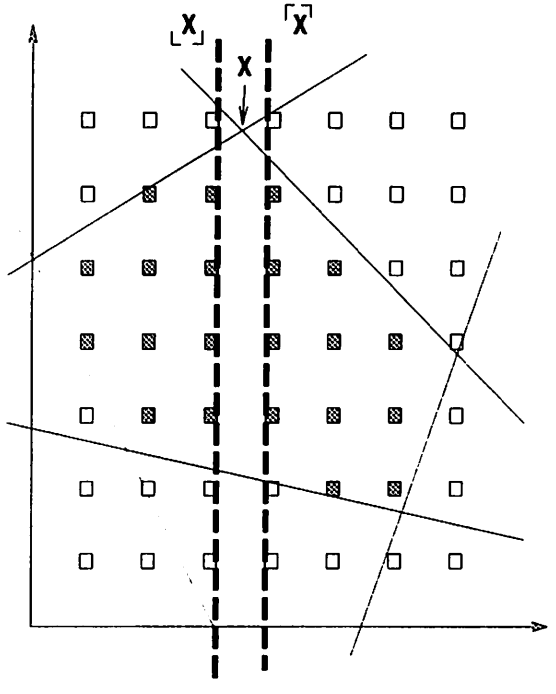

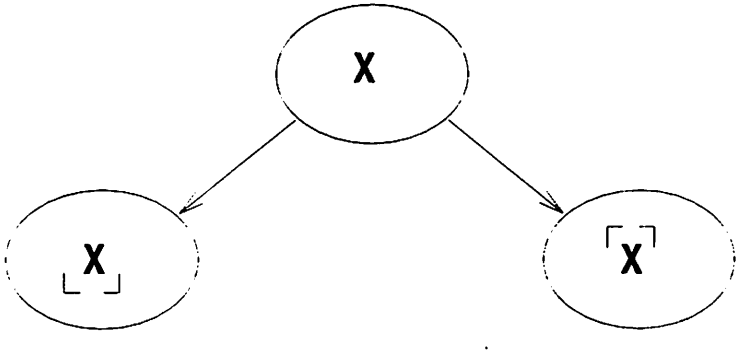

Algorithmic Smarts

Plunging

Pick the right variable

Pick the right node

Cutting planes (branch-and-cut)

10,000 lines of C code

Results

MIPLIB problems longer than 2,000 seconds on 1 processor.

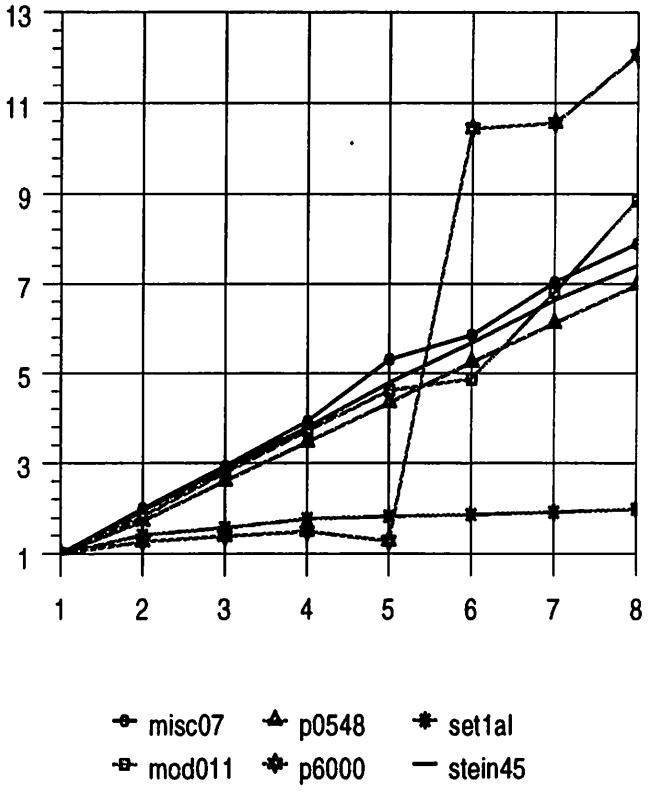

[Lee et al., 1994]

Neat Result

D. Bienstock and O. Gunluk,
Computational Experience with a Difficult
Mixed-Integer Multicommodity Flow Problem,
to appear in Mathematical Programming

Lightwave network configuration (Bellcore)

521 variables, 56 0/1 variables

664 constraints

Previously unsolved

Solved on an 8-node SP-1 (3 1/2 days)

Genetic Linkage Analysis

Disease gene location:

- biological experiments
- computational steps (linkage analysis)

Computation is bottleneck

Hours to months is normal

Better accuracy desired

A 1-Minute Intro to Genetics

Probability of recombination θ

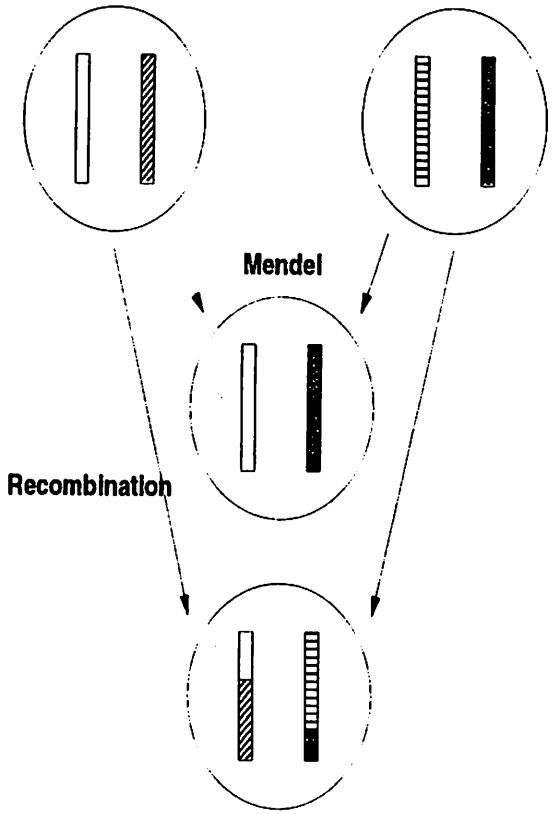

The Linkage Computation

Maximum likelihood optimization of θ

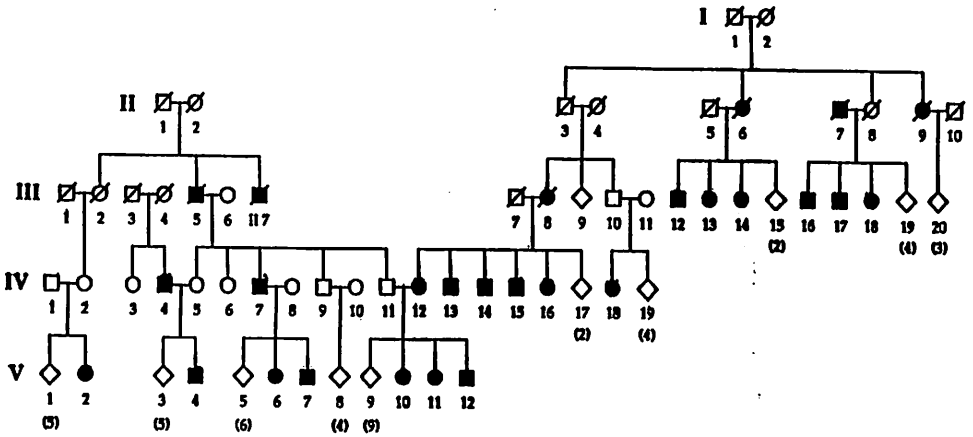

Linkage Parallelized

Optimize for θ

For each nuclear family

Split up rows of R into p sets

For each processor

Do updates for assigned rows

Synchronize

Load balancing in splitting

13,000 lines of C code

Results

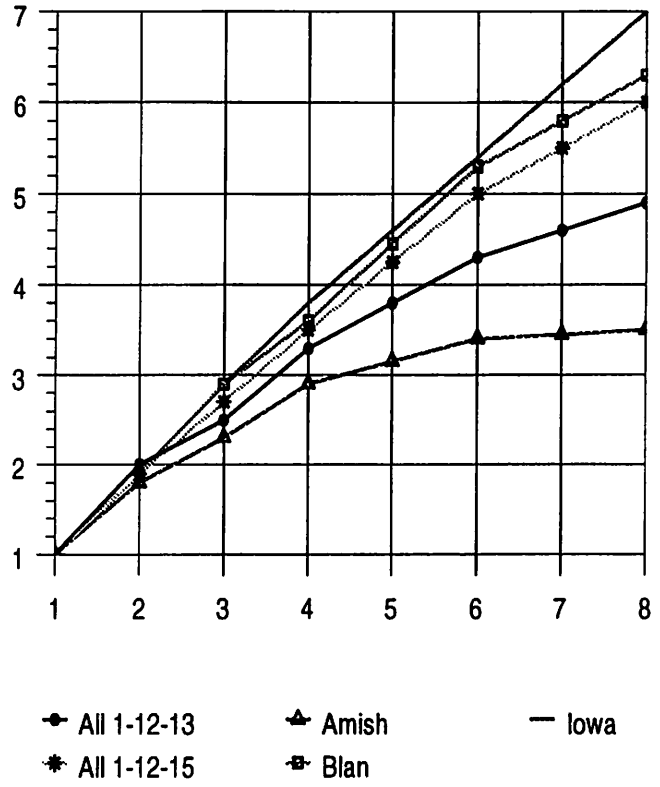

[Dwarkadas et al., 1994]

Lessons Learned

Real problems can be solved

on networks of workstations

using distributed shared memory

with reasonable efficiency

with reasonable programmer effort

Wishlist

Real payoff: solving large/unsolved problems

Need methods/tools for large problems