110043 N95-16458

Rapid Solution of Large-scale Systems Of Equations

by Olaf O. Storaasli, (O.O.Storaasli@larc.nasa.gov or 804-864-2927)

for Workshop on the Role of Computers in Langley R&D (6-15-94)

The analysis and design of complex aerospace structures requires the rapid solution of large systems of linear and nonlinear equations, eigenvalue extraction for buckling, vibration and flutter modes, structural optimization and design sensitivity calculation. Computers with multiple processors and vector capabilities can offer substantial computational advantages over traditional scalar computers for these analyses. These computers fall into two categories: shared-memory computers (e.g., Cray C-90) and distributed-memory computers (e.g., Intel Paragon, IBM SP-2).

Shared-memory computers have only a few processors (16 on a Cray C-90), which rapidly process vector instructions (simultaneous adds and multiplies) and address a large memory. Information is shared among processors by referencing a common variable in shared-memory.

Distributed-memory computers may have thousands of processors, each with limited memory. Explicit message passing commands (i.e. send, receive), are used to communicate information between processors. Such communication is time consuming, so algorithms need to be designed to run efficiently on distributed-memory computers.

This presentation will cover general-purpose, highly-efficient algorithms for: generation/assembly of element matrices, solution of systems of linear and nonlinear equations, eigenvalue and design sensitivity analysis and optimization. All algorithms are coded in FORTRAN for shared-memory computers, and many adapted to distributed-memory computers. The capability and numerical performance of these algorithms will be addressed.

O. Storaasli, D. Nguyen, M. Baddourah and J. Qin (1993), "Computational Mechanics Analysis Tools for Parallel-Vector Supercomputers", AIAA/ASME/ASCE/AHS/ASC 34th Structures, Structural Dynamics and Materials Conference Proceedings, Part 2, pp. 772-778 (Int. J. of Computing Systems in Engineering, Vol 4, No. 2-4, 1993)

Dr. Olaf Oliver Storaasli is a senior research scientist in computational mechanics at the NASA Langley Research Center, Hampton, Virginia. He began his career at Langley after receiving a Ph.D. degree in Engineering Mechanics from North Carolina State University in 1970.

Long before parallel computers were commercially available, Dr. Storaasli led a hardware, software and applications team at NASA Langley Research Center to develop one of the first parallel computers, the Finite Element Machine. He has authored over 80 works in computational structural mechanics including static and dynamic structural analysis, eigenvalue and optimization methods, interdisciplinary analysis, data management, and parallel-vector structural analysis methods on supercomputers. He received the Floyd L. Thompson Fellowship of NASA Langley Research Center for post-doctoral research at Norges Tekniske Hogskole in Trondheim, Norway, and Det Norske Veritas, Oslo, Norway, during 1984-85 and has been invited back twice since He received 5 NASA-wide and 8 Langley Achievement awards for outstanding work in Computational Structural Mechanics. These awards included significant contributions to the NASA Viking and Integrated Programs for Aerospace-Vehicle Design (IPAD) Projects as well as to the development of Relational Information Management (RIM), since developed into the commercial relational data-base software: R:BASE. In August, 1989, Cray Research selected the general-purpose matrix equation solution software, pvsolve, developed by Dr. Storaasli and his colleagues, to receive the GigaFLOP Performance Award. pvsolve was used to solve the 54,870 equations (9.2 billion floating point operations) in the Space Shuttle Solid Rocket Booster structural analysis in six seconds elapsed time. His recent research has resulted in methods to analyze a 172,400 equation (5,737 bandwidth) refined model of a high speed civil transport and a 265,000 equation automobile (3,374 bandwidth) application in less than two minutes on the Cray C-90 and a method to generate and assemble stuctural stiffness matrices on the Intel Delta at speeds 25 times that of one Cray C-90 processor.

Rapid Solution of Large Systems of Equations

Dr. Olaf Storaasli
Computational Structures Branch
Mail Stop 240
NASA Langley Research Center
Hampton, VA 23681

Email: O.O.Storaasli@larc.nasa.gov
Phone: 804-864-2927
FAX: 804-864-8912

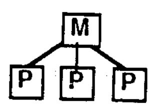

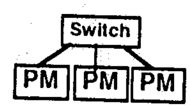

presented at

Workshop on
The Role of Computers in Langley R&D
June 15, 1994, Reid Conference Center
NASA Langley Research Center

Objective

- *Faster, cheaper, better* analysis/design of large-scale structures
  - Develop algorithms to exploit high-performance computers
  - Evaluate computational performance

Outline

- Supercomputers & Structural Models
- Structural Analysis
  - Nodal Generation and Assembly
  - Linear Equation Solvers

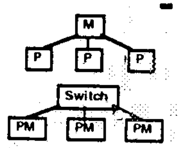

    Shared-memory computers

    Distributed-memory computers
  - Nonlinear Equation Solvers
- Structural Optimization
- Design Sensitivity

Langley Research Center

Parallel-Vector Speedup
(over sequential code)

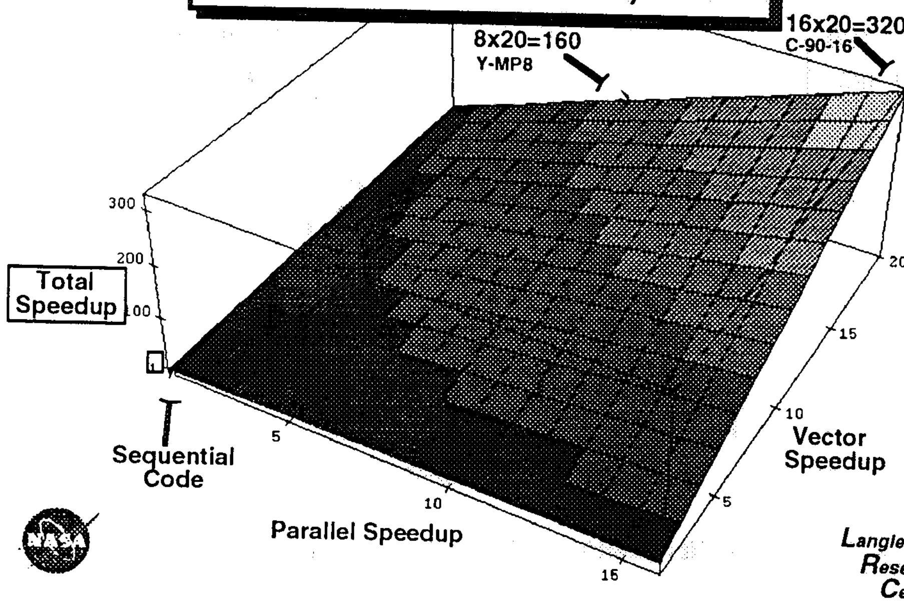

Langley Research Center

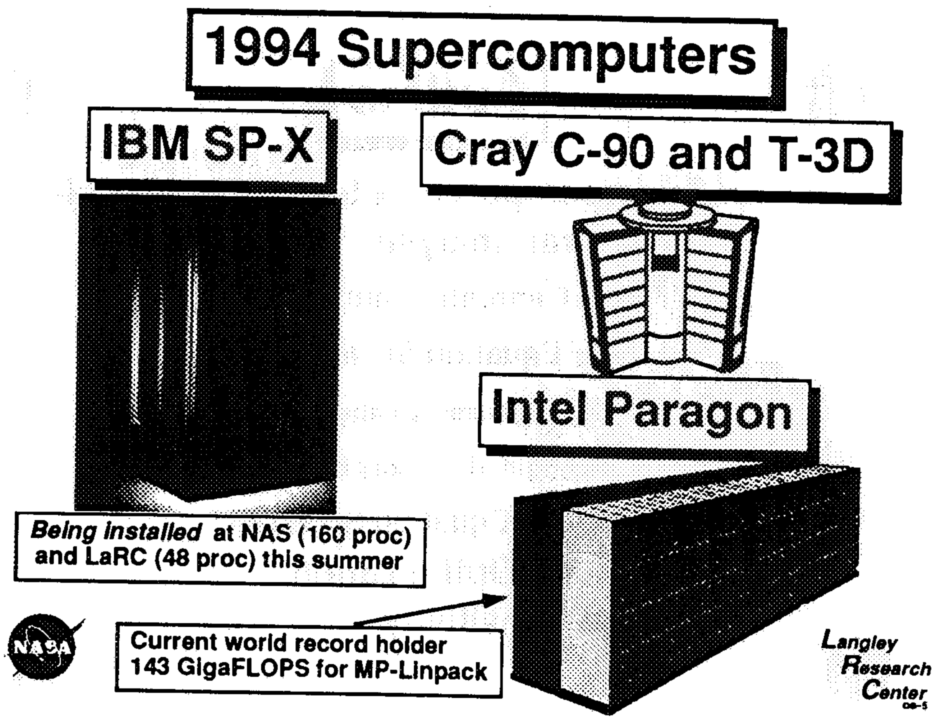
1994 Supercomputers
IBM SP-X
Cray C-90 and T-3D
Intel Paragon
Being installed at NAS (160 proc) and LaRC (48 proc) this summer
Current world record holder 143 GigaFLOPS for MP-Linpack
Langley Research Center

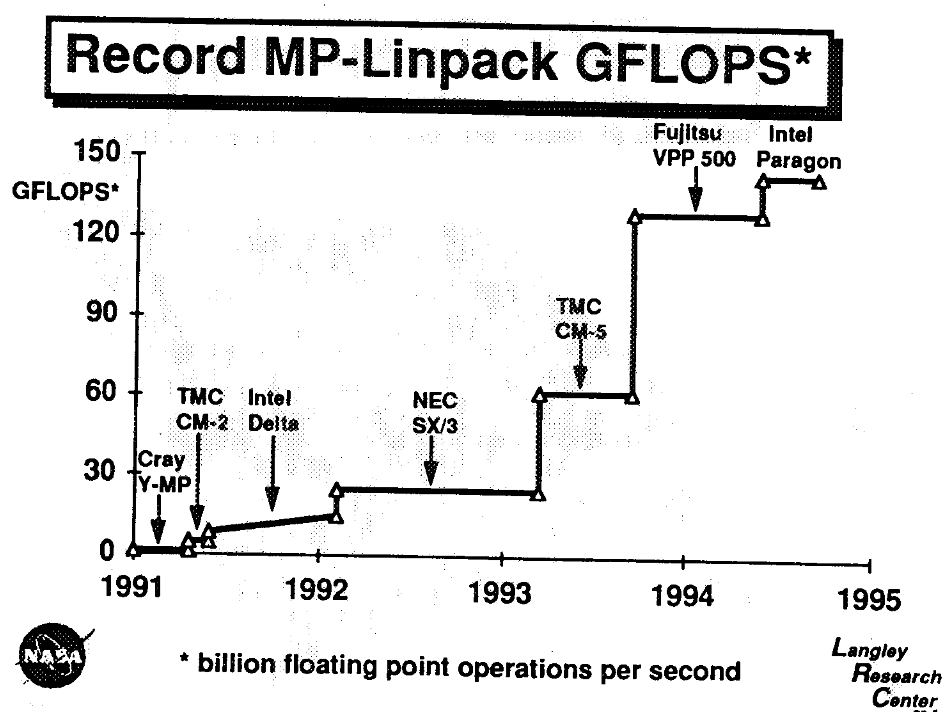
Record MP-Linpack GFLOPS*
GFLOPS*
150
120
90
60
30
0
1991
1992
1993
1994
1995
Cray Y-MP
TMC CM-2
Intel Delta
NEC SX/3
TMC CM-5
Fujitsu VPP 500
Intel Paragon
* billion floating point operations per second
Langley Research Center

Typical Structural Analysis

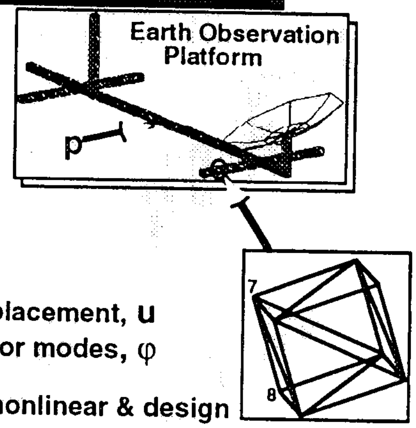

- **Generate mesh** (nodes and elements)
- **Assemble** stiffness [K], mass [M], and load {p}
- **Solve:** [K] {u} = {p} for displacement, u
  [K] {φ} = λ [M] {φ} for modes, φ
- **Repeat:** multiple analyses for nonlinear & design
- **Plot:** u, stresses and vibration modes, φ

Langley Research Center

Performance Assessment Applications

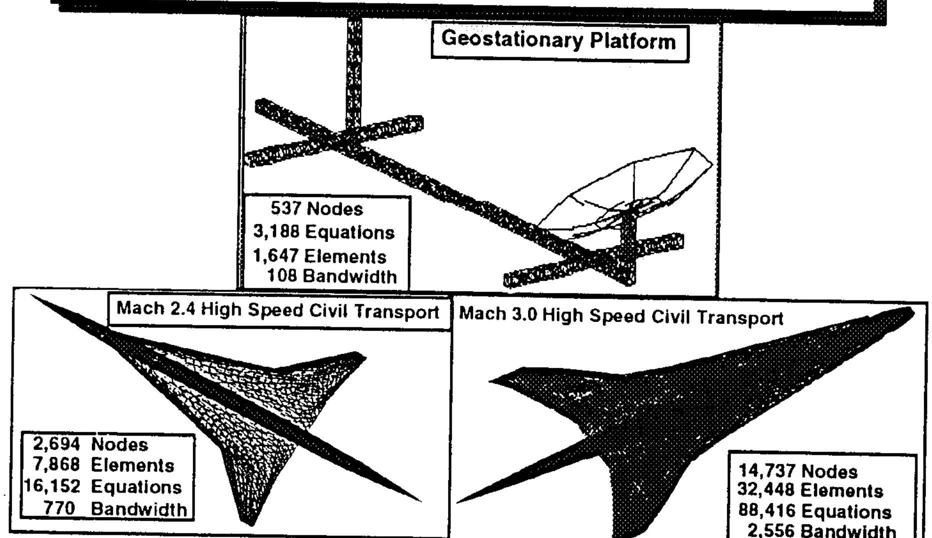

Langley Research Center

Matrix Storage Methods

2910 Equation Stiffened Panel

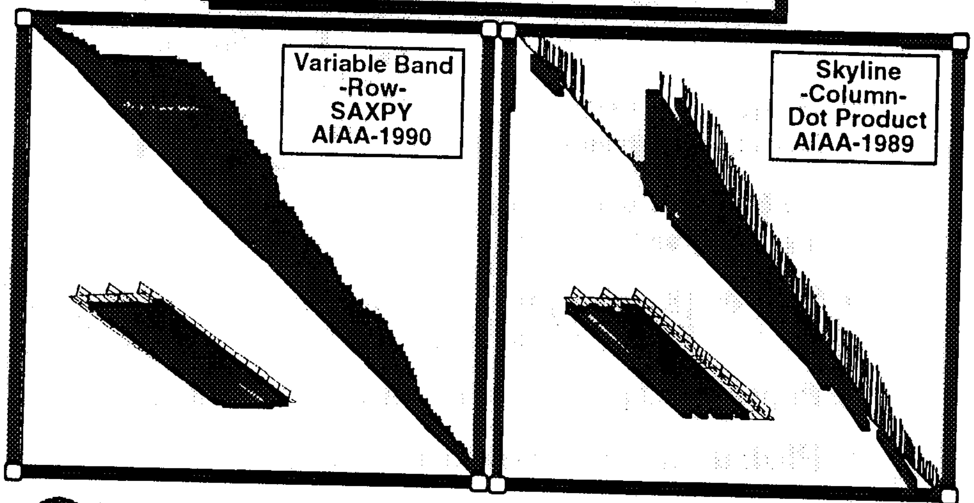

Langley Research Center

OS-9

Parallel-Vector Structures Algorithms

Static	Eigenvalue	Dynamics-Control	Flutter	Optimization
$Ku = f$	$K\varphi = \lambda M\varphi$	$M\ddot{u} + C\dot{u} + Ku = f(t)$	$K\Phi = \lambda M\Phi$	$b_{k+1} = b_k + s_k d_k$
Substructuring NL Algorithms	Subspace Lanczos	Time Integration Reduced-Order Simulate Multibody	Unsymmetric Choleski and Lanczos	Search methods Sensitivity

Matrix Assemblers

- Finite Element based
- Degree-of-Freedom based

Equation Solvers

- Direct
- Iterative
- Sparse
- SVD

Langley Research Center

OS-10

Structural Analysis Computation Time

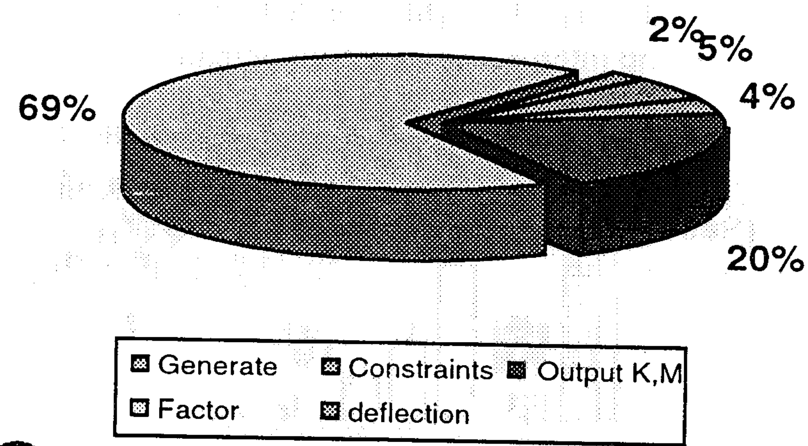

Langley Research Center

Parallel Matrix Generation and Assembly

By element: Traditional thinking
Generate $[k^{(e)}]$ on different processors
Assemble global $[K] = \sum [k^{(e)}]$

can't write elements simultaneously!

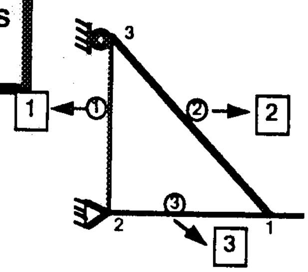

By node: New method

Nodal Connectivity

Node	Proc.	Elements
1	1	②,③
2	2	①,③
3	3	①,②

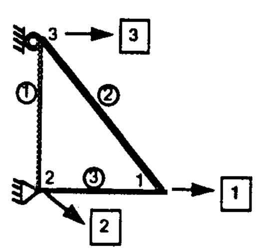

Langley Research Center

Parallel Structural Matrix Generator/ Assembler Demonstrated on HSCT

- Nearly ideal parallel speedup
- (no interprocessor communication)

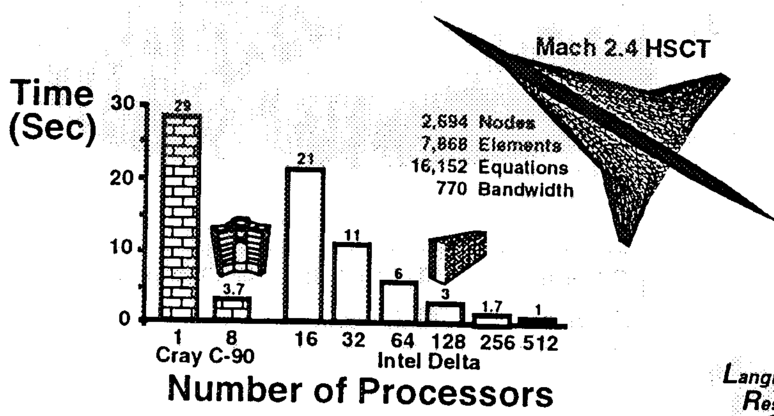

Langley Research Center

Equation Solution Issues
(Time, memory, disk space, I/O)

- Iterative or direct ?
- Banded or sparse ?
- "In-core" or "out-of-core" ?

Communication

- Broadcast or ring?
- OSF or SUNMOS?

Langley Research Center

Equation Solvers

- Iterative and ***Direct*** (function of application)
- Linpack (MP Linpack), LApack
  (needs full matrix for best performance)
- Banded Indefinite, nonsymmetric
  (requires pivoting)
- Banded Definite Symmetric
  (seldom occurs in practical structures)
- Skyline*, *Variable-band**
  (DOT-product, SAXPY operations minimize time)
- *Sparse**, Wavefront* (<5% nonzeros)

* node or equation reordering minimizes solution time

Langley Research Center

Direct Equation Solvers

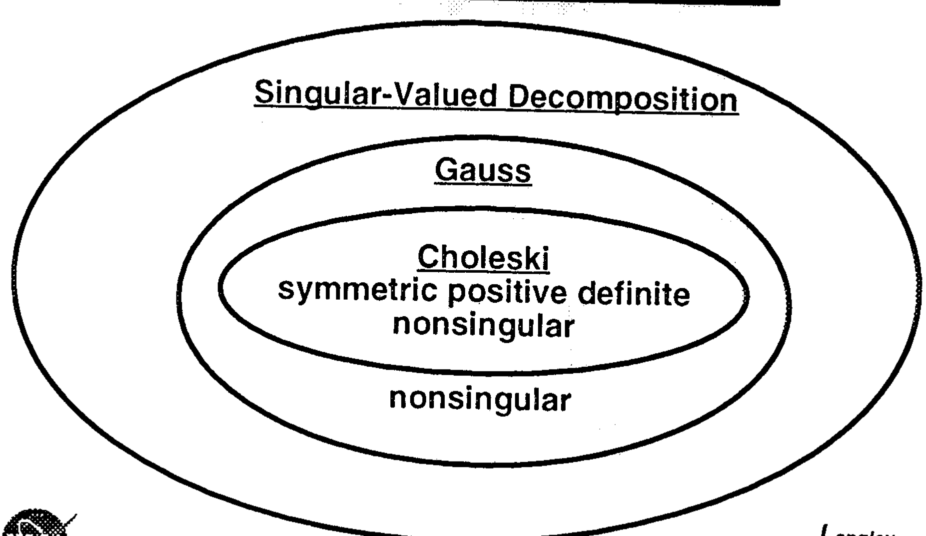

Langley Research Center

Iterative vs Direct Solvers

- Iterative slow, convergence not guaranteed
- Direct complex coding (banded, sparse)

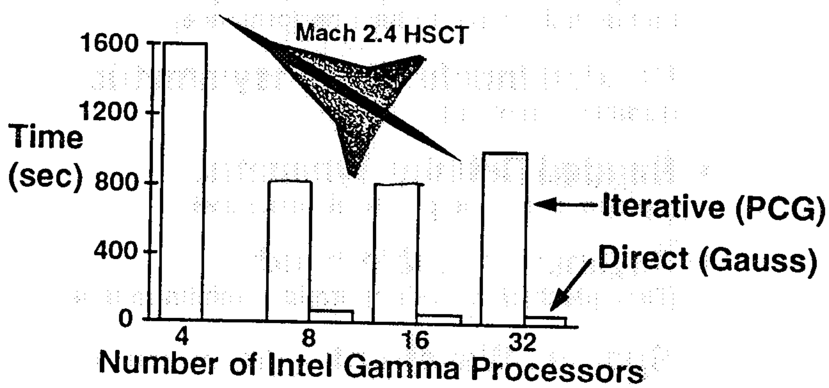

Langley Research Center

"Out-of-core" Direct Solver

- using Cray Solid State Disk -

- as fast as "in-core" solver
- memory used: 1.1 x bandwidth² (or 24 x bandwidth + 6 x neq)

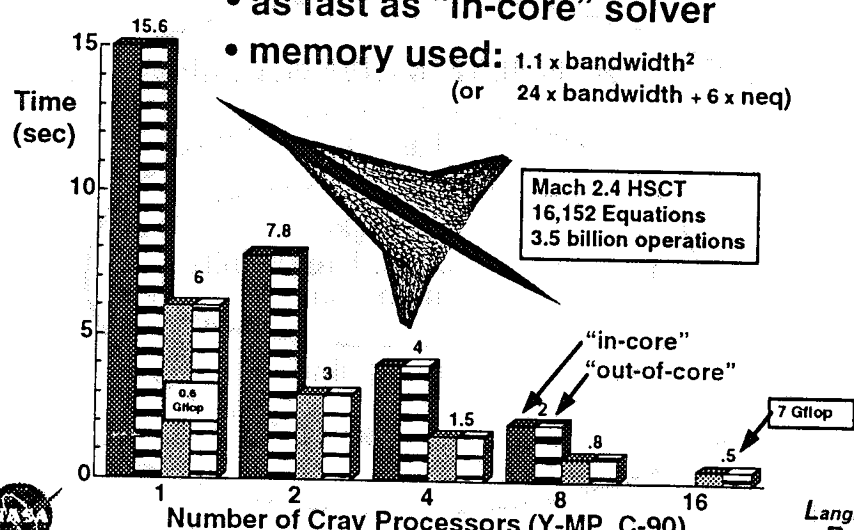

Langley Research Center

Automotive Application
of Sparse Solver

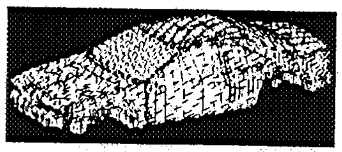

48,894 Elements
44,188 Nodes
263,574 Equations

- **Langley solution took 40 CPU sec (1 Cray C-90 processor)**
  - fastest solution known to date -

- ***Challenge: achieve even faster solution on SP-2 and Paragon!***

Langley Research Center
OS-19

Convex

Equation Solver Results
- 10 Finite Element Models -

- **Iterative slowest, Sparse fastest!**

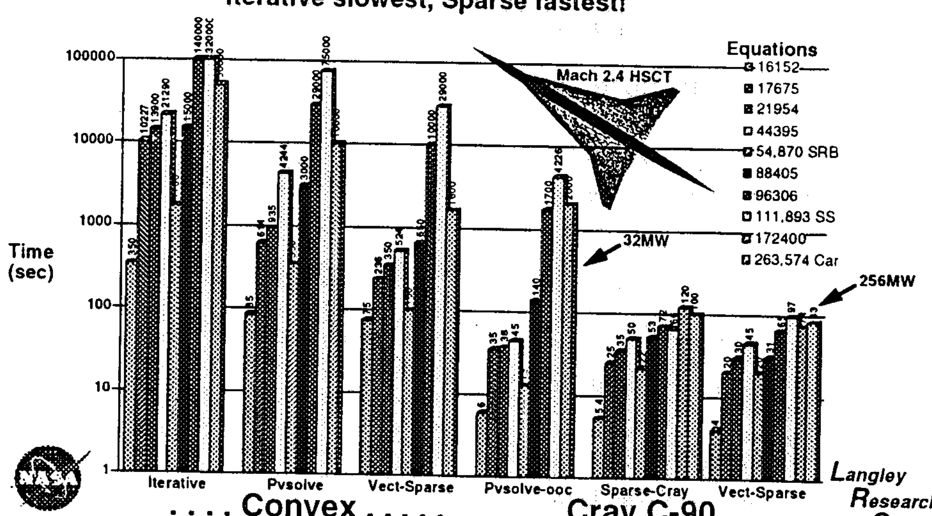

Langley Research Center
OS-20

Industrial-Strength Equation Solvers for $[A]\{x\}=\{b\}$ and $[A]\{x\}=\lambda[M]\{x\}$

Solver	Application	Memory	Parallel(shared)	Parallel(Distributed)
PVSOLVE	Symmetric + Def*	equations X Bandwidth	Yes (Cray C-90, etc)	Yes (Intel Paragon, IBM SP-1)
PVSOLVE-OOC	"	1.1 x Bandwidth^2	Yes (Cray C-90, etc)	Not yet
PVSOLVE-OOC+	"	24 x Bandwidth	No (Cray C-90, etc)	Not yet
VSS (Vector Sparse)	"	function of sparsity	No (Cray C-90, etc)	Not yet
PCG(Iterative)	"	~ matrix nonzeros	Yes (Cray C-90, etc)	Yes (Intel Paragon, IBM SP-1)
LANZ(Eigensolver)	"	equations X bandwidth	Yes (Cray C-90, etc)	Yes (Intel Paragon)

NOTE: These solvers have been evaluated on real applications with up to 263,574 equations and larger matrices with several million equations. PCG is slowest, VSS is fastest (for large, sparse problems) and PVSOLVE-OOC is the best all-around parallel-vector solver. PVSOLVE-OOC exploits Cray solid-state disk.

* special versions of PVSOLVE for unsymmetric and negative coefficient matrices solve panel flutter, CFD, nonlinear and optimization problems

Langley Research Center

08-21

Parallel-Vector Equation Solver (PVSOLVE)

Shared Memory — *Cray GigaFLOP award*

- "in-core skyline and variable-band versions
- "out-of-core" versions: memory ~ 1.2 bandwidth^2 and 24 x bandwidth
- tuned for Crays (or shared memory computer/workstation)

Distributed Memory

- "in-core" skyline - Intel i/860 or Paragon
- "in-core" row version - Intel 860 or Paragon, IBM SP-1
- Conversion underway to TMC CM-5, Convex SSP-1 and Cray T-3D

Use

- COMET, Ford, U. Virginia, IBM, Princeton, LLNL, NSF sites Convex, COMCO, NASA Lewis + several dozen sites

Langley Research Center

08-22

Parallel Geometric Nonlinear Methods

- Newton-Raphson fastest on parallel-vector computers

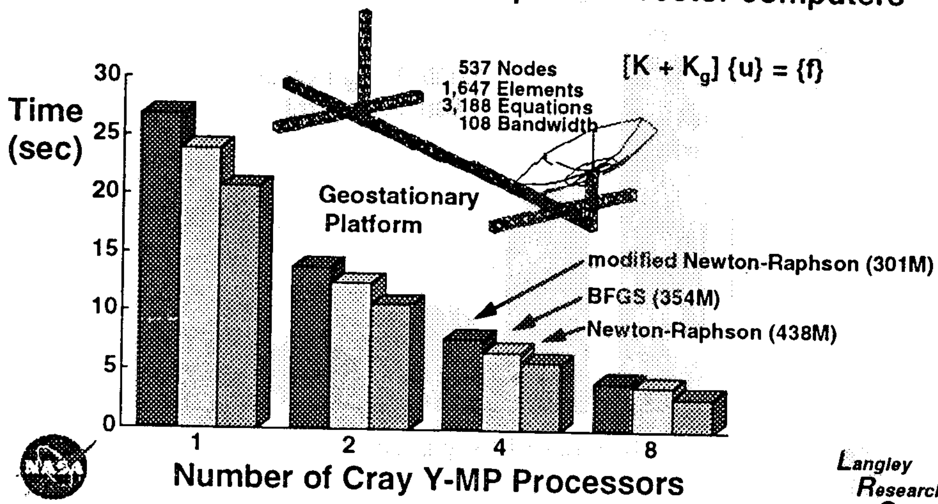

Langley Research Center

Optimization Procedure

- Find aircraft minimum weight subject to displacement and stress constraints
- Nonlinear constrained optimization finds:
  - Direction: BFGS, Simplex-Linear Programming
  - Step size: Golden Block

$b_{k+1} = b_k + s_k d_k$

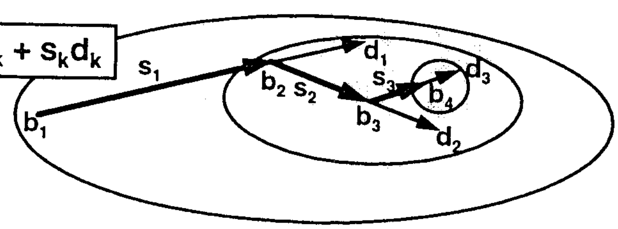

Langley Research Center

Parallel Simplex Method

Linear Programming

- Scalable time reduction

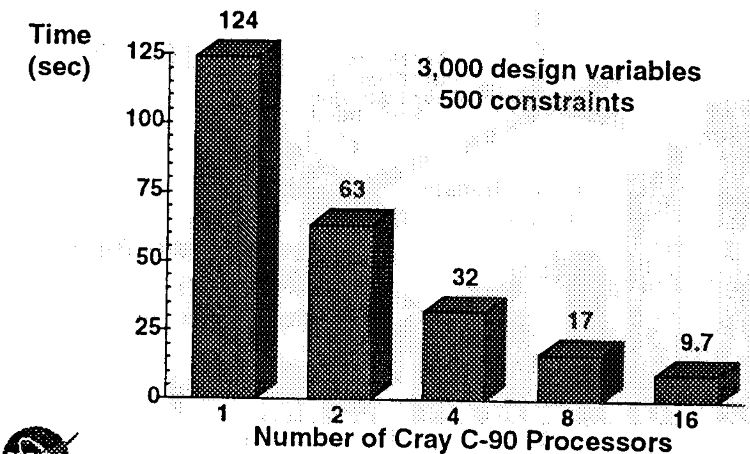

Langley Research Center

Parallel BFGS Optimization

Minimize $F(x_1, x_2, \ldots, x_n)$ is equivalent to

$$\left.\begin{array}{l} F_1(x_1, x_2, \ldots, x_n) = 0 \\ F_2(x_1, x_2, \ldots, x_n) = 0 \\ \vdots \qquad\qquad\qquad \vdots \\ F_n(x_1, x_2, \ldots, x_n) = 0 \end{array}\right\}$$

For 11,000 nonlinear equations:

or

Min $(F_1^2 + F_2^2 + \ldots F_n^2)$

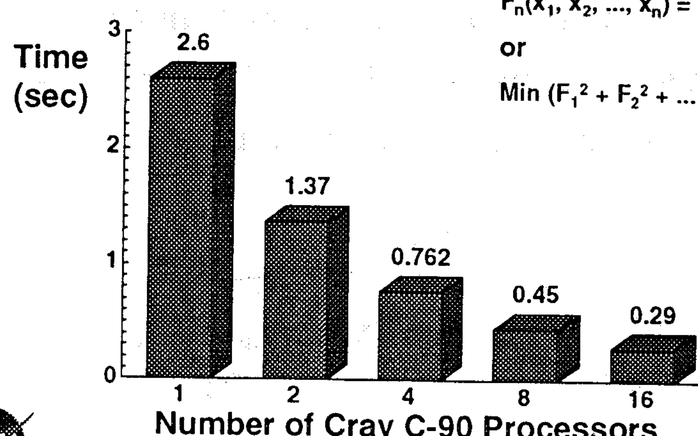

Langley Research Center

Displacement Sensitivity Analysis by Automatic Differentiation (ADIFOR)

– 2-D Truss (80 bays x 190 stories)

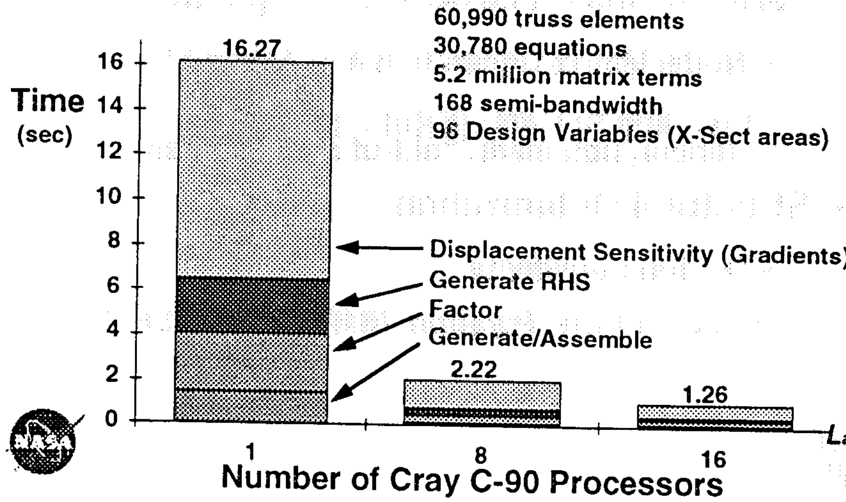

Langley Research Center

Design Sensitivity Analysis Methods for Mach 2.4 HSCT

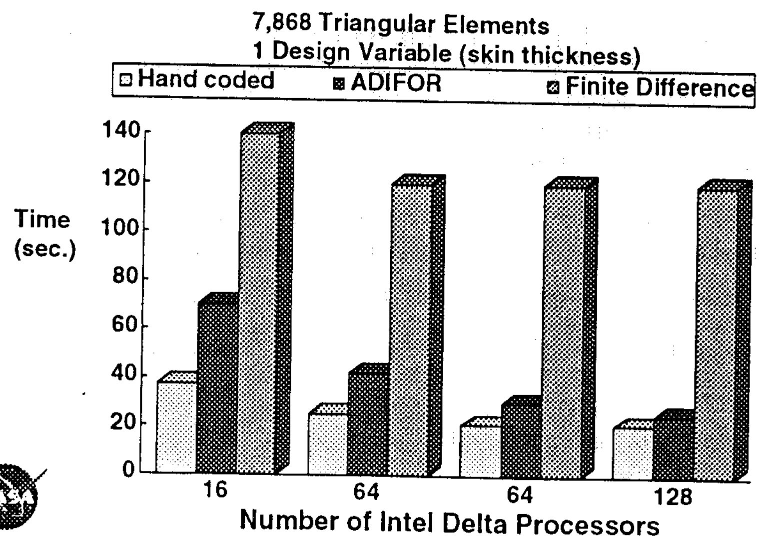

Langley Research Center

Concluding Remarks

- New algorithms for high-performance computers
- Perform well on large-scale applications:
  - Nodal Matrix Generation and Assembly
  - Equation Solvers: [K]{u} = {p}
    (linear, nonlinear, "out-of core",sparse)
- Structural Optimization
  - Design Sensitivity
- *Operate on Cray, Paragon, IBM SP-1 and SP-2!*

Langley
Research
Center

References

- Storaasli, O., Nguyen, D., Baddourah, M. and Qin, J.; Computational Mechanics Analysis Tools for Parallel-Vector Supercomputers",*AIAA/ASME/ASCE/AHS/ASC 34th Structures, Structural Dynamics and Materials Conference Proceedings*, Part 2, pp. 772-778, April 1993.
- also *International Journal of Computing Systems in Engineering*, Vol. 4, No. 2-4, 1993 pp. 349-354
- on MOSAIC-WWW (Langley Technical Report Server)

- Questions: O.O.Storaasli@larc.nasa.gov
- Free Videotape from: shuguez@nas.nasa.gov
  (Santa Huguez at 415-604-4632)

Langley
Research
Center