Parallel Numerical Algorithms for Heterogeneous Parallel Computers

Antonio M. Vidal Maciá

(in collaboration with Pedro Alonso, Miguel O. Bernabeu, Victor M. García)

Departamento de Sistemas Informáticos y Computación
Universidad Politécnica de Valencia
Valencia, Spain

Supported by Spanish MCYT and FEDER under Grant TIC 2003-08238-C02-02 and
by the "Programa de Incentivo a la Investigación" of the Universidad Politécnica de Valencia.

Outline

- Introduction
- Heterogeneous distributed memory multicomputers
- The eigenvalue problem to solve
- Classical solutions
- New algorithmic schemes
- Results

Introduction (i)

- Computational problems in signal processing applications:
  - **Implementation of spectral multiresolution analysis/synthesis methods for 3D audio:**
    - Cross-talk cancelers design, Multichannel adaptive filters, Massive multichannel convolutions, ...
  - **Study and evaluation of optimal and quasi-optimal detection algorithms in Multiple Input-Multiple Output (MIMO) communication systems:**
    - Detection algorithms, precodification algorithms, ...
  - **Practical design of passive components for radio communication systems (wireless systems, mobile communication):**
    - BI-RME technique formulation for the accurate and efficient computation of arbitrarily shaped waveguide modes.

Introduction (ii)

- Numerical Linear Algebra addressed problems:
  - To solve structured linear systems (Toeplitz, block-Toeplitz, Toeplitz by blocks, blocks, . . .).
  - To solve structured least squares problems (Toeplitz, block-Toeplitz, Toeplitz by blocks, blocks, . . .).
  - To compute generalized and ordinary eigenvalues and eigenvectors (some or all) of structured matrices.

Introduction (iii)

- Requirements
  - Large and structured matrices.
  - Conventional computers or clusters of PCs.
  - Current libraries (LAPACK, ScaLAPACK) don't provide good performance.
  - Parallel computing must be used with some caution.
  - Heterogeneous parallel computing can be a solution.
- Consequences
  - Methods for computing eigenvalues and eigenvectors must be carefully selected.
  - Algorithms should be restructured.
- **Objective of the presentation**
  - **To analyze methods for solving structured eigenvalue problems on heterogeneous parallel computers.**

Heterogeneous distributed memory multicomputers (i)

- Formally: Set of processors with different computing and communication capabilities that *work together* closely and can be viewed as a single computer.
- Alternative to expensive tightly-coupled supercomputers.
- Great performance-cost ratio.
- Typical scenarios:
  - Clusters of legacy PCs and workstations.
  - LANs of PCs in a university department or company.
  - Homogeneous clusters and supercomputers connected through a LAN.

Heterogeneous distributed memory multicomputers (ii)

- Heterogeneous parallel architectures and numerical linear algebra libraries:
  - There does not exist any numerical linear algebra library specifically designed for heterogeneous parallel architectures.
  - Some authors (Beaumont, Kalinov, Lastovetsky, . . .) have proposed successful techniques to adapt current homogeneous libraries (like ScaLAPACK).
  - Few numerical kernels have been specifically designed for heterogeneous architectures.

Heterogeneous distributed memory multicomputers (iii)

- Our heterogeneous cluster consist of 6 nodes with 22 cores:
  - 1 Intel Pentium IV at 1.6 GHz with 256 KB of L2 cache and 1 GB of RAM
  - 1 Intel Pentium IV at 1.7 GHz with 256 KB of L2 cache and 1 GB of RAM
  - 2 Intel Xeon two-processors at 2.2 GHZ with 512 KB of L2 cache and 4 GB of RAM.
  - 2 Intel Itanium II Montecito four-processors dual-core at 1.4 GHZ with 1 MB of instructions L2 cache and 256 KB of data L2 cache and 8 GB of RAM
- Nodes are linked through a switched Gigabit Ethernet network.

The problem to solve

- An increasing number of real passive waveguide components (filters, multiplexers, . . .) are composed of the cascaded connection of arbitrarily shaped waveguides.

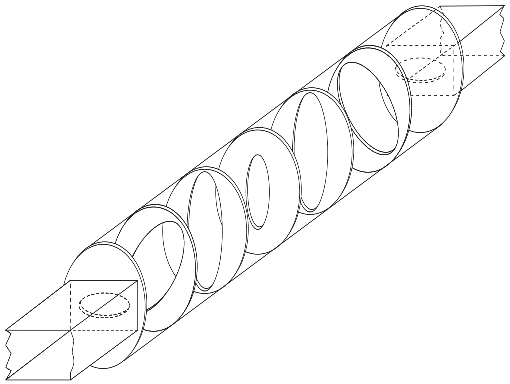

- Different techniques have been proposed for the accurate analysis and design of such components (finite elements method, transmission line matrix, . . .).
- Strong requirements on CPU time and memory storage.

The problem to solve (ii)

- In this work, the modal computation of arbitrary waveguides is based on the Boundary Integral - Resonant Mode Expansion (BI-RME) method [a].
- This technique provides the modal cut-off frequencies of an arbitrary waveguide from the solution of two generalized eigenvalue problems

$$Ax = \lambda Bx$$

with some specific characteristics:

  - Matrices A and B are structured and highly sparse.
  - Only the real positive eigenvalues contained in a $[0, \beta]$ interval are needed.

[a] Conciauro G., Bressan M., Zuffada C.: Waveguide modes via an integral equation leading to a linear matrix eigenvalue problem; IEEE Transactions on Microwave Theory and Techniques. (1984)

The problem to solve (iii)

- Structured matrices *A* and *B* for a ridge waveguide

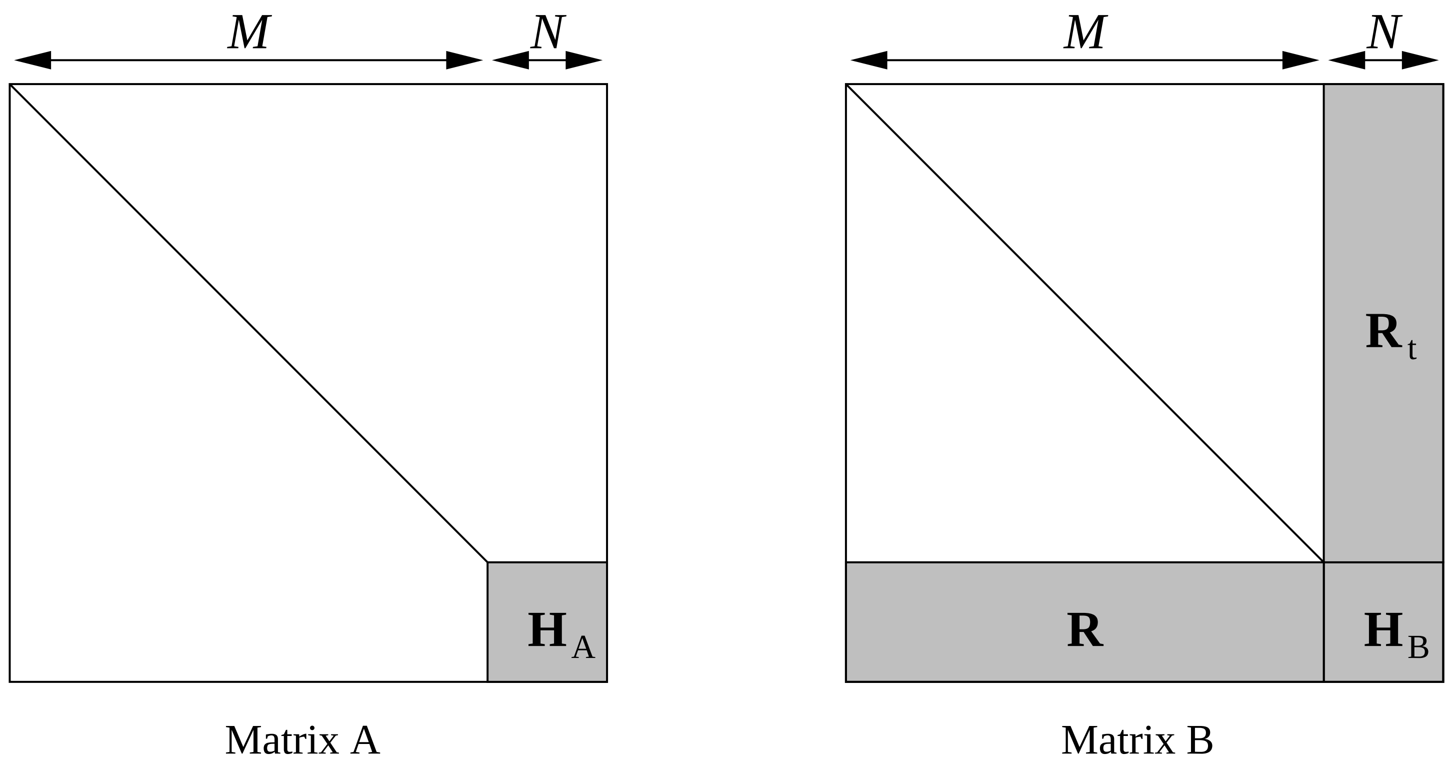

$$M \gg N$$

Classical Approach

- The standard algorithm for generalized eigenvalue problems ($Ax = \lambda Bx$) is the *QZ* algorithm:
  - It is not possible to take advantage of the matrix structure in order to improve its performance.
  - Under certain conditions (symmetric A and symmetric positive definite B) the problem can be transformed into a standard eigenvalue problem ($Cy = \lambda y$).
    - Using the Cholesky or the LDL^T factorization.
  - Once the transformation is done the *QR* iteration or other classic algorithm can be applied.

Classical Approach (ii)

- For a classic eigenvalue algorithm:
  - Its temporal cost is of the form:

$$\alpha + \sum_{i=1}^{n} \beta_i \quad or \quad \alpha + \beta$$

    - $\alpha \equiv$ cost of the matrix tridiagonalization.
    - $\beta_i \equiv$ cost of extracting the i-*th* eigenvalue/eigenvector.
    - $\beta \equiv$ cost of extracting all the eigenvalues/eigenvectors.
  - Properties
    - $\alpha \gg \beta_i$.
    - Parallel tridiagonalization is a highly-coupled parallel problem.
    - Not suitable for structured matrices (filling, structure loss and misuse of the structure for optimization)

New algorithmic scheme

- Our proposal is to implement algorithms for heterogeneous parallel computers, which temporal cost is of the form:

$$\delta + \sum_{i=1}^{m} \varepsilon_i$$

  - $\delta \equiv$ cost of splitting the problem into m independent sub-problems.
  - $\varepsilon_i \equiv$ cost of solving the i-*th* sub-problem sequentially.
- Properties
  - $\delta \ll \varepsilon_i$.
  - $\forall\, i, j:\ \varepsilon_i \simeq \varepsilon_j$
  - Algorithms should take advantage of the structure of the matrices (if any).

New algorithmic scheme applied to eigenproblems

- **We propose to implement a modified version of the Lanczos' algorithm for the solution of eigenproblems in heterogeneous multicomputers.**
- Splitting of the original problem: based on spectrum partitioning.
  - $\lambda(C)$: the set of all the eigenvalues of C (spectrum).
  - An upper and a lower bound (*lb* and *ub*) of the set can be computed by means of the Gershgorin Circle Theorem.

$$\lambda_i \in \lambda(C) \rightarrow \lambda_i \in [lb, ub]$$

  - The idea is to partition [*lb*,*ub*] into m subsets containing the same number of eigenvalues (approx.).

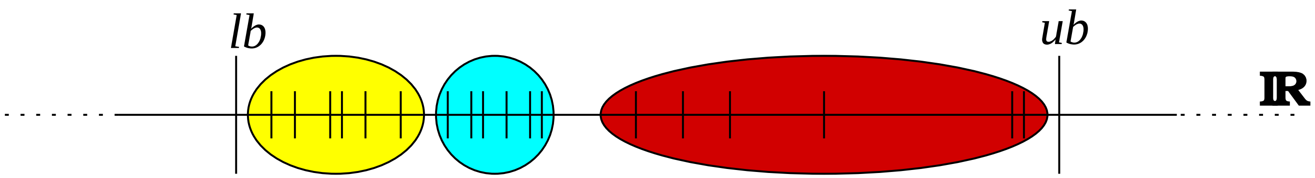

New algorithmic scheme applied to eigenproblems (ii)

- Partitioning [*lb*,*ub*]: Inertia Theorem
  - Let $L^{\alpha}D^{\alpha}L_t^{\alpha}$ and $L^{\beta}D^{\beta}L_t^{\beta}$ be the LDL_t decomposition of $A-\alpha B$ and $A-\beta B$, respectively.
  - The number of eigenvalues in $[\alpha,\beta]$ is

$$\nu(D^{\beta})-\nu(D^{\alpha}),$$

    where $\nu(D)$ denotes the number of negative elements in the diagonal D.
  - LDL_t decompositions can be computed with a moderated cost, taking profit from the structure of the matrices.
- Based on the Inertia and the Gershgorin circle theorem we have developed a bisection-like algorithm that performs the spectrum partitioning.

New algorithmic scheme applied to eigenproblems (iii)

- Solving the sub-problems: the “Shift-and-Invert” version of the Lanczos’ method
  - Basic Lanczos’ algorithm allows the computation of a reduced number of extremal eigenvalues (largest or smallest in magnitude).
  - Given a real number σ (the shift), Lanczos’ algorithm can be applied to the matrix

$$W = (A - \sigma B)^{-1} B$$

    to extract the eigenvalues of the original problem closer to the shift σ.
  - This variation requires the solution of several linear systems, with $A - \sigma B$ as coefficient matrix.
  - System solution cost can be reduced taking profit from the structure of the matrices.

Parallelization of the algorithmic scheme

- The parallelization of the previous algorithm is quite straightforward:
  1. Apply the bisection-like algorithm to divide the original problem into m sub-problems.
  2. Distribute the sub-problems among the p available processors and solve them sequentially.
- The way the sub-problems are distributed will determine the work-load balance of the algorithm.
  - Statically: processor P_i gets a number of sub-problems proportional to its *relative power*.
  - Dynamically: sub-problems are assigned on demand to the processors (master-slave).

Results

- We have implemented the previous parallel algorithm to solve the waveguide analysis problem described before.
- In addition we have implemented it for other kinds of structured matrices:
  - Toeplitz
  - Tridiagonal
- Note that all of them imply the development of linear system solvers optimized for the matrix structure.
- Several publications have been produced:

1. V.M.García, A.Vidal, V.E.Boria and A.M.Vidal, Efficient and accurate waveguide mode computation using BI-RME and Lanczos methods. INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN ENGINEERING. 2006; 65:1773
2. A.M.Vidal, A.Vidal, V.E.Boria and V.M.García, Parallel computation of arbitrarily shaped waveguide modes using BI-RME and Lanczos methods. COMMUNICATIONS IN NUMERICAL METHODS IN ENGINEERING. 2006; 23-4:273-284

Results (ii)

3 Miguel Oscar Bernabeu, Mariam Taroncher, Víctor M. Garcia, Ana Vidal: Parallel Implementation in PC Clusters of a Lanczos-based Algorithm for an Electromagnetic Eigenvalue Problem. ISPDC 2006: 296-300

4 Miguel O. Bernabeu, Antonio M. Vidal: The symmetric Tridiagonal Eigenvalue Problem: a Heterogeneous Parallel Approach. WSEAS TRANSACTIONS ON MATHEMATICS. 2007; 4-6: 587-594

5 Antonio M. Vidal, Víctor M. García, Pedro Alonso, Miguel O. Bernabeu: Parallel Computation of the Eigenvalues of Symmetric Toeplitz Matrices through Iterative Methods. JOURNAL OF PARALLEL AND DISTRIBUTED COMPUTING. Under revision.

6 P.Alonso and J.M.Badía and A. M.Vidal, An Efficient and Stable Parallel Solution for Non-Symmetric Toeplitz Linear Systems, LNCS 3402:685-692, 2005.

7 P.Alonso and J.M.Badía and A. M.Vidal, An Efficient Parallel Algorithm to Solve Block-Toeplitz systems, The Journal of Supercomputing 32:251-278, 2005.

8 P.Alonso and A.L.Lastovetsky and A.M.Vidal, A Parallel Algorithm for the Solution of the Deconvolution Problem on Heterogeneous Networks, HeteroPar’06: Fifth International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Networks, IEEE, online, 2006.

Results (iii)

- Some conclusions extracted from the previous citations:
  - The method parallelizes extremely well, achieving close to optimum speedups [5]:

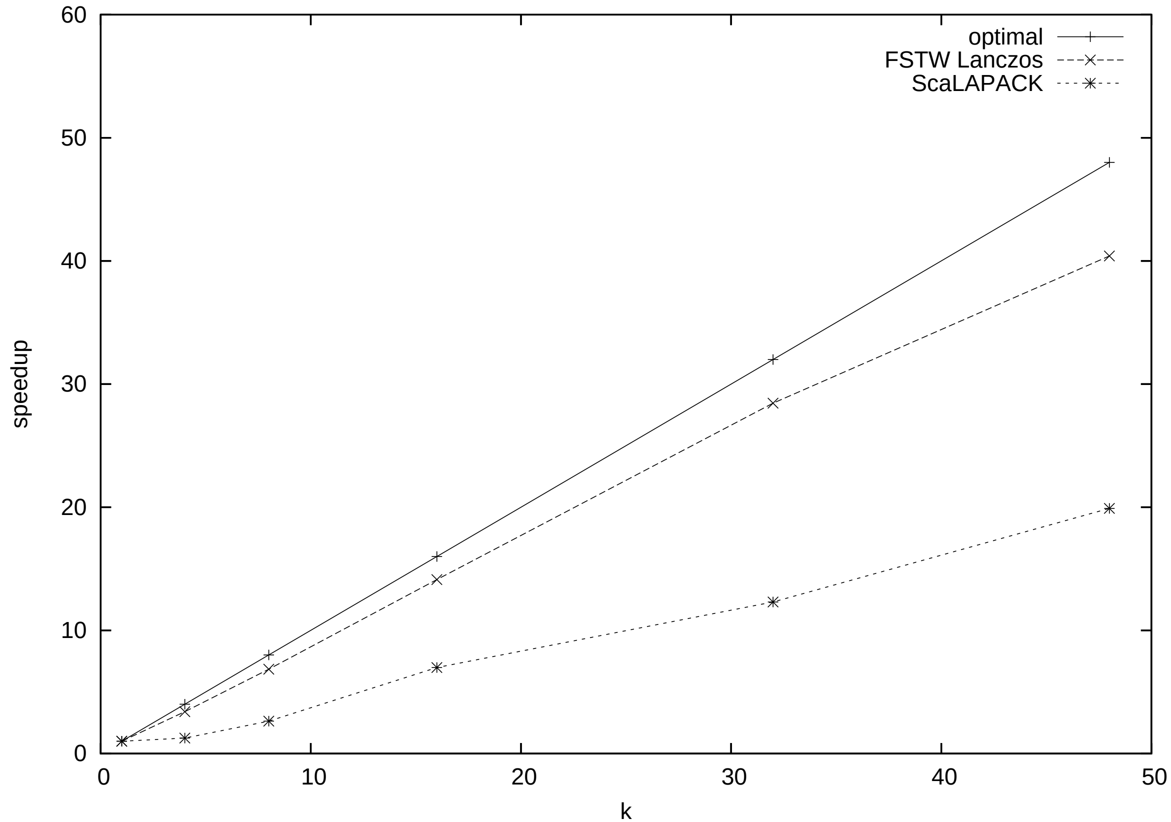

Scaled speedup of FSTW Lanczos' parallel algorithm[5] solving Toeplitz Eigenproblems.

Results (iv)

- Some conclusions extracted from the previous citations:
  - Due to the optimal use of the structure of matrices, our implementations can solve larger problems that current libraries (LAPACK, ScaLAPACK) cannot [2].
  - Based on the cost model $\delta + \sum_{i=1}^{m} \varepsilon_i$ of our parallel algorithm [4]:
    - If $\forall\, i, j : \varepsilon_i \simeq \varepsilon_j$ both static and dynamic work-load balance algorithms achieve good performance.
    - If $\exists\, i, j : \varepsilon_i \gg \varepsilon_j$ only the dynamic algorithm can ensure a correct work-load balance.
    - These situations will depend on the distribution of the eigenvalues along the spectrum (uniform distribution, clusters of eigenvalues, hidden eigenvalues, ...)

Parallel Numerical Algorithms for Heterogeneous Parallel Computers

Antonio M. Vidal Maciá

(in collaboration with Pedro Alonso, Miguel O. Bernabeu, Victor M. García)

Departamento de Sistemas Informáticos y Computación
Universidad Politécnica de Valencia
Valencia, Spain

Supported by Spanish MCYT and FEDER under Grant TIC 2003-08238-C02-02 and
by the "Programa de Incentivo a la Investigación" of the Universidad Politécnica de Valencia.