RECSY —
A High Performance Library for Sylvester-Type Matrix Equations

Isak Jonsson, Bo Kågström

Department of Computing Science
Umeå University, SE-901 87 Umeå, Sweden

{isak,bokg}@cs.umu.se

Euro-Par 2003, Klagenfurt, Austria
August 26--29, 2003

Outline

- Blocked algorithms and memory hierarchies.
- Sylvester-type matrix equations.
- Recursive blocked solvers for triangular Sylvester-type matrix equations.
- Implementation issues.
- The RECSY library.
- Questions.

Deep Memory Hierarchies

Architecture Evolution: HPC systems with multiple SMP nodes, several levels of caches and more functional units per CPU.

Key to performance: Understand the algorithm and architecture interaction.

Goal: Maintain 2-dim data locality at every level of the 1-dim tiered memory structure.

- Hierarchical blocking.
- Matching an algorithm and its data structure.

1–dim ↔ 2–dim: Blocking

1. Explicit multi-level blocking

- Each loop set matches a specific level of the memory hierarchy.
- Deep knowledge of architecture characteristics needed.
- Needs a blocking parameter for each level.
- Two-level blocked matrix multiply (tuned for L1 and L2 cache).

2. Automatic blocking via recursion

- Recursion: key concept for matching an algorithm and its data structure.
- Recursive algorithms – divide and conquer style.
- Automatic hierarchical blocking – variable and "squarish".
- Only tuning parameter is L1 cache.

Sylvester-Type Matrix Equations

Appear in different control theory applications: stability problems, model reduction, balancing, H_∞ control.

One-sided:

- Sylvester (SYCT): $AX - XB = C$, A, B and C general
- Lyapunov (LYCT): $AX + XA^T = C$, A general, $C = C^T$ (semi)definite
- Generalized (coupled) Sylvester (GCSY):

$$\begin{aligned} AX - YB &= C \\ DX - YE &= F \end{aligned}$$

Two-sided:

- Discrete Sylvester (SYDT): $AXB^T - X = C$
- Discrete Lyapunov (or Stein) (LYDT): $AXA^T - X = C$
- Generalized Sylvester (GSYL): $AXB^T - CXD^T = F$
- Generalized Lyapunov

$$\begin{aligned} &\text{(GLYCT):}\ \ AXE^T + EXA^T = F \\ &\text{(GLYDT):}\ \ AXA^T - EXE^T = F \end{aligned}$$

Sylvester-Type Matrix Equations

Bartels-Stewart-type of algorithms $\Longrightarrow$
second major step in the solution is to solve a triangular matrix equation.

Our blocked recursive technique works for all! Here

- triangular discrete-time Sylvester and Lyapunov
- triangular generalized Sylvester and Lyapunov

Great source of triangular matrix equation problems from:

- Condition estimation
  - of the matrix equations themselves,
  - and in various eigenspace problems including reordering of eigenvalues.
- Computing functions of matrices.

Recursive Triangular Sylvester Solvers

$op(A) \cdot X \pm X \cdot op(B) = \beta \cdot C$, $C \leftarrow X \ (M \times N)$, where
$A(M \times M)$ and $B(N \times N)$ upper quasi-triangular.

$transA =$'N', $transB =$'N', $sign = -$, $\beta = 1$:

Case 1 ($1 \leq N \leq M/2$): Split A and C (by rows)

$$\begin{bmatrix} A_{11} & A_{12} \\ & A_{22} \end{bmatrix} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix} - \begin{bmatrix} X_1 \\ X_2 \end{bmatrix} B = \begin{bmatrix} C_1 \\ C_2 \end{bmatrix}$$

$$\begin{aligned} A_{11}X_1 - X_1B &= C_1 - A_{12}X_2 \\ A_{22}X_2 - X_2B &= C_2 \end{aligned}$$

1. SYLV('N', 'N', A_{22}, B, C_2)
2. GEMM('N', 'N', $\alpha = -1$, A_{12}, C_2, C_1)
3. SYLV('N', 'N', A_{11}, B, C_1)

Case 2 ($1 \leq M \leq N/2$): Split B and C (by columns)

Recursive Triangular Sylvester Solvers

Case 3 ($N/2 < M < 2N$): Split A, B and C

$$\begin{bmatrix} A_{11} & A_{12} \\ & A_{22} \end{bmatrix} \begin{bmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{bmatrix} - \begin{bmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{bmatrix} \begin{bmatrix} B_{11} & B_{12} \\ & B_{22} \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix}$$

$$\begin{aligned} A_{11}X_{11} - X_{11}B_{11} &= C_{11} - A_{12}X_{21} \\ A_{11}X_{12} - X_{12}B_{22} &= C_{12} - A_{12}X_{22} + X_{11}B_{12} \\ A_{22}X_{21} - X_{21}B_{11} &= C_{21} \\ A_{22}X_{22} - X_{22}B_{22} &= C_{22} + X_{21}B_{12} \end{aligned}$$

1. SYLV('N', 'N', A_{22}, B_{11}, C_{21})

2a. GEMM('N', 'N', $\alpha = +1$, C_{21}, B_{12}, C_{22})

2b. GEMM('N', 'N', $\alpha = -1$, A_{12}, C_{21}, C_{11})

3a. SYLV('N', 'N', A_{22}, B_{22}, C_{22})

3b. SYLV('N', 'N', A_{11}, B_{11}, C_{11})

4. GEMM('N', 'N', $\alpha = -1$, A_{12}, C_{22}, C_{12})

5. GEMM('N', 'N', $\alpha = +1$, C_{11}, B_{12}, C_{12})

6. SYLV('N', 'N', A_{11}, B_{22}, C_{12})

Operations 2a, 2b can be executed in parallel, as well as Operations 3a, 3b.

Implementation issues

Two alternatives for doing the recursive splits:

1. Always split the largest dimension in two (Cases 1 and 2).
2. Split both dimensions simultaneously (Case 3) when the dimensions are within a factor 2 from each other.

2. $\Longrightarrow$ a shorter but wider recursion tree, which offers more "parallel tasks".

Use of BLAS routines

The recursive approach gives algorithms that call level 3 BLAS routines (GEMM, TRSM, SYRK, etc.) with square blocks. This enables the best performance from (blocked) BLAS routines.

Sylvester Kernel

For problems smaller than the block size, the dimension is split in two until the problem is of size 2×2 – 4×4, when subsystems are solved using

$$(B \otimes A - I_n \otimes I_m)\mathsf{vec}(X) = \mathsf{vec}(C)$$

Impact of Kernel Solvers

Kernel solvers execute $O(N^2)$ flops out of the total $O(N^3)$.
Model of the *overall computation rate*:

$$S = \frac{\text{Total number of flops}}{\text{GEMM time} + \text{Kernel time}} = \frac{2N^3}{\frac{2N^3-4N^2}{G} + \frac{4N^2}{K}}$$

where G = DGEMM perf., and K = kernel solver perf.

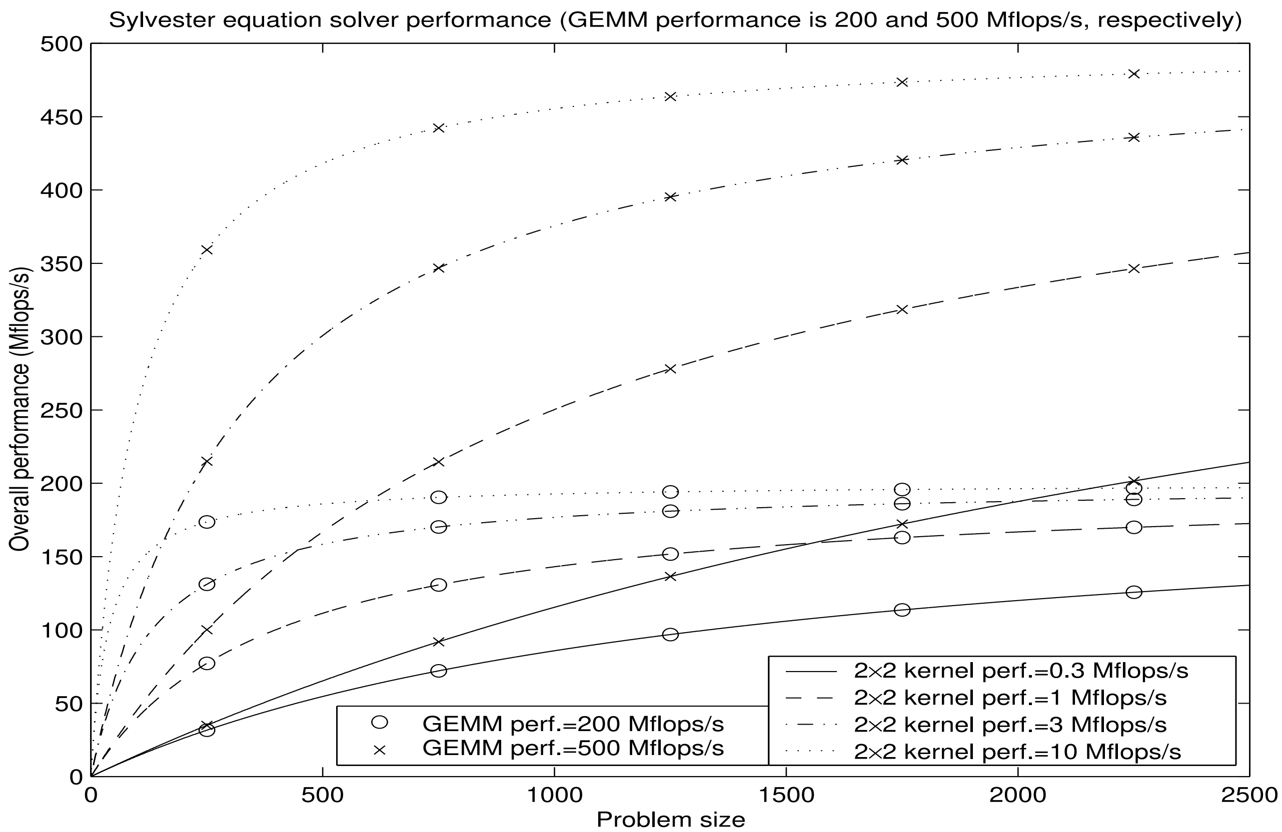

Impact of Kernel Solvers

Modelling results:

- N = 2500; G = 500, K = 0.3 Mflops/s $\Longrightarrow$
  Overall performance is at most approaching 50% of GEMM peformance.
- N = 2500; G = 500, K = 3.0 or 10 Mflops/s $\Longrightarrow$
  Overall performance rather quickly approaches 80–90% of GEMM performance.
- N = 750; G = 500, K = 0.3 Mflops/s $\Longrightarrow$
  Kernel flops = 0.3%, Kernel time = 80%
  (e.g. LAPACK DTGSY2 on modern RISC/CISC processors)
  DTGSY2 designed with the primary goal of producing high-accuracy results and signaling ill-conditioning (complete pivoting, overflow guarding).

Trade-off between ROBUSTNESS, RELIABILITY AND SPEED.

Kernel performance decisive for the overall performance of matrix equation solver.

Optimized Superscalar Kernels

Our design approach:

- One single routine to solve a kernel problem using Kronecker product matrix representation $\Longrightarrow$
  great potential for register reuse.
- LU with partial pivoting and overflow guarding, forward and backward substitutions using complete loop unrolling.
- Matrix multiply kernel "lite" using register blocking techniques (complete outer-loop unrolling and fusion).

Facilitates for the compiler to look ahead and schedule the resulting assembly instructions optimally.

Two-sided Matrix Equations
Cost and Execution Order

Recursive blocked algorithms require both extra workspace (two-sided matrix products) and more flops compared to the standard elementwise algorithms.

Matrix equation	Overall cost in flops	Flop ratio ($M = N$)
SYDT	$\frac{6}{4}M^2N + 2MN^2 \quad (M \leq N)$	1.1667
	$2M^2N + \frac{6}{4}MN^2 \quad (M > N)$	
LYDT	$\frac{25}{12}N^3$	1.5625
GSYL	$3M^2N + 4MN^2 \quad (M \leq N)$	1.1667
	$4M^2N + 3MN^2 \quad (M > N)$	
GLYCT	$\frac{21}{6}N^3$	1.3125
GLYDT	$\frac{25}{6}N^3$	1.5625

Outperform the standard algorithms for large enough problems—due to GOOD DATA LOCALITY.

DATA REFERENCE PATTERNS: Order in which they access data and how big chunks and how many times the data is moved in the memory hierarchy. CRITICAL!

SMP Parallelism

SMP parallelism is obtained by solving independent equations as different OpenMP sections.
Recursion:

- Shows what parts can be solved in parallel.
- Creates problems that are large ⇒ coarse granularity.

Also, due to the coarse granularity, SMP versions of DGEMM run well.

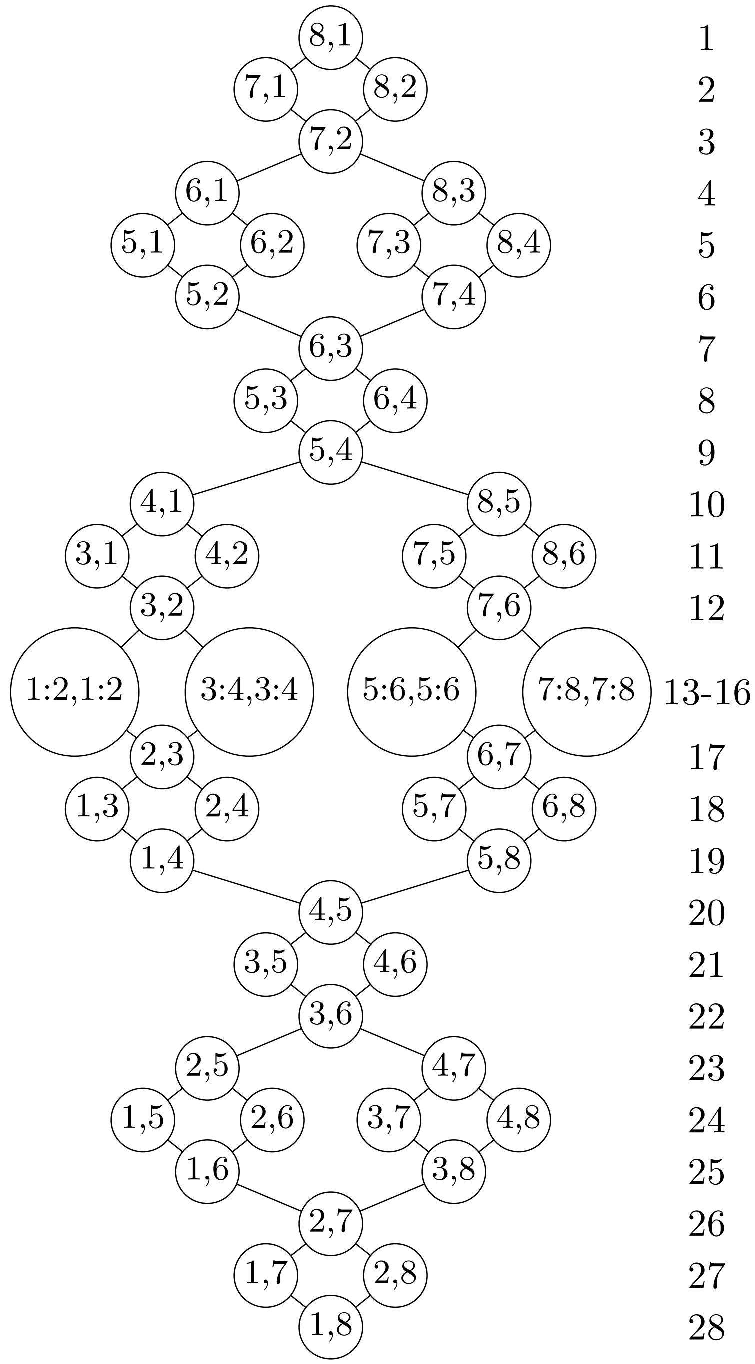

Coupled Sylvester – SMP Performance

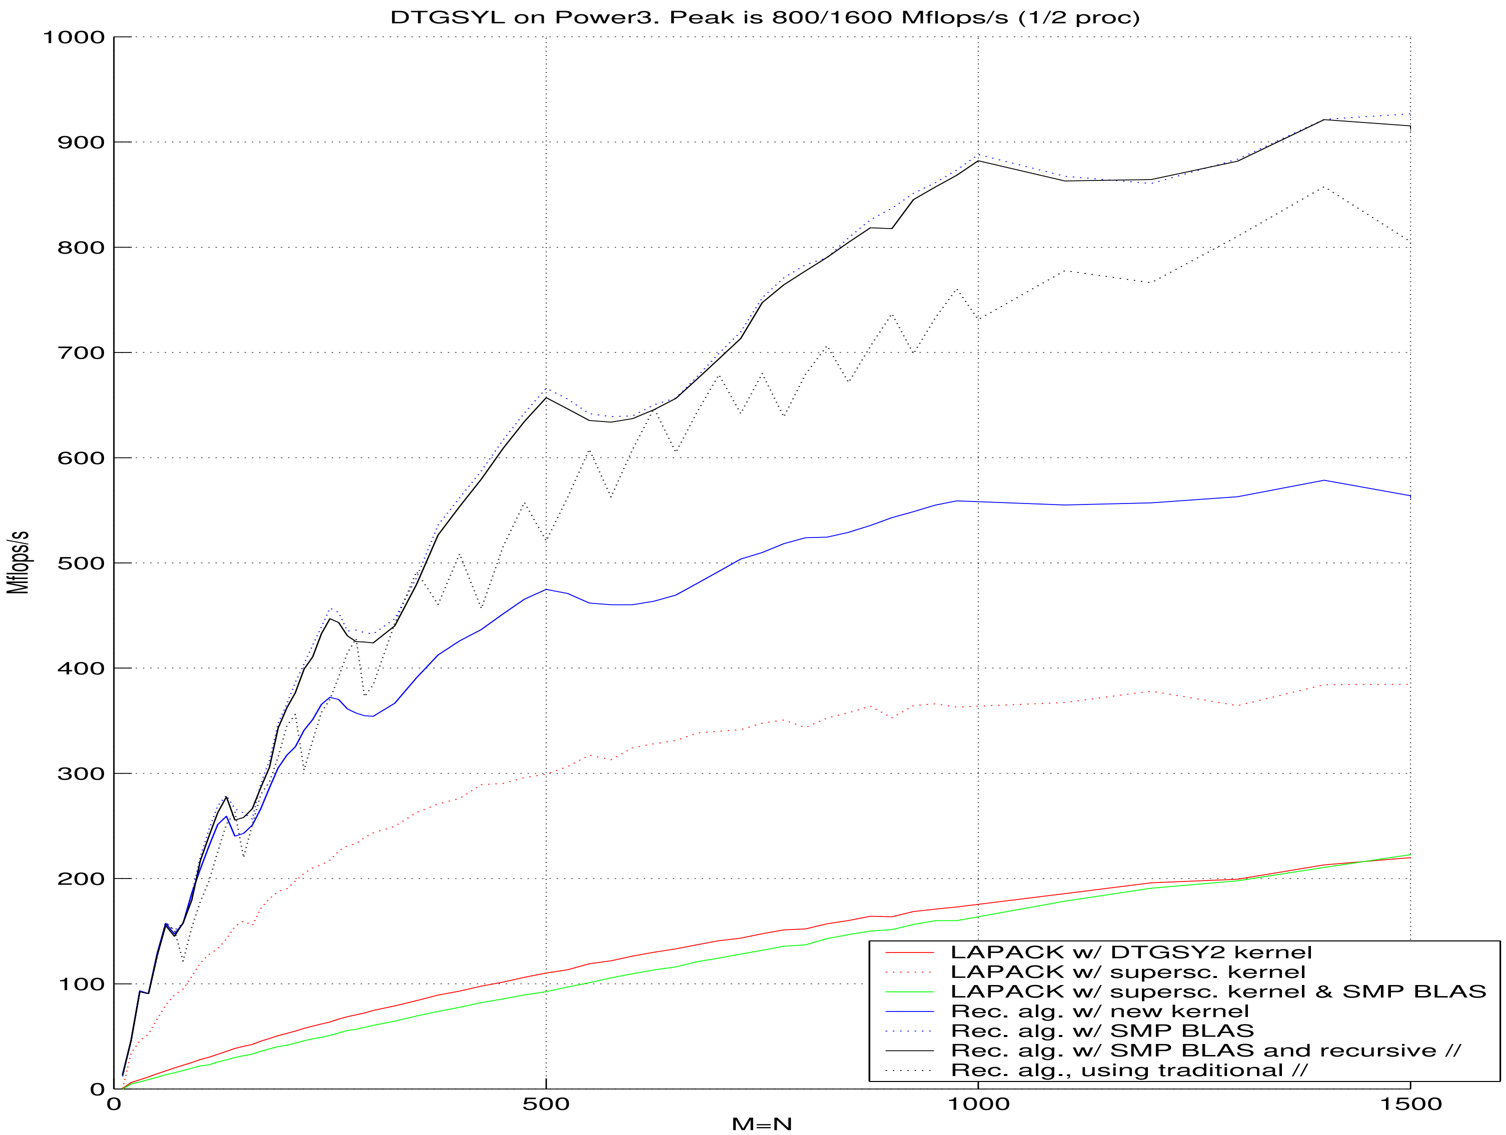

TRIANGULAR $AXB^T - X = C$ (SYDT)
IBM POWER 3

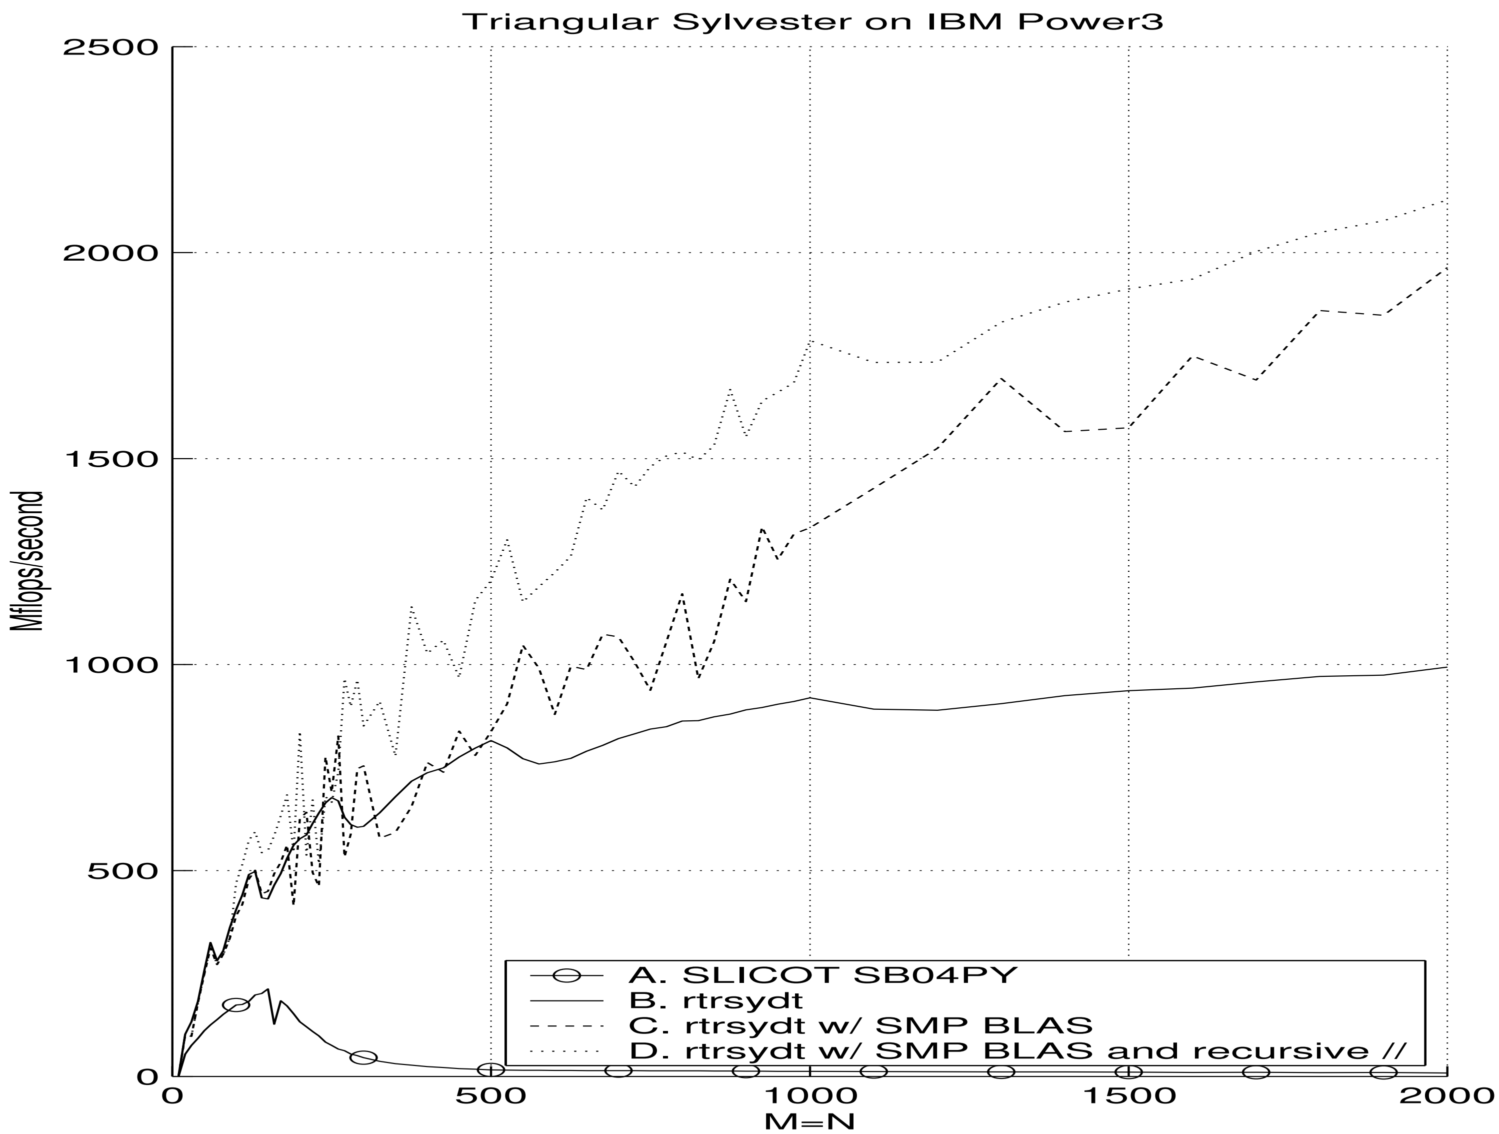

RECSY

- A library that encompasses all eight mentioned matrix equations, with serial and in SMP parallel versions.
- Recursion is done using Fortran 90 `RECURSIVE SUBROUTINES`.
- Extra memory buffered are dynamically allocated, or can be provided by the user.
- SMP versions are available on all platforms with OpenMP compilers.
- F77 Wrappers for SLICOT and LAPACK routines provided with library. No need to recompile "legacy" code, only relink.
- Fall-back routines provide the same accuracy and stability as the replaced routines.
- Source publicly available at `http://www.cs.umu.se/~isak/recsy` for review.

RECSY – Uniprocessor Routines

- `RECSYCT(UPLOSIGN, SCALE, M, N, A, LDA, B, LDB, C, LDC, INFO, MACHINE )` (LAPACK: DTRSYL)
- `RECLYCT(UPLO, SCALE, M, A, LDA, C, LDC, INFO, MACHINE )` (SLICOT: SB03MY)
- `RECGCSY(UPLOSIGN, SCALE, M, N, A, LDA, B, LDB, C, LDC, D, LDD, E, LDE, F, LDF, INFO, MACHINE )` (LAPACK: DTGSYL)
- `RECSYDT(UPLOSIGN, SCALE, M, N, A, LDA, B, LDB, C, LDC, INFO, MACHINE, WORKSPACE, WKSIZE )` (SLICOT: SB04PY)
- `RECLYDT(UPLO, SCALE, M, A, LDA, C, LDC, INFO, MACHINE, WORKSPACE, WKSIZE )` (SLICOT: SB03MX)
- `RECGSYL(UPLOSIGN, SCALE, M, N, A, LDA, B, LDB, C, LDC, D, LDD, E, LDE, INFO, MACHINE, WORKSPACE, WKSIZE )` (No equivalent in SLICOT or LAPACK!)
- `RECGLYDT(UPLO, SCALE, M, A, LDA, E, LDE, C, LDC, INFO, MACHINE, WORKSPACE, WKSIZE )` (SLICOT: SG03AX)
- `RECGLYCT(UPLO, SCALE, M, A, LDA, E, LDE, C, LDC, INFO, MACHINE, WORKSPACE, WKSIZE )` (SLICOT: SG03AY)

RECSY – Multiprocessor Routines

- RECSYCT_P(PROCS, UPLOSIGN, SCALE, M, N, A, LDA, B, LDB, C, LDC, INFO, MACHINE)
- RECLYCT_P(PROCS, UPLO, SCALE, M, A, LDA, C, LDC, INFO, MACHINE)
- RECGCSY_P(PROCS, UPLOSIGN, SCALE, M, N, A, LDA, B, LDB, C, LDC, D, LDD, E, LDE, F, LDF, INFO, MACHINE)
- RECSYDT_P(PROCS, UPLOSIGN, SCALE, M, N, A, LDA, B, LDB, C, LDC, INFO, MACHINE, WORKSPACE, WKSIZE)
- RECLYDT_P(PROCS, UPLO, SCALE, M, A, LDA, C, LDC, INFO, MACHINE, WORKSPACE, WKSIZE)
- RECGSYL_P(PROCS, UPLOSIGN, SCALE, M, N, A, LDA, B, LDB, C, LDC, D, LDD, E, LDE, INFO, MACHINE, WORKSPACE, WKSIZE)
- RECGLYDT_P(PROCS, UPLO, SCALE, M, A, LDA, E, LDE, C, LDC, INFO, MACHINE, WORKSPACE, WKSIZE)
- RECGLYCT_P(PROCS, UPLO, SCALE, M, A, LDA, E, LDE, C, LDC, INFO, MACHINE, WORKSPACE, WKSIZE)

Unreduced Two-sided: $AXA^T - EXE^T = C$
Solving and Sep[GLYDT]-estimaton

a) N	SG03AD using SG03AX Total time	Solver	SG03AD using RECSY Total time	Solver	S
50	0.0277	49.9 %	0.0185	20.1 %	1.50
100	0.180	51.2 %	0.0967	9.0 %	1.86
250	2.89	46.8 %	1.62	4.7 %	1.79
500	59.0	42.3 %	34.5	1.5 %	1.71
750	303.4	42.0 %	177.5	0.9 %	1.71
1000	646.6	44.6 %	361.8	1.0 %	1.79
50	0.117	87.6 %	0.0263	45.6 %	4.44
100	0.709	87.3 %	0.152	40.6 %	4.68
250	9.98	84.5 %	2.08	25.4 %	4.81
500	178.6	80.9 %	37.8	9.4 %	4.73
750	924.1	80.9 %	184.4	4.5 %	5.01
1000	2076.6	82.7 %	391.8	8.4 %	5.30

We get another 2x speedup ($N > 500$) by replacing LAPACK routines DGGHRD and DHGEQZ by Dackland-Kågström's blocked Hessenberg-triangular reduction and QZ algorithms (ACM TOMS'99) for transforming (A, E) to generalized Schur form.

Conclusions – so far

- State-of-the-art HPC systems have **deep memory hierarchies**.
- **Recursion** efficiently provides automatic variable blocking for all levels of the memory hierarchy.
- Recursive blocking $\Longrightarrow$ Temporal locality
- Our recursive blocked implementations with optimized kernels
  - are GEMM-rich, and
  - show significant performance improvements (30% to 400+%).
- Code at `http://www.cs.umu.se/~isak/recsy`
  - Uses F90 for recursion, dynamic memory allocation
  - Uses (nested) OpenMP for SMP parallelism
  - Overloads LAPACK and SLICOT routines for Sylvester-type matrix equations