

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 caches and more functional units per CPU.

KEY TO PERFORMANCE: Understand the algorithm and architecture.

GOAL: Maintain 2-dim data locality at every level of the 1-dim tree.

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

1-DIM \leftrightarrow 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 d
- Recursive algorithms – divide and conquer style.
- Automatic HIERARCHICAL BLOCKING – variable and “suaris
- Only tuning parameter is L1 cache.

SYLVESTER-TYPE MATRIX EQUATIONS

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

ONE-SIDED:

- Sylvester (SYCT): $AX - XB = C$, A , B and C general
- Lyapunov (LYCT): $AX + XA^T = C$, A general, $C = C^T$ (symmetric)
- 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

$$\text{(GLYCT): } AXE^T + EXA^T = F$$

$$\text{(GLYDT): } AXA^T - EXE^T = F$$

SYLVESTER-TYPE MATRIX EQUATIONS

Bartels-Stewart-type of algorithms \implies
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 EQUATIONS**
- **TRIANGULAR GENERALIZED SYLVESTER AND LYAPUNOV EQUATIONS**

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 S

$op(A) \cdot X \pm X \cdot op(B) = \beta \cdot C$, $C \leftarrow X$ ($M \times N$), 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)

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

$$A_{11}X_1 - X_1B = C_1 - A_{12}X_2$$

$$A_{22}X_2 - X_2B = C_2$$

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 S

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

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

$$A_{11}X_{11} - X_{11}B_{11} = C_{11} - A_{12}X_{21}$$

$$A_{11}X_{12} - X_{12}B_{22} = C_{12} - A_{12}X_{22} + X_{11}B_{12}$$

$$A_{22}X_{21} - X_{21}B_{11} = C_{21}$$

$$A_{22}X_{22} - X_{22}B_{22} = C_{22} + X_{21}B_{12}$$

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 Operation

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 dimension is a factor 2 from each other.
2. \implies a shorter but wider recursion tree, which offers more “parallelism”

USE OF BLAS ROUTINES

The recursive approach gives algorithms that call level 3 BLAS routines (e.g., SYRK, etc.) with square blocks. This enables the best performance by using the BLAS routines.

SYLVESTER KERNEL

For problems smaller than the block size, the dimension is split in two. For example, if the dimension is of size $2 \times 2 - 4 \times 4$, when subsystems are solved using

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

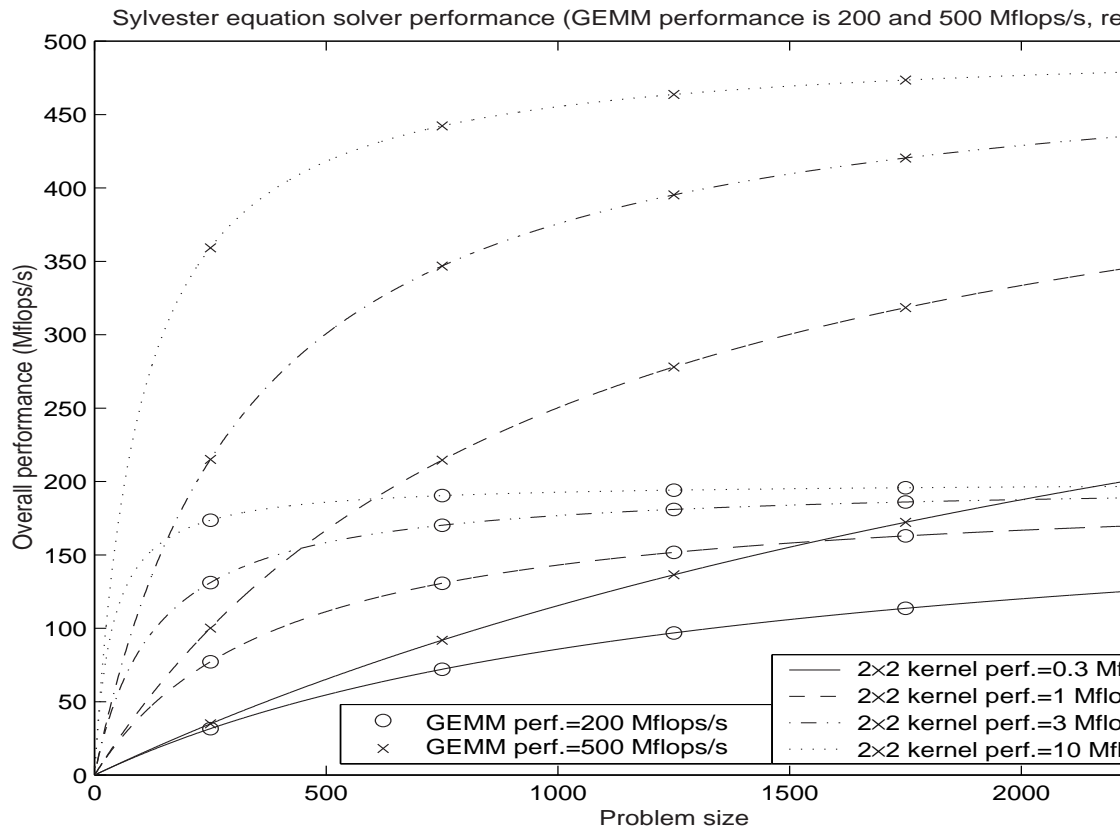
IMPACT OF KERNEL SOLVERS

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 $G = \text{DGEMM perf.}$, and $K = \text{kernel solver perf.}$



IMPACT OF KERNEL SOLVERS

Modelling results:

- $N = 2500$; $G = 500$, $K = 0.3$ Mflops/s \implies
Overall performance is **at most approaching 50%** of GEMM performance
- $N = 2500$; $G = 500$, $K = 3.0$ or 10 Mflops/s \implies
Overall performance rather quickly **approaches 80–90%** of GEMM performance
- $N = 750$; $G = 500$, $K = 0.3$ Mflops/s \implies
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 solutions while
signaling ill-conditioning (complete pivoting, overflow guarding)

Trade-off between **ROBUSTNESS, RELIABILITY AND SPEED**.

Kernel performance decisive for the overall performance of matrix solvers

OPTIMIZED SUPERSCALAR KERNEL

Our design approach:

- One **single routine** to solve a kernel problem using Kronecker product representation \implies **great potential for register reuse.**
- LU with partial pivoting and overflow guarding, forward and backward substitution using **complete loop unrolling.**
- **Matrix multiply kernel “lite”** using **register blocking techniques** (e.g., loop unrolling and fusion).

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

TWO-SIDED MATRIX EQUATIONS COST AND EXECUTION ORDER

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

Matrix equation	Overall cost in flops	Flop ratio ($M = N$)
SYDT	$\frac{6}{4}M^2N + 2MN^2$ ($M \leq N$)	1.166
	$2M^2N + \frac{6}{4}MN^2$ ($M > N$)	
LYDT	$\frac{25}{12}N^3$	1.562
GSYL	$3M^2N + 4MN^2$ ($M \leq N$)	1.166
	$4M^2N + 3MN^2$ ($M > N$)	
GLYCT	$\frac{21}{6}N^3$	1.312
GLYDT	$\frac{25}{6}N^3$	1.562

Outperform the standard algorithms for large enough problems—**LOCALITY**.

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

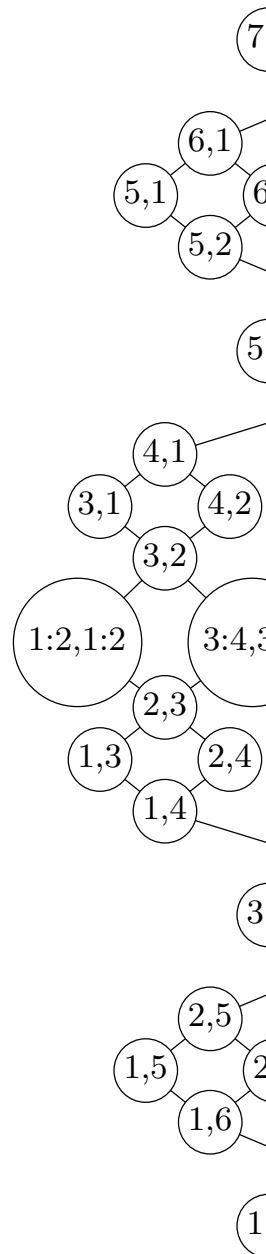
SMP PARALLELISM

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

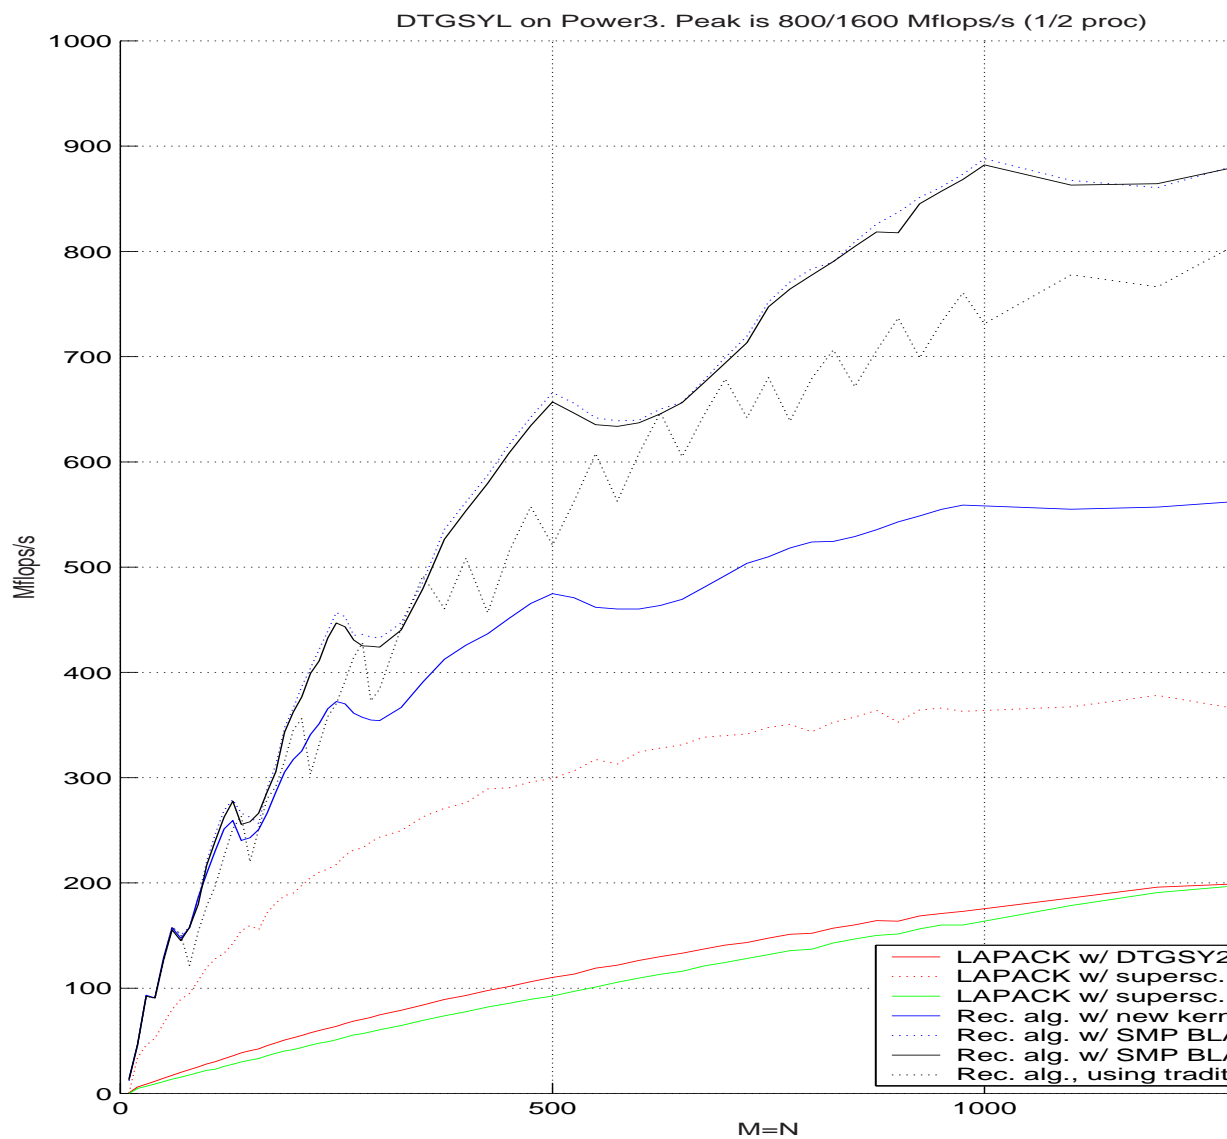
Recursion:

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

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

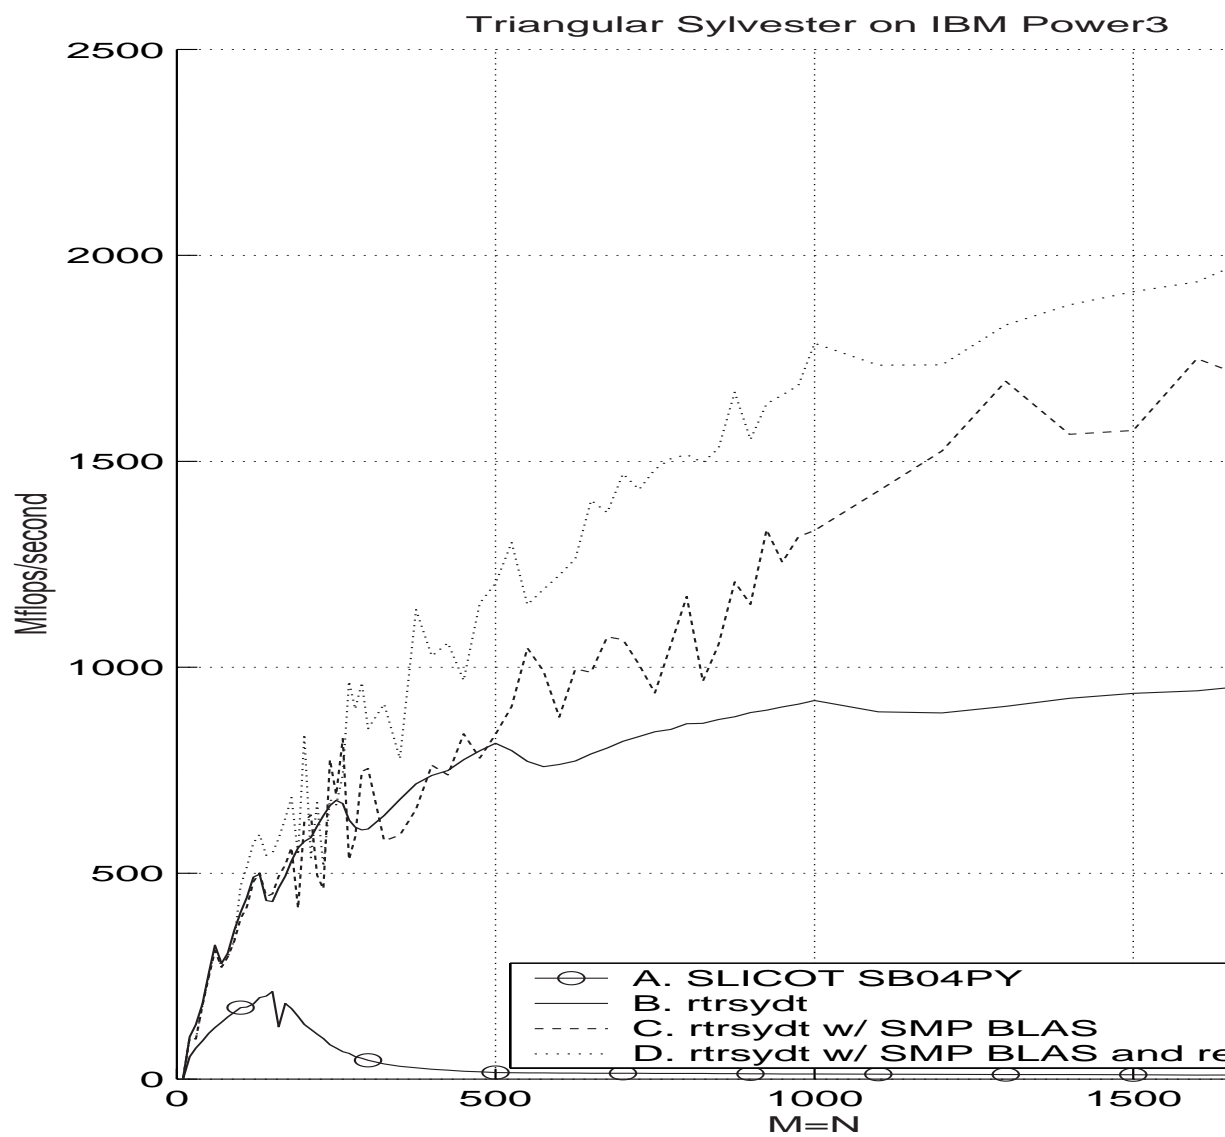


COUPLED SYLVESTER – SMP PERFORM



TRIANGULAR $AXB^T - X = C$ (SYD)

IBM POWER 3



RECSY

- A library that encompasses all eight mentioned matrix equations and their SMP parallel versions.
- Recursion is done using Fortran 90 `RECURSIVE SUBROUTINES`.
- Extra memory buffered are dynamically allocated, or can be pre-allocated.
- SMP versions are available on all platforms with OpenMP compiler.
- F77 Wrappers for SLICOT and LAPACK routines provided with wrappers. To use, recompile “legacy” code, only relink.
- Fall-back routines provide the same accuracy and stability as the main routines.
- Source publicly available at <http://www.cs.umu.se/~isak/>

RECSY – UNIPROCESSOR ROUTINE

- RECSYCT(UPLOSIGN, SCALE, M, N, A, LDA, B, LDB, C, LDC, INFO, WORKSPACE, WKSIZE) (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, L, INFO, MACHINE) (LAPACK: DTGSYL)
- RECSYDT(UPLOSIGN, SCALE, M, N, A, LDA, B, LDB, C, LDC, INFO, 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, L, 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 ROUT

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

UNREDUCED TWO-SIDED: $AXA^T - EX$ SOLVING AND SEP[GLYDT]-ESTIMATION

a)	SG03AD using SG03AX		SG03AD using RECO	
N	Total time	Solver	Total time	Solver
50	0.0277	49.9 %	0.0185	20.7 %
100	0.180	51.2 %	0.0967	9.0 %
250	2.89	46.8 %	1.62	4.7 %
500	59.0	42.3 %	34.5	1.5 %
750	303.4	42.0 %	177.5	0.9 %
1000	646.6	44.6 %	361.8	1.0 %
50	0.117	87.6 %	0.0263	45.6 %
100	0.709	87.3 %	0.152	40.6 %
250	9.98	84.5 %	2.08	25.4 %
500	178.6	80.9 %	37.8	9.4 %
750	924.1	80.9 %	184.4	4.5 %
1000	2076.6	82.7 %	391.8	8.4 %

We get **another 2x speedup** ($N > 500$) by replacing LAPACK routines DHGEEZ and DHGEEQZ by **Dackland-Kågström's blocked Hessenberg-triangular algorithms** (**ACM TOMS'99**) for transforming (A, E) to generalized

CONCLUSIONS – SO FAR

- State-of-the-art HPC systems have **deep memory hierarchie**
- **Recursion** efficiently provides automatic variable blocking for hierarchy.
- Recursive blocking \implies 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-typ