


Performance Modeling for Ranking Blocked Algorithms

Elmar Peise

Aachen Institute for Advanced Study in
Computational Engineering Science

27.4.2012



Inversion of a triangular matrix $L \leftarrow L^{-1} \in \mathbb{R}^{n \times n}$ 

Variant 1

$$\begin{aligned} L_{10} &\leftarrow L_{10} L_{00} \\ L_{10} &\leftarrow -L_{11}^{-1} L_{10} \\ L_{11} &\leftarrow L_{11}^{-1} \end{aligned}$$

Variant 2

$$\begin{aligned} L_{21} &\leftarrow L_{22}^{-1} L_{21} \\ L_{21} &\leftarrow -L_{21} L_{11}^{-1} \\ L_{11} &\leftarrow L_{11}^{-1} \end{aligned}$$


Variant 3

$$\begin{aligned} L_{21} &\leftarrow -L_{21} L_{11}^{-1} \\ L_{20} &\leftarrow L_{20} + L_{21} L_{10} \\ L_{10} &\leftarrow L_{11}^{-1} L_{10} \\ L_{11} &\leftarrow L_{11}^{-1} \end{aligned}$$

Variant 4

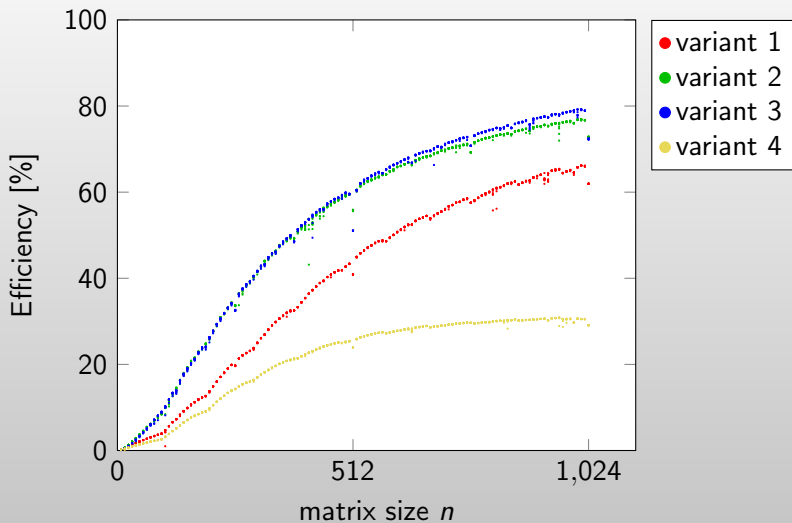
$$\begin{aligned} L_{21} &\leftarrow -L_{22}^{-1} L_{21} \\ L_{20} &\leftarrow L_{20} - L_{21} L_{10} \\ L_{10} &\leftarrow L_{10} L_{00} \\ L_{11} &\leftarrow L_{11}^{-1} \end{aligned}$$

Inversion of a triangular matrix $L \leftarrow L^{-1} \in \mathbb{R}^{n \times n}$



Intel Harpertown E5450 @ 2.99 GHz — 1 core — Intel MKL BLAS

Inversion of a triangular matrix $L \leftarrow L^{-1} \in \mathbb{R}^{n \times n}$



Intel Harpertown E5450 @ 2.99 GHz — 1 core — Intel MKL BLAS

Inversion of a triangular matrix $L \leftarrow L^{-1} \in \mathbb{R}^{n \times n}$

3rd (●)

2nd (●)

1st (●)

4th (●)

Variant 1

$$\begin{aligned} L_{10} &\leftarrow L_{10}L_{00} \\ L_{10} &\leftarrow -L_{11}^{-1}L_{10} \\ L_{11} &\leftarrow L_{11}^{-1} \end{aligned}$$

Variant 2

$$\begin{aligned} L_{21} &\leftarrow L_{22}^{-1}L_{21} \\ L_{21} &\leftarrow -L_{21}L_{11}^{-1} \\ L_{11} &\leftarrow L_{11}^{-1} \end{aligned}$$

Variant 3

$$\begin{aligned} L_{21} &\leftarrow -L_{21}L_{11}^{-1} \\ L_{20} &\leftarrow L_{20} + L_{21}L_{10} \\ L_{10} &\leftarrow L_{11}^{-1}L_{10} \\ L_{11} &\leftarrow L_{11}^{-1} \end{aligned}$$

Variant 4

$$\begin{aligned} L_{21} &\leftarrow -L_{22}^{-1}L_{21} \\ L_{20} &\leftarrow L_{20} - L_{21}L_{10} \\ L_{10} &\leftarrow L_{10}L_{00} \\ L_{11} &\leftarrow L_{11}^{-1} \end{aligned}$$

Can we rank the algorithms based on their update statements?

No!

- ① Sampling
- ② Modeling
- ③ Prediction and Ranking

$$B \leftarrow 0.37 B L^{-1} \quad \text{with} \quad B \in \mathbb{R}^{128 \times 96}, L \in \mathbb{R}^{96 \times 96}$$

Routine call

```
dtrsm(R, L, N, U, 128, 96, 0.37, L, 128, B, 128)
```

↓ Sampling ↓

Performance counters (PAPI)

<i>ticks</i>	<i>flops</i>	<i>L1misses</i>	<i>...</i>
887491	595968	4	

Routine call

```
dtrsm(R, L, N, U, 128, 96, 0.37, L, 128, B, 128)
```

Performance is independent of L and B .
⇒ Assign (random) memory chunk of sufficient size.

Input

```
(dtrsm, R, L, N, U, 128, 96, 0.37,  $128 \times 96$ , 128,  $128 \times 96$ , 128)
```


Input


```
(dtrsm, R, L, N, U, 128, 96, 0.37, 12288, 128, 12288, 128)
(dtrsm, R, L, N, U, 128, 96, 0.37, 12288, 128, 12288, 128)
(dtrsm, R, L, N, U, 128, 96, 0.37, 12288, 128, 12288, 128)
(dtrsm, R, L, N, U, 128, 96, 0.37, 12288, 128, 12288, 128)
⋮
```

↓ Sampling ↓

Performance counters

<i>ticks</i>	<i>flops</i>	<i>L1misses</i>	<i>...</i>
12755926	595976	5172	
926324	595968	27	
887491	595968	4	
882572	595968	1	
			⋮

dtrsm(R, L, N, U, 128, 96, 0.37, L, 128, B, 128)




AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core


- ① Sampling
- ② Modeling
- ③ Prediction and Ranking

- Understanding performance
- Model structure
- Model generation
- Modeling results

$$\boxed{B} \leftarrow \alpha \begin{array}{|c|} \hline \triangle \\ \hline \end{array} A^{-1} \boxed{B}$$




`dtrsm(side, uplo, transA, diag, 256, 256, 0.5, A, 256, B, 256)`



AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core

Small scale


`dgemm(N, N, m, n, k, 0.5, A, m, B, k, 0.5, C, m)`



AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core


Large scale

```
dgemm(N, N, m, n, k, 0.5, A, m, B, k, 0.5, C, m)
```



AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core

`dgemm(N, N, 256, 256, 256, alpha, A, 256, B, 256, beta, C, 256)`




AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core

```
dgemm(N, N, 256, 256, 256, 0.5, A, 256, B, 256, 0.5, C, 256)
```

No performance dependency

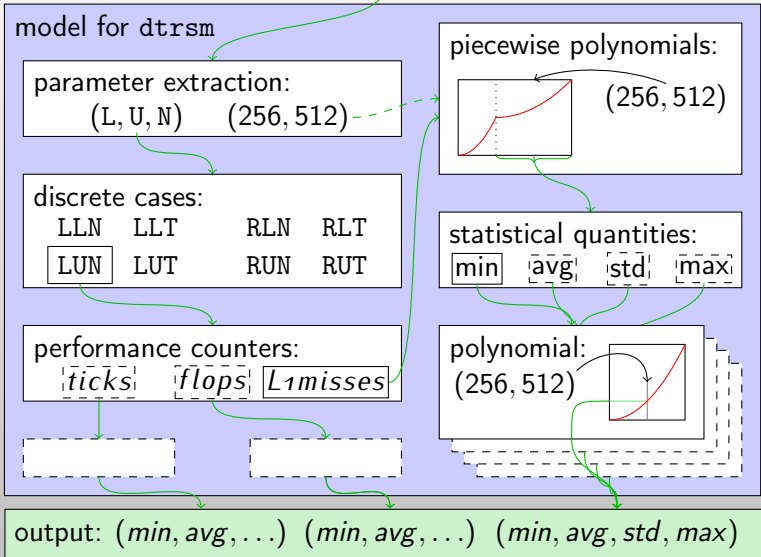
`dgemm(N, N, 128, 128, 128, 0.5, A, 1da, B, 1dB, 0.5, C, 1dC)`



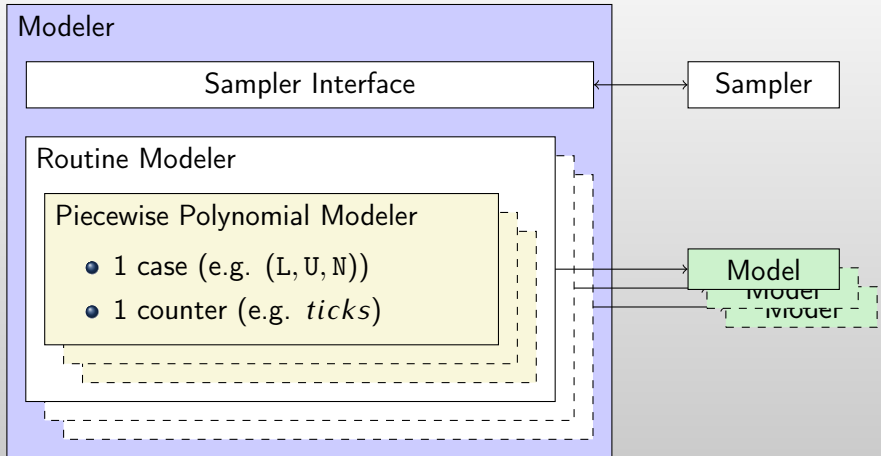
AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core

- Understanding performance
- **Model structure**
- Model generation
- Modeling results

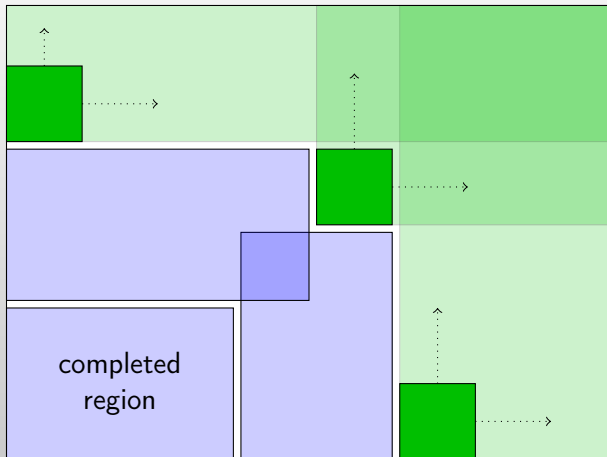
input: (dtrsm, L, U, N, N, 256, 512, 0.7, A, 512, B, 512)

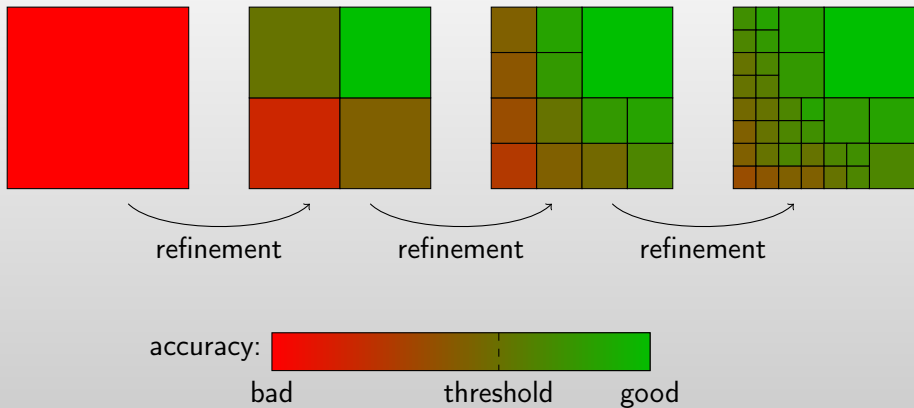


- Understanding performance
- Model structure
- **Model generation**
- Modeling results




- Model Expansion
- Adaptive Refinement






- Understanding performance
- Model structure
- Model generation
- Modeling results

```
dtrsm(L, L, N, N, m, n, .5, L, 2500, B, 2500)
```



AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core — GotoBLAS2

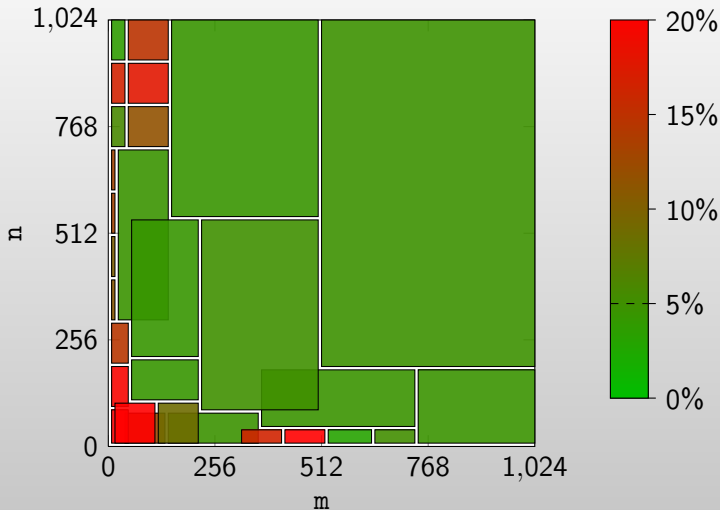
```
dtrsm(L, L, N, N, m, n, .5, L, 2500, B, 2500)
```



AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core — GotoBLAS2

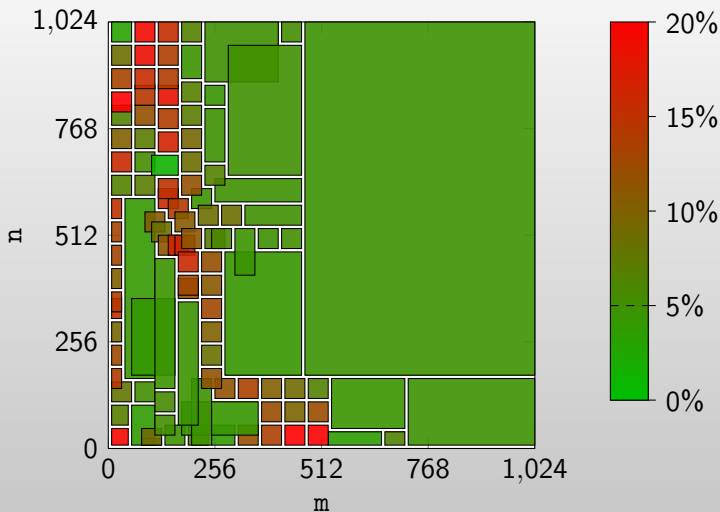
Smaller approximation error bound

```
dtrsm(L, L, N, N, m, n, .5, L, 2500, B, 2500)
```




AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core — GotoBLAS2

```
dtrsm(L, L, N, N, m, n, .5, L, 2500, B, 2500)
```




AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core — GotoBLAS2

```
dtrsm(L, L, N, N, m, n, .5, L, 2500, B, 2500)
```



AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core — GotoBLAS2



```
dtrsm(L, L, N, N, m, n, .5, L, 2500, B, 2500)
```



AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core — GotoBLAS2

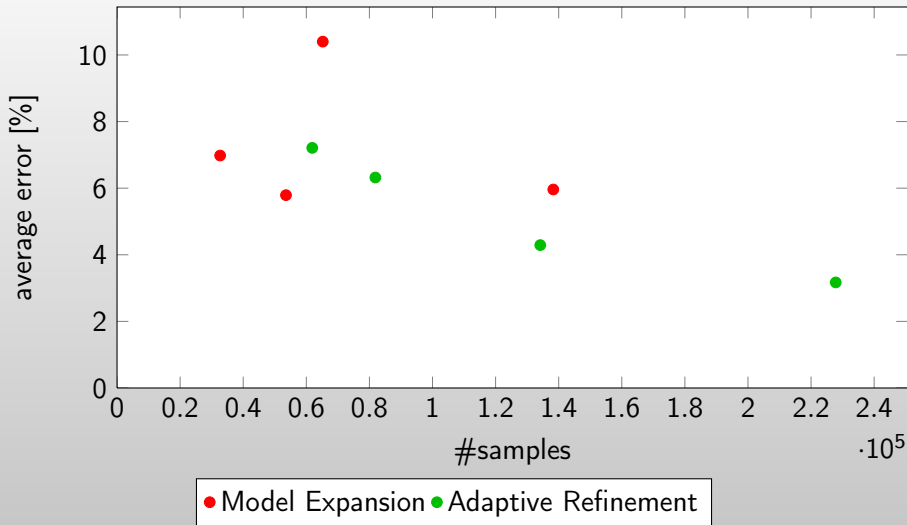
Smaller approximation error bound

```
dtrsm(L, L, N, N, m, n, .5, L, 2500, B, 2500)
```



AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core — GotoBLAS2


```
dtrsm(L, L, N, N, m, n, .5, L, 2500, B, 2500)
```



AMD Opteron 8356 (Barcelona) @ 2.3 GHz — 1 core — GotoBLAS2

- ① Sampling
- ② Modeling
- ③ Prediction and Ranking

Blocked algorithm revisited: $L \leftarrow L^{-1}$



Variant 1

$$\begin{aligned} L_{10} &\leftarrow L_{10}L_{00} \\ L_{10} &\leftarrow -L_{11}^{-1}L_{10} \\ L_{11} &\leftarrow L_{11}^{-1} \end{aligned}$$

Variant 2

$$\begin{aligned} L_{21} &\leftarrow L_{22}^{-1}L_{21} \\ L_{21} &\leftarrow -L_{21}L_{11}^{-1} \\ L_{11} &\leftarrow L_{11}^{-1} \end{aligned}$$

Variant 3

$$\begin{aligned} L_{21} &\leftarrow -L_{21}L_{11}^{-1} \\ L_{20} &\leftarrow L_{20} + L_{21}L_{10} \\ L_{10} &\leftarrow L_{11}^{-1}L_{10} \\ L_{11} &\leftarrow L_{11}^{-1} \end{aligned}$$

Variant 4

$$\begin{aligned} L_{21} &\leftarrow -L_{22}^{-1}L_{21} \\ L_{20} &\leftarrow L_{20} - L_{21}L_{10} \\ L_{10} &\leftarrow L_{10}L_{00} \\ L_{11} &\leftarrow L_{11}^{-1} \end{aligned}$$

```

int trinv1_(char* diag, int* n, double* A, int* ldA, int* bsize) {
    if (*n == 1) {
        if (diag[0] == 'N')
            *A = 1 / *A;
        return 0;
    }

    int ione = 1; double one = 1; double mone = -1;

    for (int p = 0; p < *n; p += *bsize) {
        int b = *bsize;
        if (p + b > *n)
            b = *n - p;
#define A00 (A)
#define A10 (A + p)
#define A11 (A + *ldA * p + p)

        dtrmm_("R", "L", "N", diag, &b, &p, &one, A00, ldA, A10, ldA);
        dtrsm_("L", "L", "N", diag, &b, &p, &mone, A11, ldA, A10, ldA);
        trinv1_(diag, &b, A11, ldA, &ione);
    }
    return 0;
}

```


Variant 1

$$\begin{aligned}
 L_{10} &\leftarrow L_{10}L_{00} \\
 L_{10} &\leftarrow -L_{11}^{-1}L_{10} \\
 L_{11} &\leftarrow L_{11}^{-1}
 \end{aligned}$$


- Input: `trinv1(N, 300, A, 300, 100)`
- Compute routine invocations

update	routine invocation
$L_{10} \leftarrow L_{10}L_{00}$	<code>(dtrmm, R, L, N, N, 100, 0, 1, ., 300, ., 300)</code>
$L_{10} \leftarrow -L_{11}^{-1}L_{10}$	<code>(dtrsm, L, L, N, N, 100, 0, -1, ., 300, ., 300)</code>
$L_{11} \leftarrow L_{11}^{-1}$	<code>(trinv1, N, 100, ., 300, 1)</code>
$L_{10} \leftarrow L_{10}L_{00}$	<code>(dtrmm, R, L, N, N, 100, 100, 1, ., 300, ., 300)</code>
$L_{10} \leftarrow -L_{11}^{-1}L_{10}$	<code>(dtrsm, L, L, N, N, 100, 100, -1, ., 300, ., 300)</code>
$L_{11} \leftarrow L_{11}^{-1}$	<code>(trinv1, N, 100, ., 300, 1)</code>
$L_{10} \leftarrow L_{10}L_{00}$	<code>(dtrmm, R, L, N, N, 100, 200, 1, ., 300, ., 300)</code>
$L_{10} \leftarrow -L_{11}^{-1}L_{10}$	<code>(dtrsm, L, L, N, N, 100, 200, -1, ., 300, ., 300)</code>
$L_{11} \leftarrow L_{11}^{-1}$	<code>(trinv1, N, 100, ., 300, 1)</code>


- Evaluate performance Models
- Accumulate *ticks*




Intel Harpertown E5450 @ 2.99 GHz — 1 core — Intel MKL BLAS



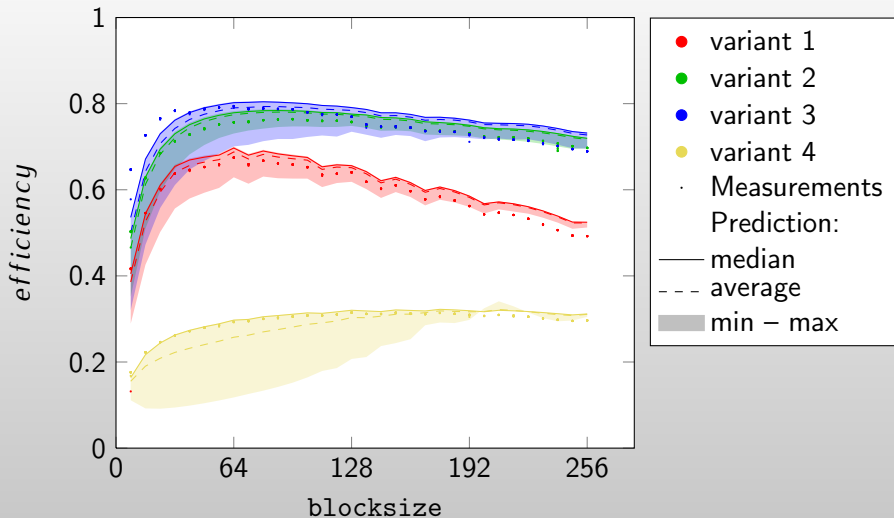
Intel Harpertown E5450 @ 2.99 GHz — 1 core — Intel MKL BLAS



Intel Harpertown E5450 @ 2.99 GHz — 1 core — Intel MKL BLAS




Intel Harpertown E5450 @ 2.99 GHz — 1 core — Intel MKL BLAS



Intel Harpertown E5450 @ 2.99 GHz — 1 core — Intel MKL BLAS

- ① Sampling
 - Performance measurement tool
 - Applicable to various DLA routines
- ② Modeling
 - Automatic modeling system
 - Flexible and accurate
- ③ Prediction and Ranking
 - Accurate predictions
 - Correct ranking



Funding from DFG is gratefully acknowledged

Deutsche
Forschungsgemeinschaft

DFG