


High Performance Implementation of LTL^T Decomposition of a Skew-Symmetric Matrix



Chao Yin, Ishna Satyarth, Ruqing Xu, Devin Matthews

Sep. 26th, 2024

Austin, TX, USA




Blocked algorithms

Blocked Left Looking Algorithm (BLL)

$$\left(\begin{array}{c|c} x_{21} & X_{22} \\ \chi_{31} & x_{32}^T \\ \hline x_{41} & X_{42} \end{array} \right) := \left(\begin{array}{c|c} x_{21} & X_{22} \\ \chi_{31} & x_{32}^T \\ \hline x_{41} & X_{42} \end{array} \right) - \left(\begin{array}{c|c} L_{20} & l_{21} \\ \hline l_{30}^T & \lambda_{31} \\ \hline L_{40} & l_{41} \end{array} \right) \left(\begin{array}{c|c} X_{00} & \star \\ \hline x_{10}^T & 0 \end{array} \right) \left(\begin{array}{c|c} l_{10} & L_{20}^T \\ \hline 1 & l_{21}^T \end{array} \right)$$

Skew-Triangular-Matrix-Multiplication (SK-TRI-GEMM)

$$\left[\left(\begin{array}{c|c} \chi_{11} & \star \\ x_{21} & X_{22} \\ \chi_{31} & x_{32}^T \\ \hline x_{41} & X_{42} \end{array} \right), \left(\begin{array}{c|c} L_{22} & 0 \\ \hline l_{32}^T & 1 \\ \hline L_{42} & l_{43} \end{array} \right) \right] := \text{LTLT_UNB} \left(\begin{array}{c|c} \chi_{11} & \star \\ x_{21} & X_{22} \\ \chi_{31} & x_{32}^T \\ \hline x_{41} & X_{42} \end{array} \right)$$



Blocked Right Looking Algorithm (BRL)

$$\left[\left(\begin{array}{c|c} \chi_{11} & \star \\ x_{21} & X_{22} \\ \chi_{31} & x_{32}^T \\ \hline x_{41} & X_{42} \end{array} \right), \left(\begin{array}{c|c} L_{22} & 0 \\ \hline l_{32}^T & 1 \\ \hline L_{42} & l_{43} \end{array} \right) \right] := \text{LTLT_UNB_0} \left(\begin{array}{c|c} \chi_{11} & \star \\ x_{21} & X_{22} \\ \chi_{31} & x_{32}^T \\ \hline x_{41} & X_{42} \end{array} \right)$$

$$\left(\begin{array}{c|c} \chi_{33} & \star \\ x_{43} & X_{44} \end{array} \right) := \left(\begin{array}{c|c} \chi_{33} & \star \\ x_{43} & X_{44} \end{array} \right) - \left(\begin{array}{c|c} l_{32}^T & 1 \\ \hline L_{42} & l_{43} \end{array} \right) \left(\begin{array}{c|c} X_{22} & \star \\ \hline x_{32}^T & 0 \end{array} \right) \left(\begin{array}{c|c} l_{32} & L_{42}^T \\ \hline 1 & l_{43}^T \end{array} \right)$$

Skew-Triangular-Matrix-Multiplication (Triangular) (SK-TRI-GEMMT)

$$X_{44} := X_{44} + (l_{43}x_{43}^T - x_{43}l_{43}^T)$$

Skew-Symmetric Rank-2 (SKR2)

Background

Unblocked algorithms

Unblocked Left Looking Algorithm (UBLL)

$$\begin{pmatrix} \chi_{21} \\ x_{31} \end{pmatrix} := \begin{pmatrix} \chi_{21} \\ x_{31} \end{pmatrix} - \begin{pmatrix} l_{20}^T & \lambda_{21} \\ L_{30} & l_{31} \end{pmatrix} \begin{pmatrix} X_{00} & -x_{10} \\ x_{10}^T & 0 \end{pmatrix} \begin{pmatrix} l_{10} \\ 1 \end{pmatrix}$$

Skew-Triangular-Matrix-Vector-Multiplication (SKTRI-GEMV)

$$l_{32} := x_{31} / \chi_{21}$$

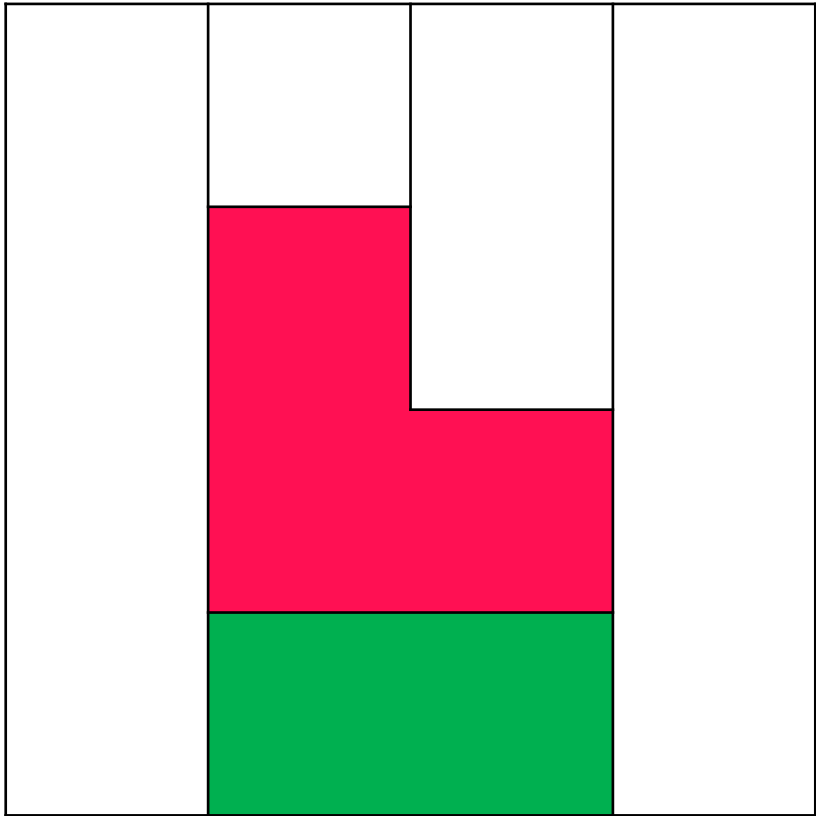
$$x_{31} := 0$$

Unblocked Right Looking Algorithm (UBRL)

$$l_{32} := x_{31} / \chi_{21}$$

$$x_{31} := 0$$


$$X_{33} := X_{33} + (l_{32}x_{32}^T - x_{32}l_{32}^T)$$



Skew-Symmetric Rank-2 (SKR2)
 General Matrix rank-2 update (GER2)


Implementation (in-place or out-of-place)

T Matrix in-place



T Matrix out-of-place

$$\left(\begin{array}{c|c} \chi_{33} & \star \\ \hline x_{43} & X_{44} \end{array} \right) := \left(\begin{array}{c|c} \chi_{33} & \star \\ \hline x_{43} & X_{44} \end{array} \right) - \left(\begin{array}{c|c} l_{32}^T & 1 \\ \hline L_{42} & l_{43} \end{array} \right) \left(\begin{array}{c|c} X_{22} & \star \\ \hline x_{32}^T & 0 \end{array} \right) \left(\begin{array}{c|c} l_{32} & L_{42}^T \\ \hline 1 & l_{43}^T \end{array} \right)$$



Implementation (Level-3 BLAS)

SKTRI-GEMM

C

+=

A

T

0	-0.5	0	0
0.5	0	-0.2	0
0	0.2	0	1.3
0	0	-1.3	0

B

SKTRI-GEMMT

C

+=

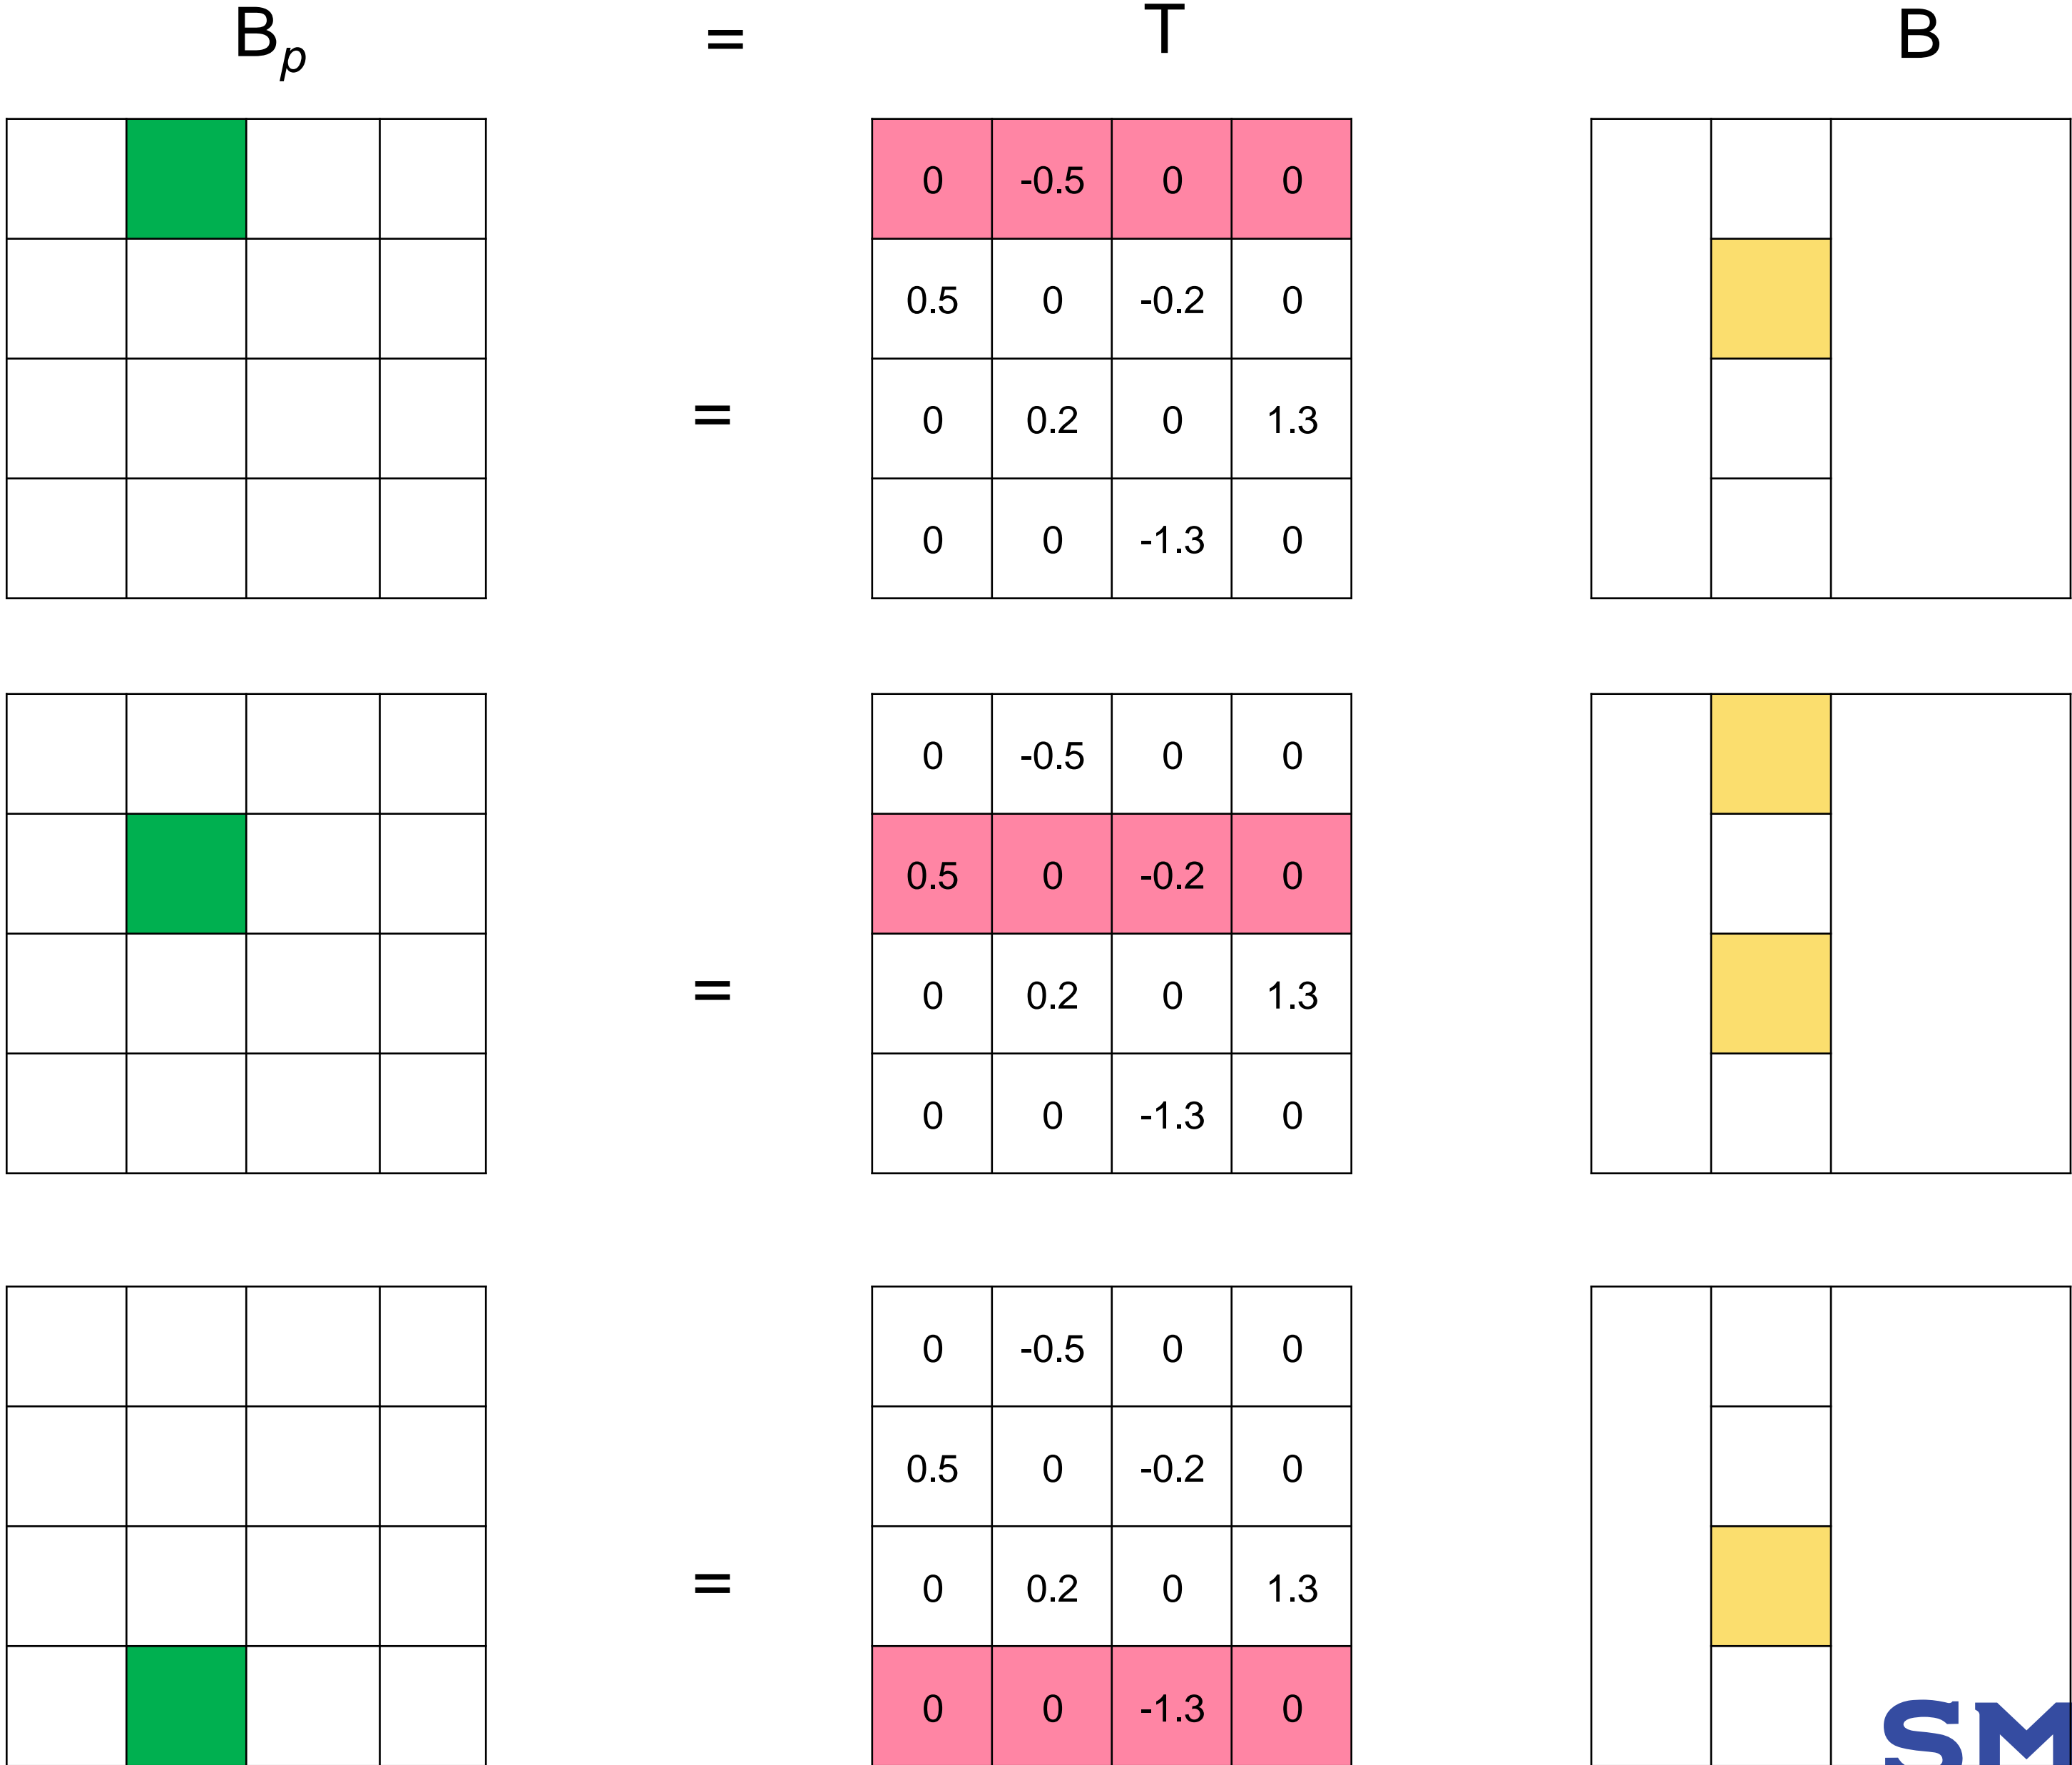

A

T

0	-0.5	0	0
0.5	0	-0.2	0
0	0.2	0	1.3
0	0	-1.3	0

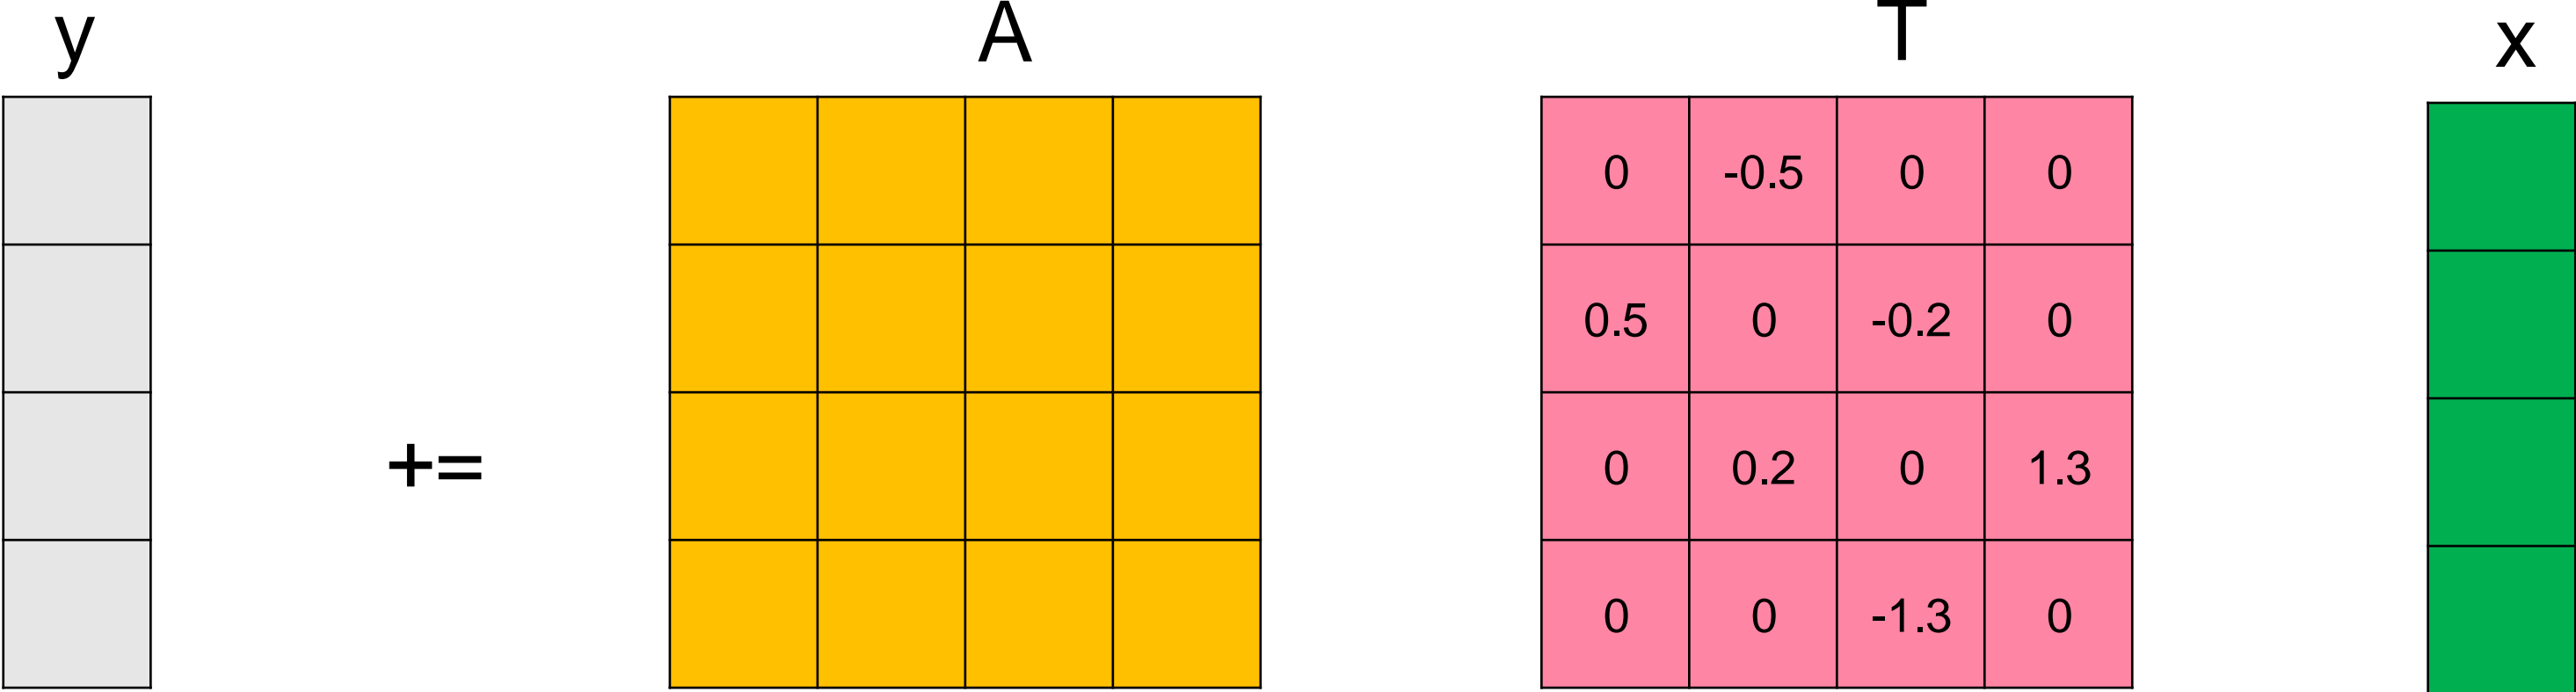
A^T

Implementation (Level-3 BLAS)




Implementation (Level-2 BLAS)


SKTRI-GEMV



SKR2

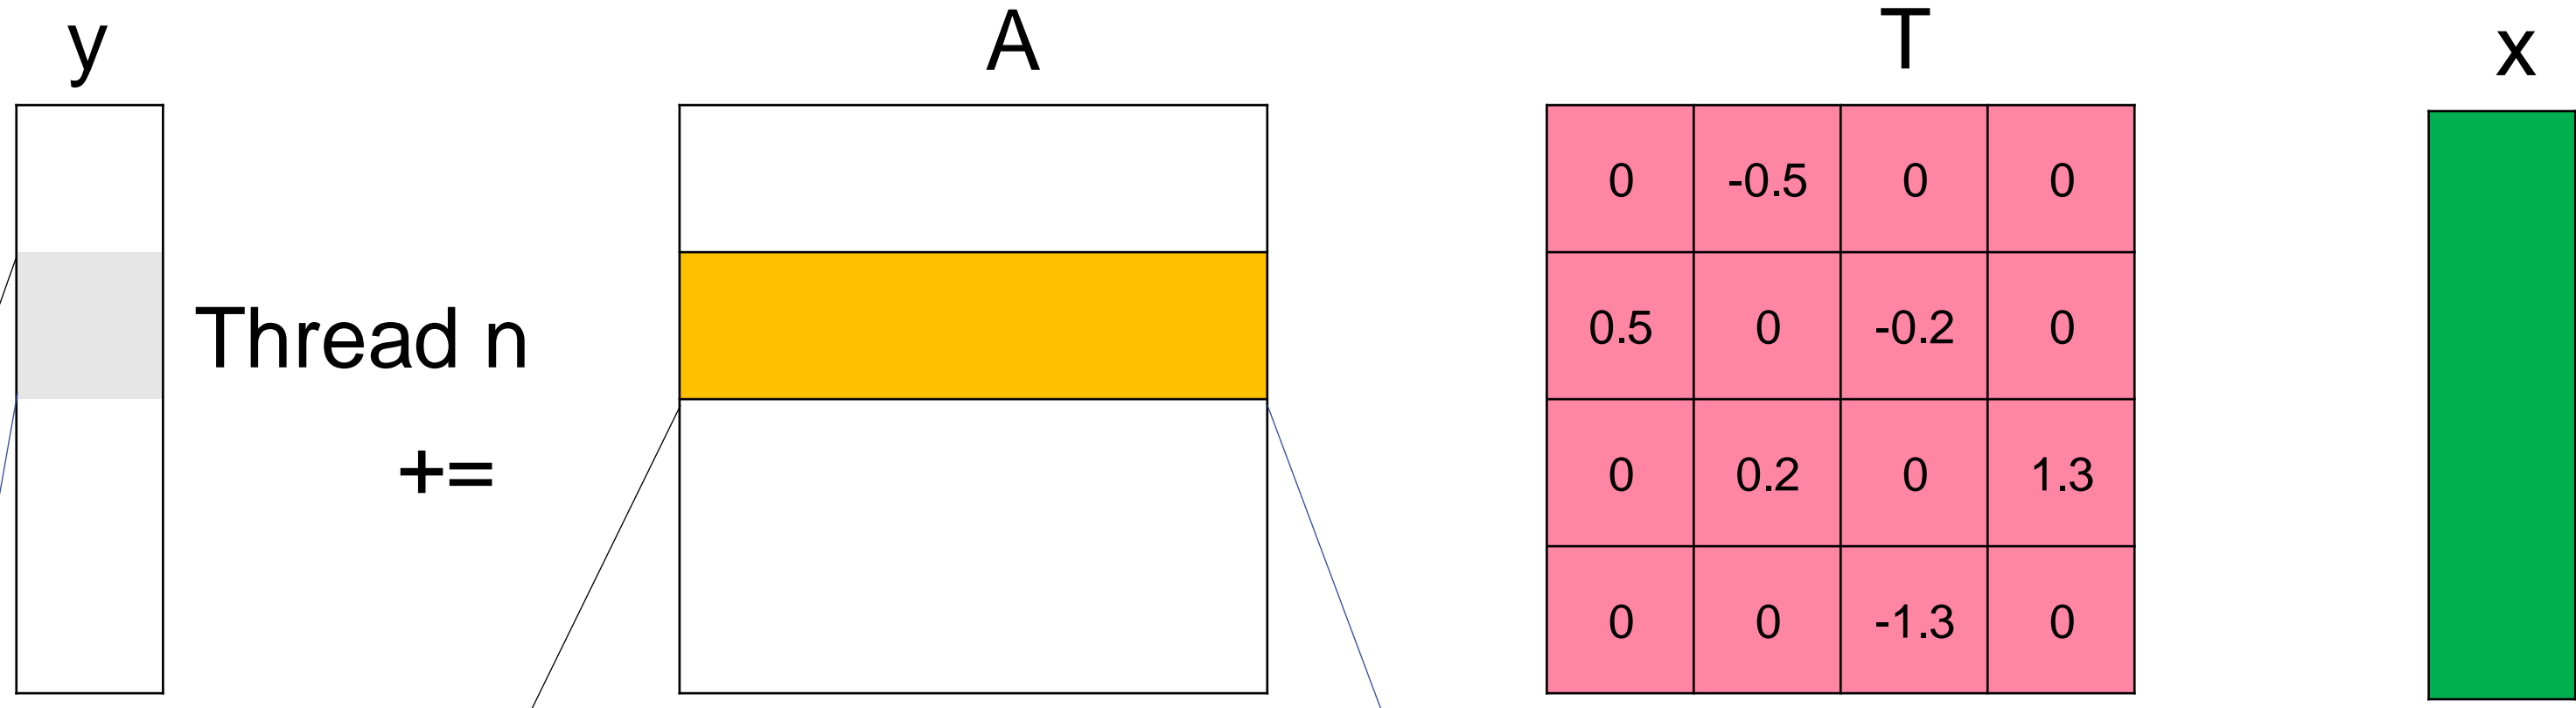


GER2




Implementation (SK-TRI-GEMV, Row-Major)

SKTRI-GEMV



0	-0.5	0	0
0.5	0	-0.2	0
0	0.2	0	1.3
0	0	-1.3	0

1. Parallelization




x'

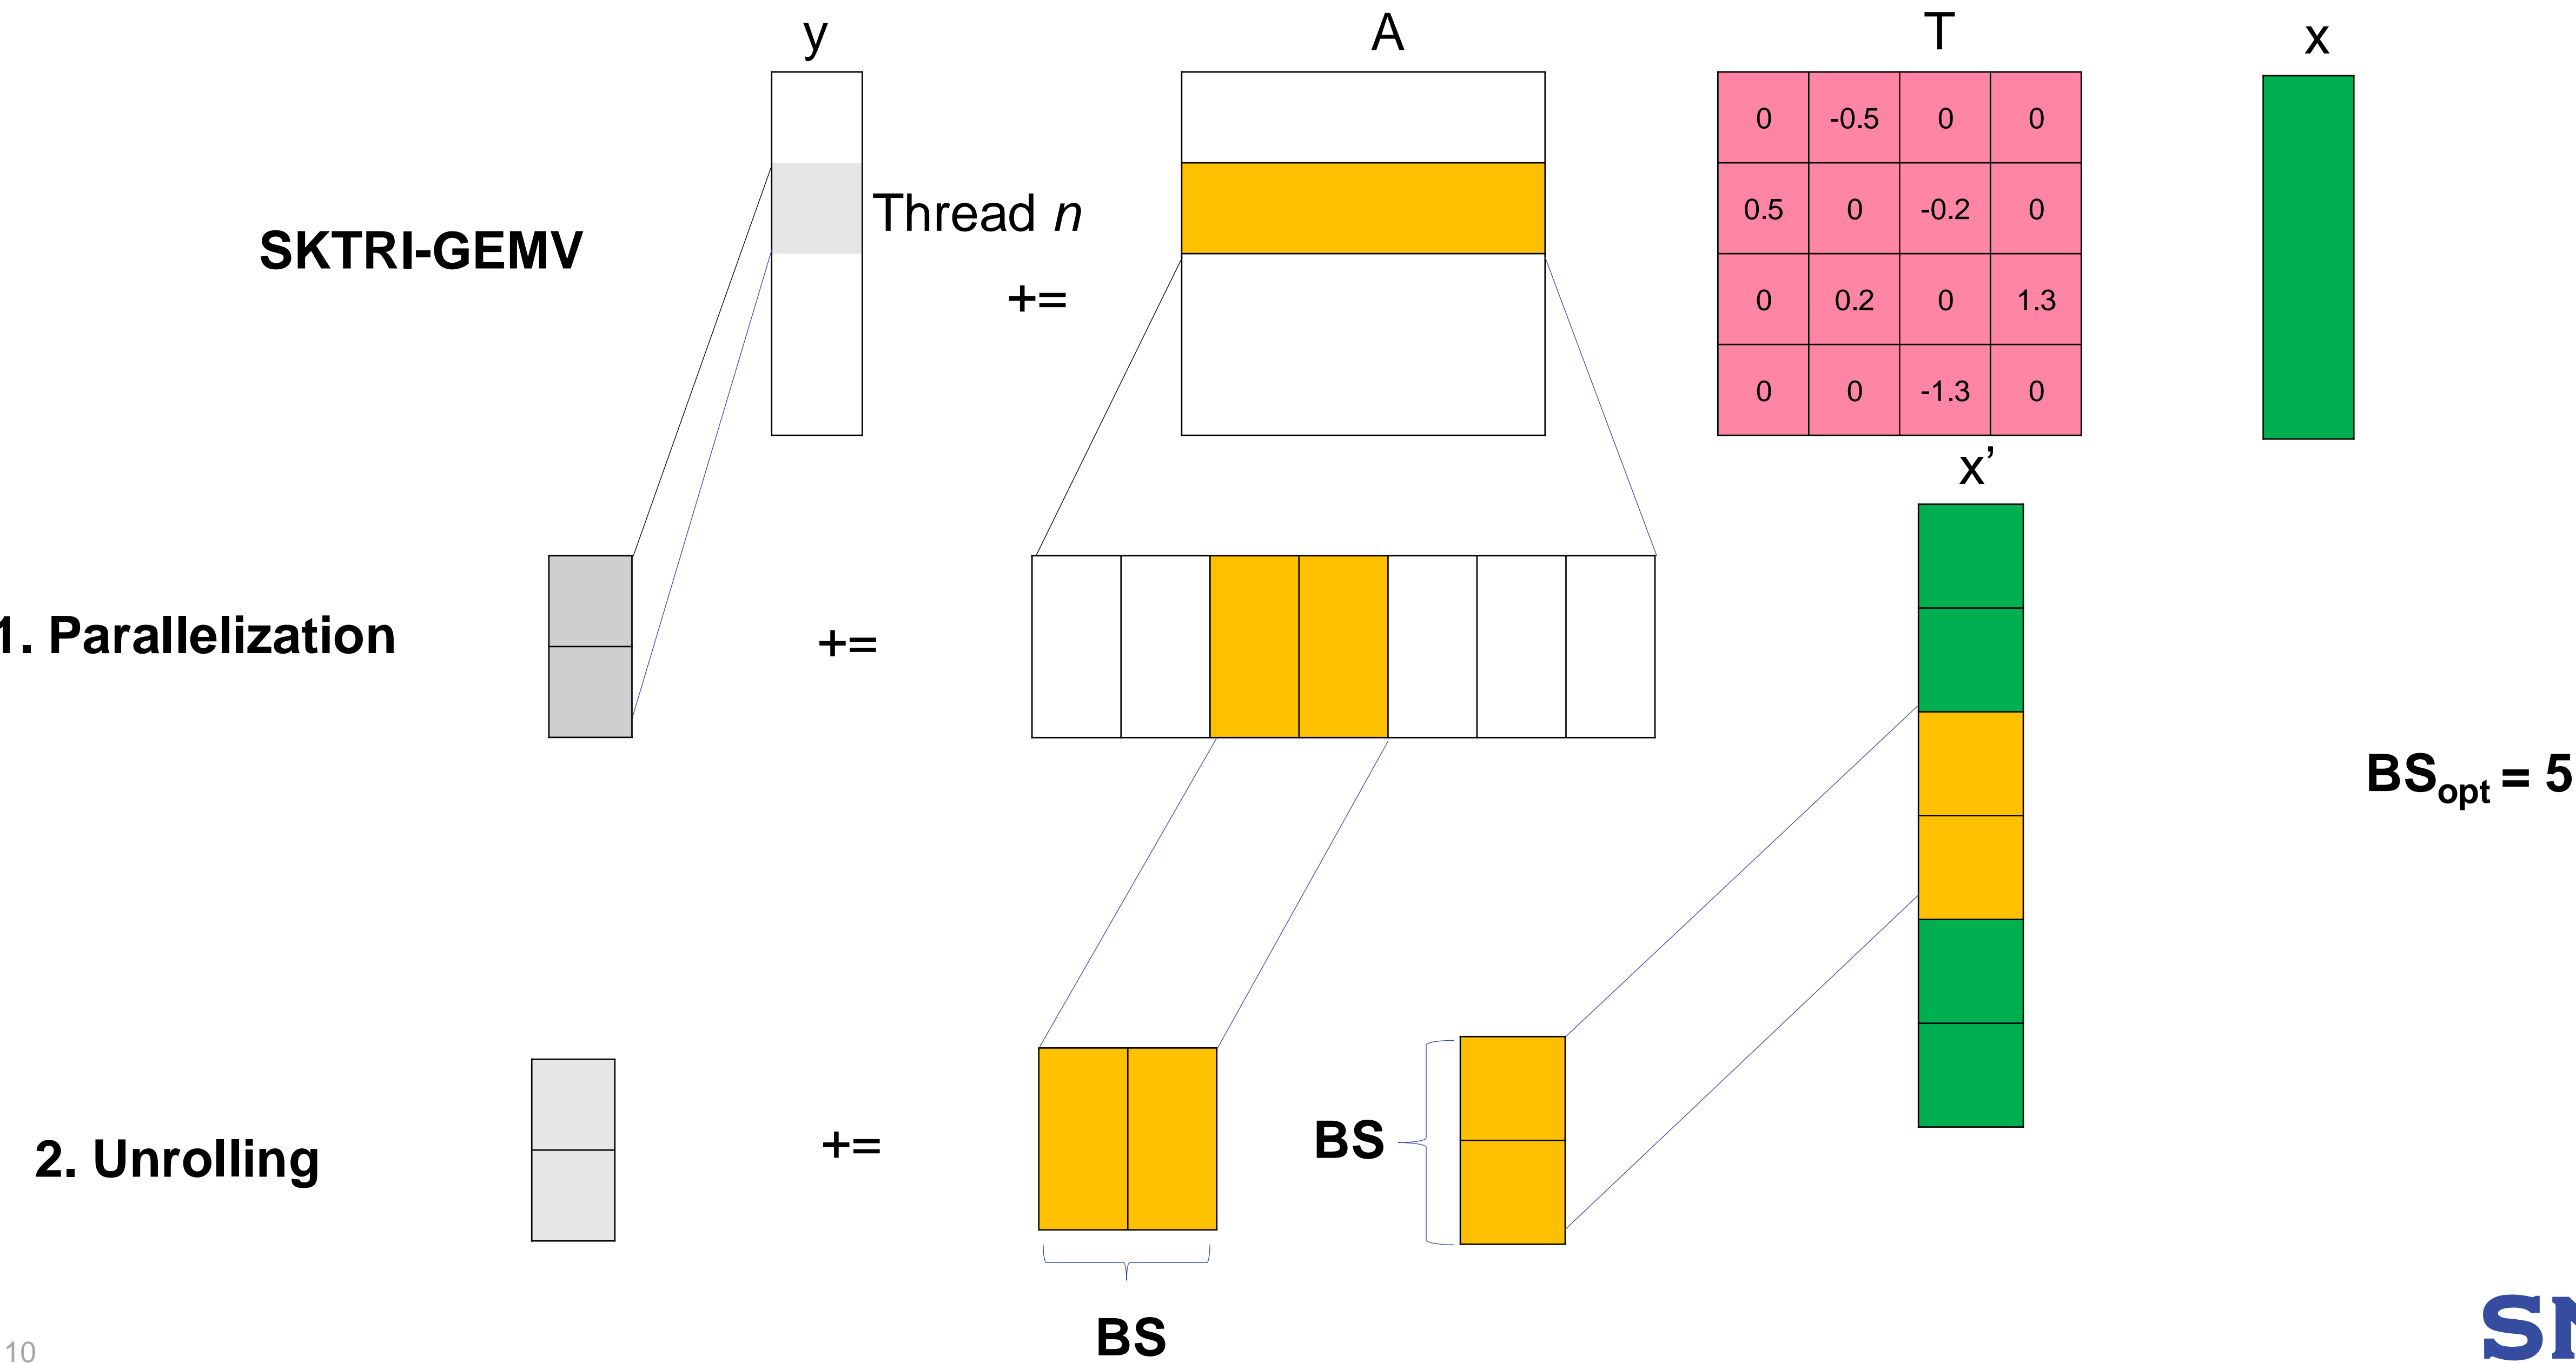


BS_{opt} = 5

2. Unrolling




Implementation (SK-TRI-GEMV, Column-Major)



Implementation (SKR2)



SKR2





1)

2)



Parallelization




Implementation (SKR2)





Implementation (GER2, Column-Major)



1. Parallelization



2. Unrolling




$BS_{opt} = 5$

Result (Different block size)

BLL + UBLL (SKTRI-GEMV, SKTRI-GEMM)

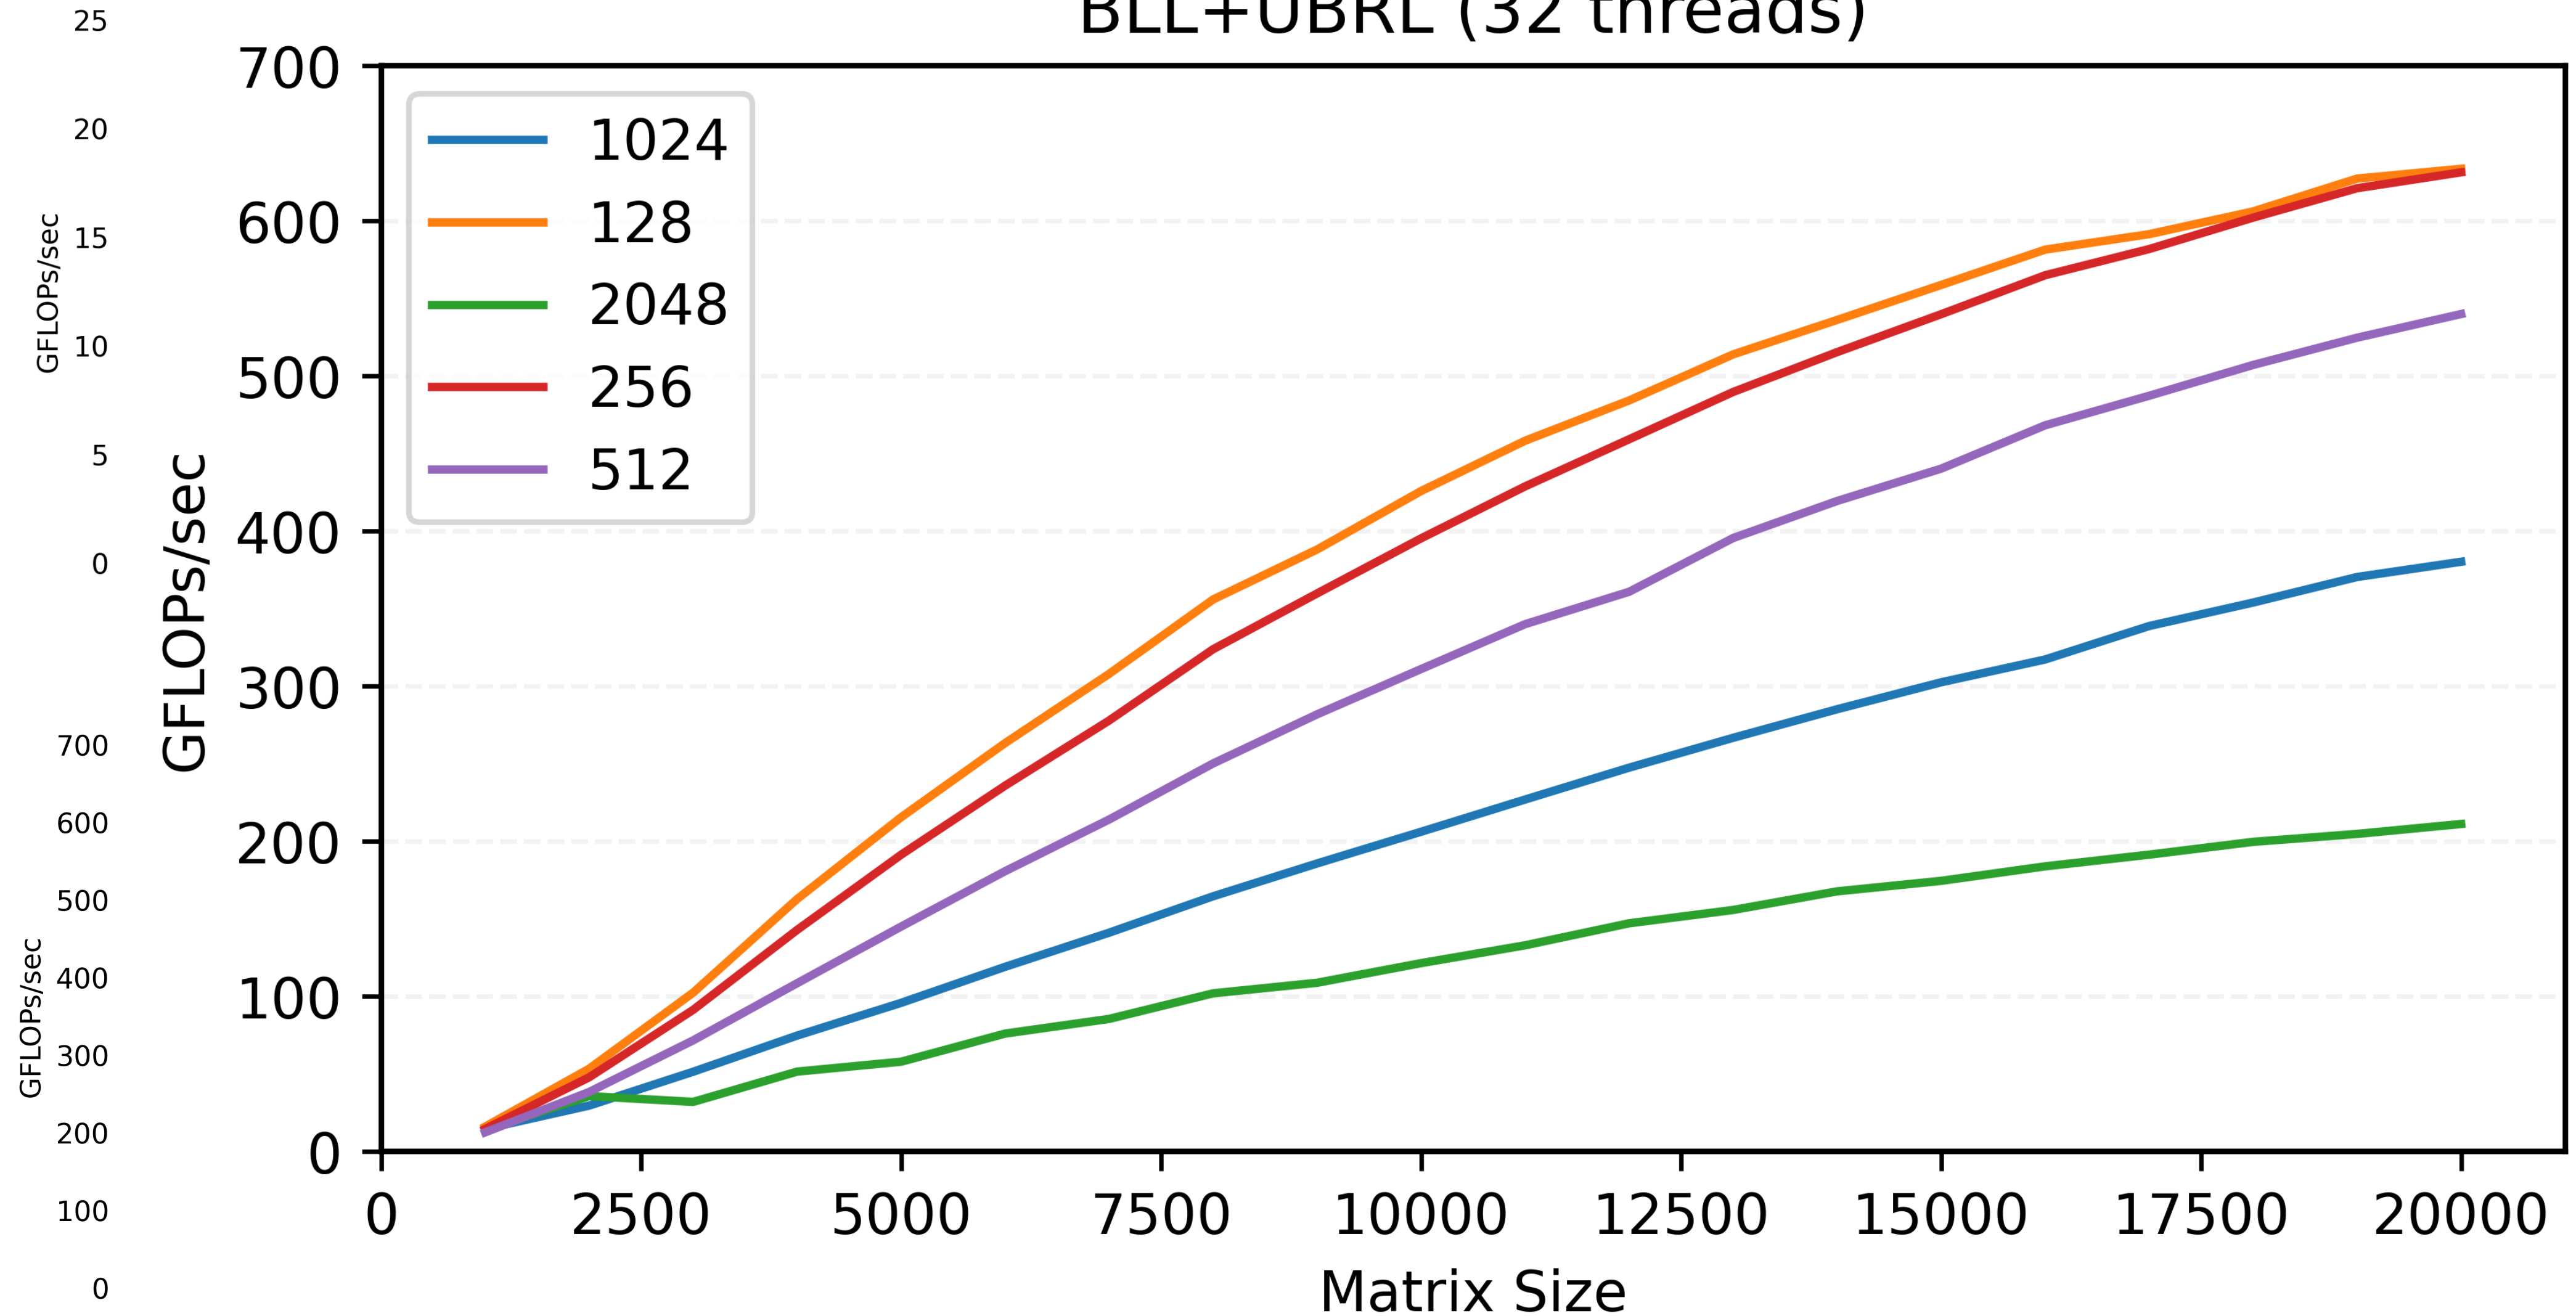
BLL+UBLL (32 threads)



Result (Different block size)

BLL + UBRL (SKR2, GER2, SKTRI-GEMM)

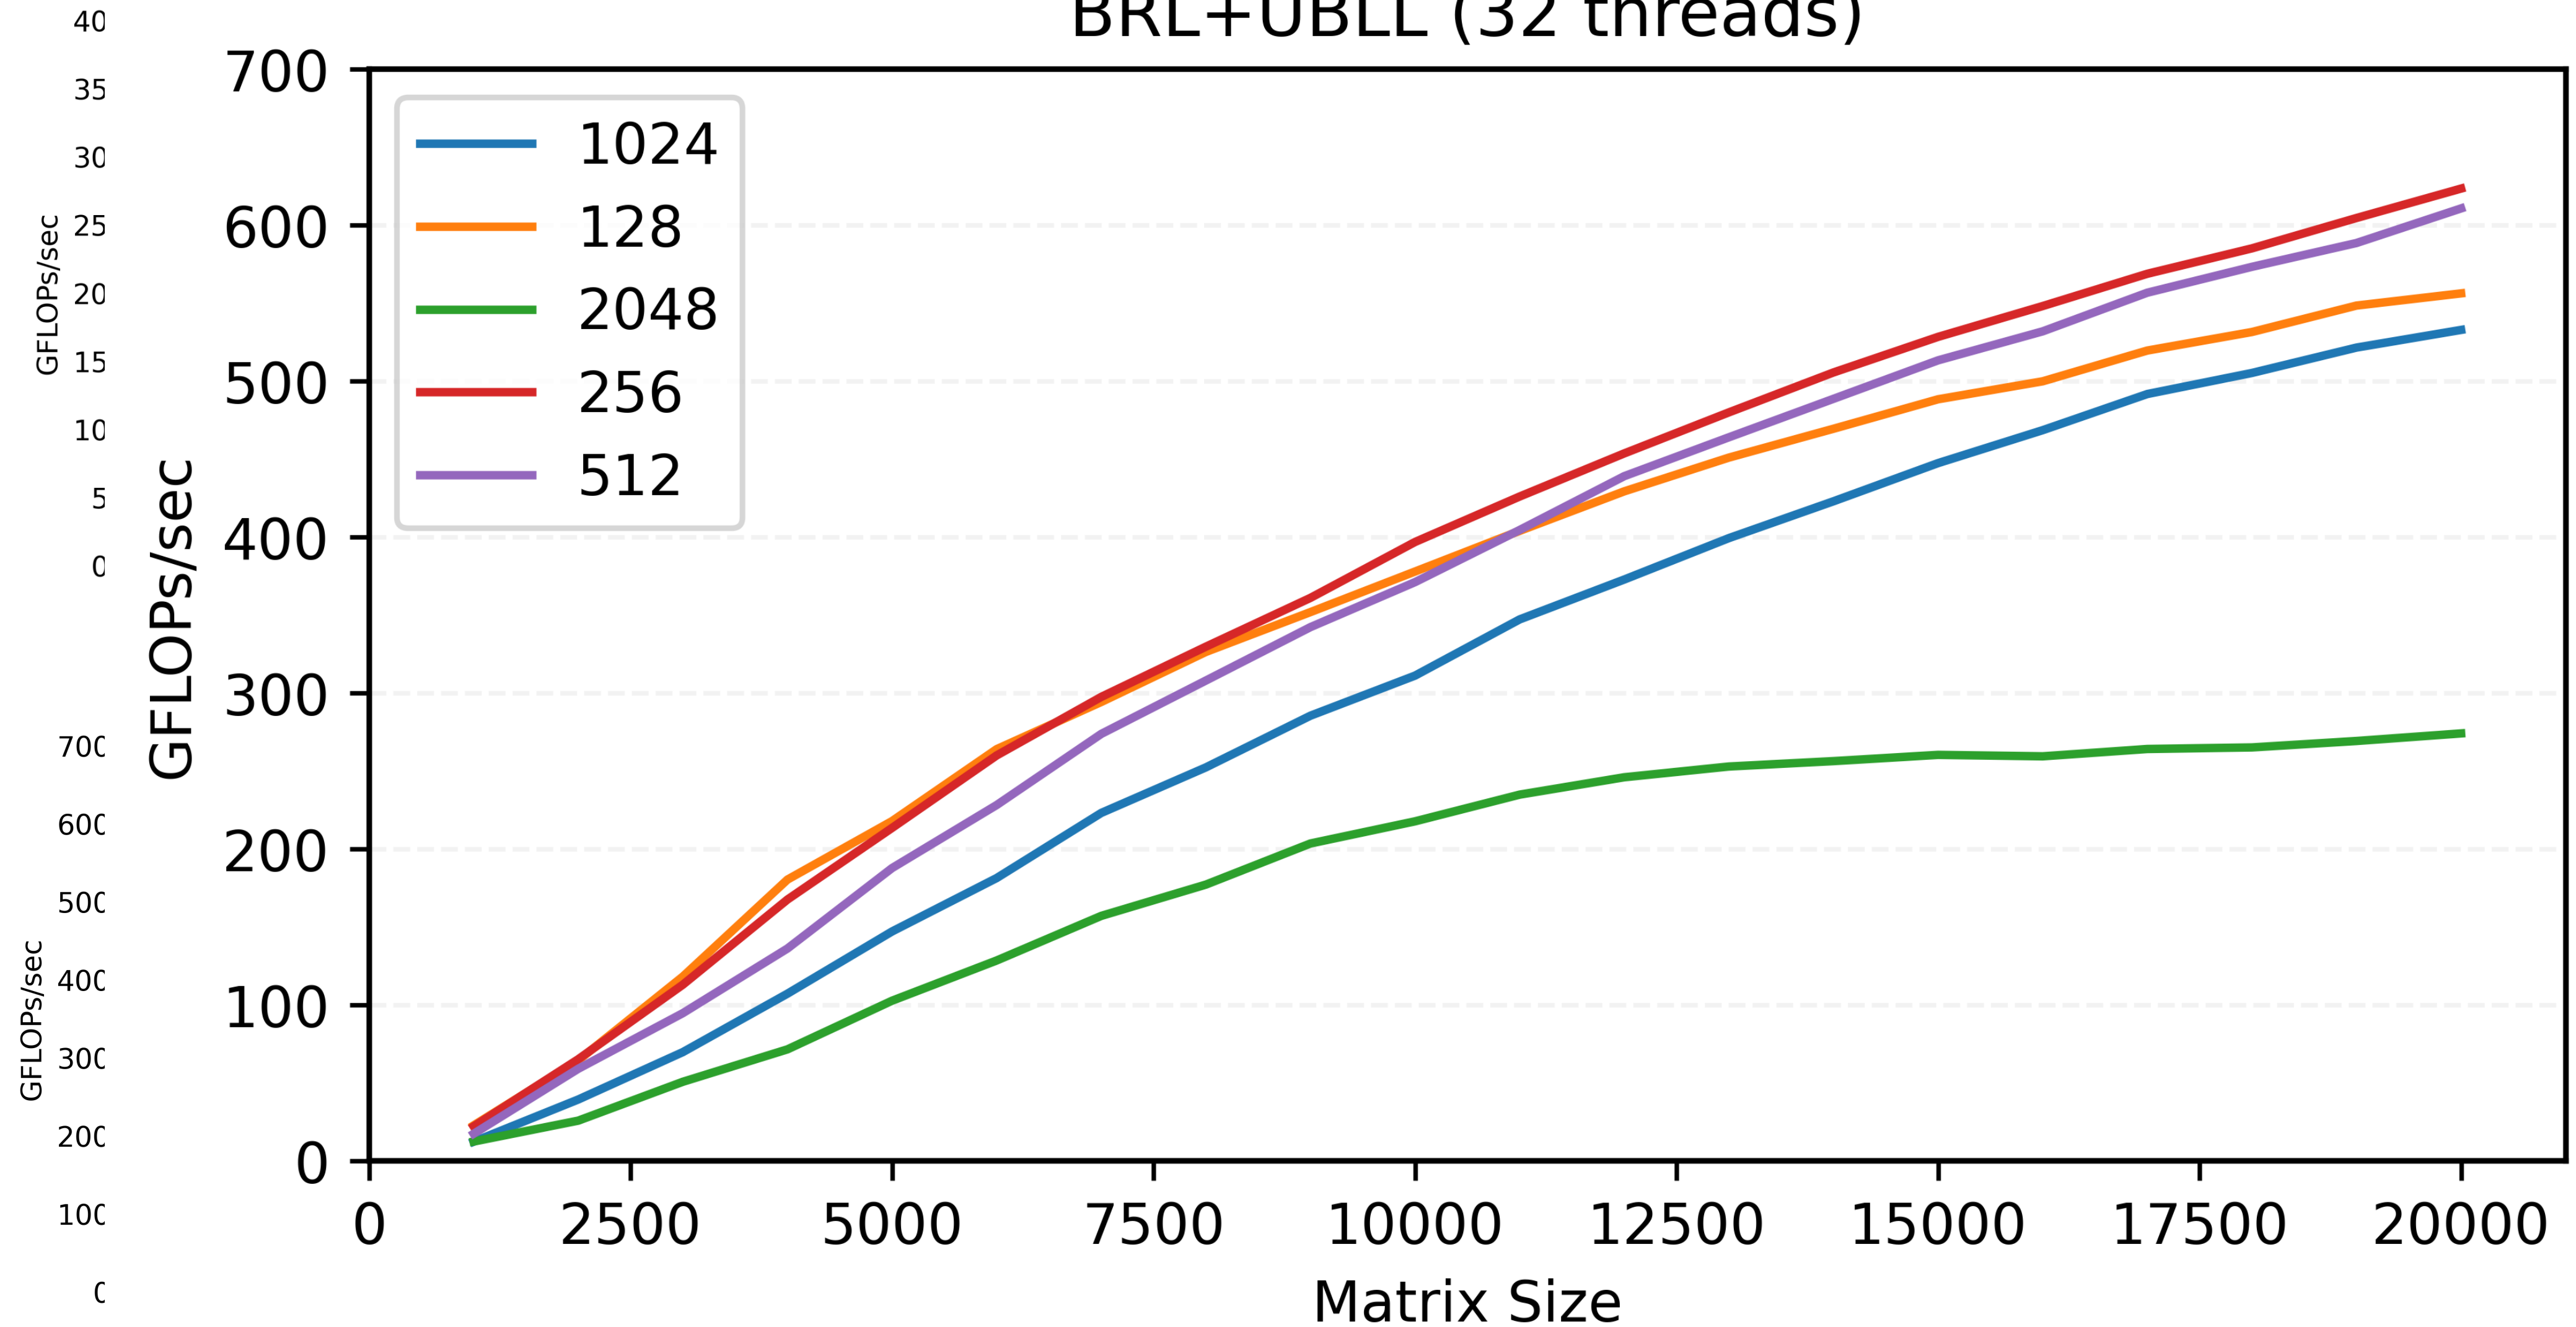
BLL+UBRL (32 threads)



Result (Different block size)

BRL + UBLL (SKTRI-GEMV, SKR2, SKTRI-GEMMT)


BRL+UBLL (32 threads)



Result (Different block size)

BRL + UBRL (GER2, SKR2, SKTRI-GEMMT)


BRL+UBRL (32 threads)



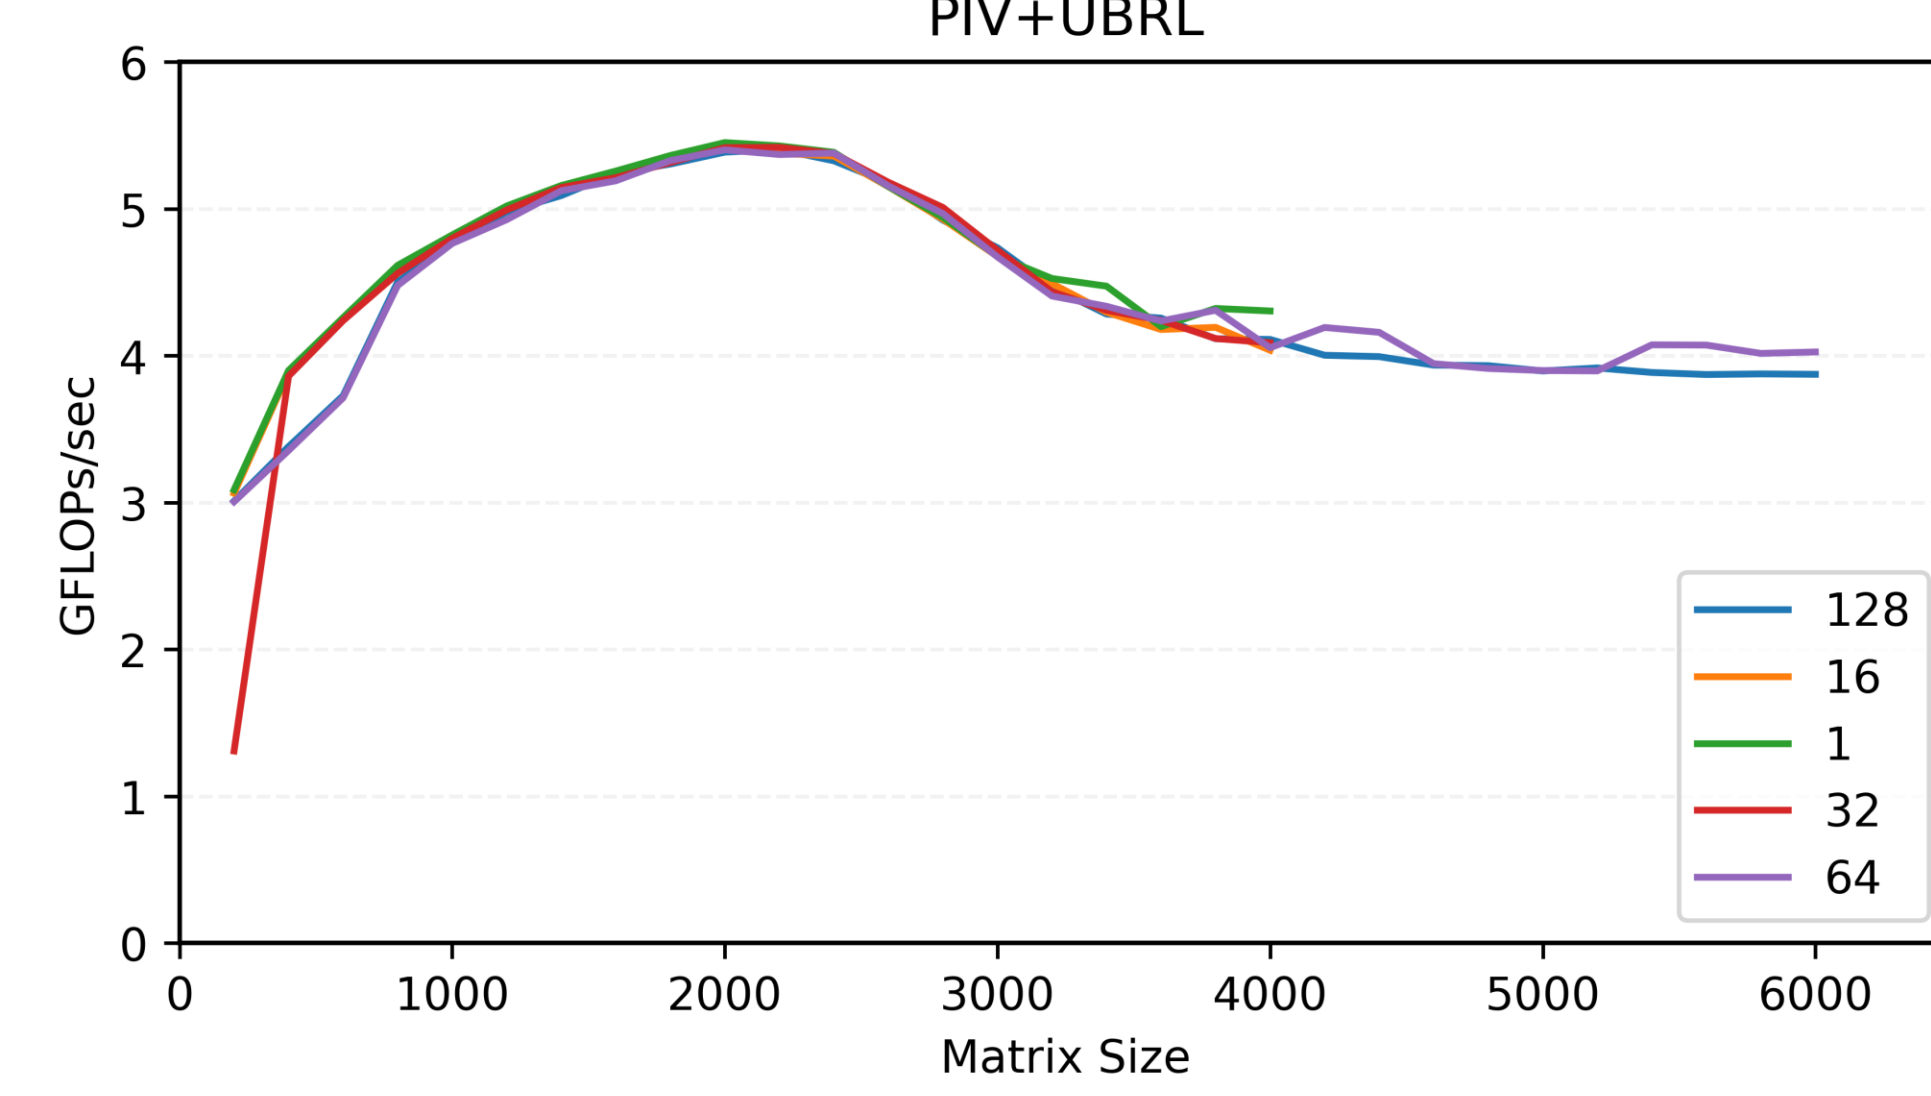
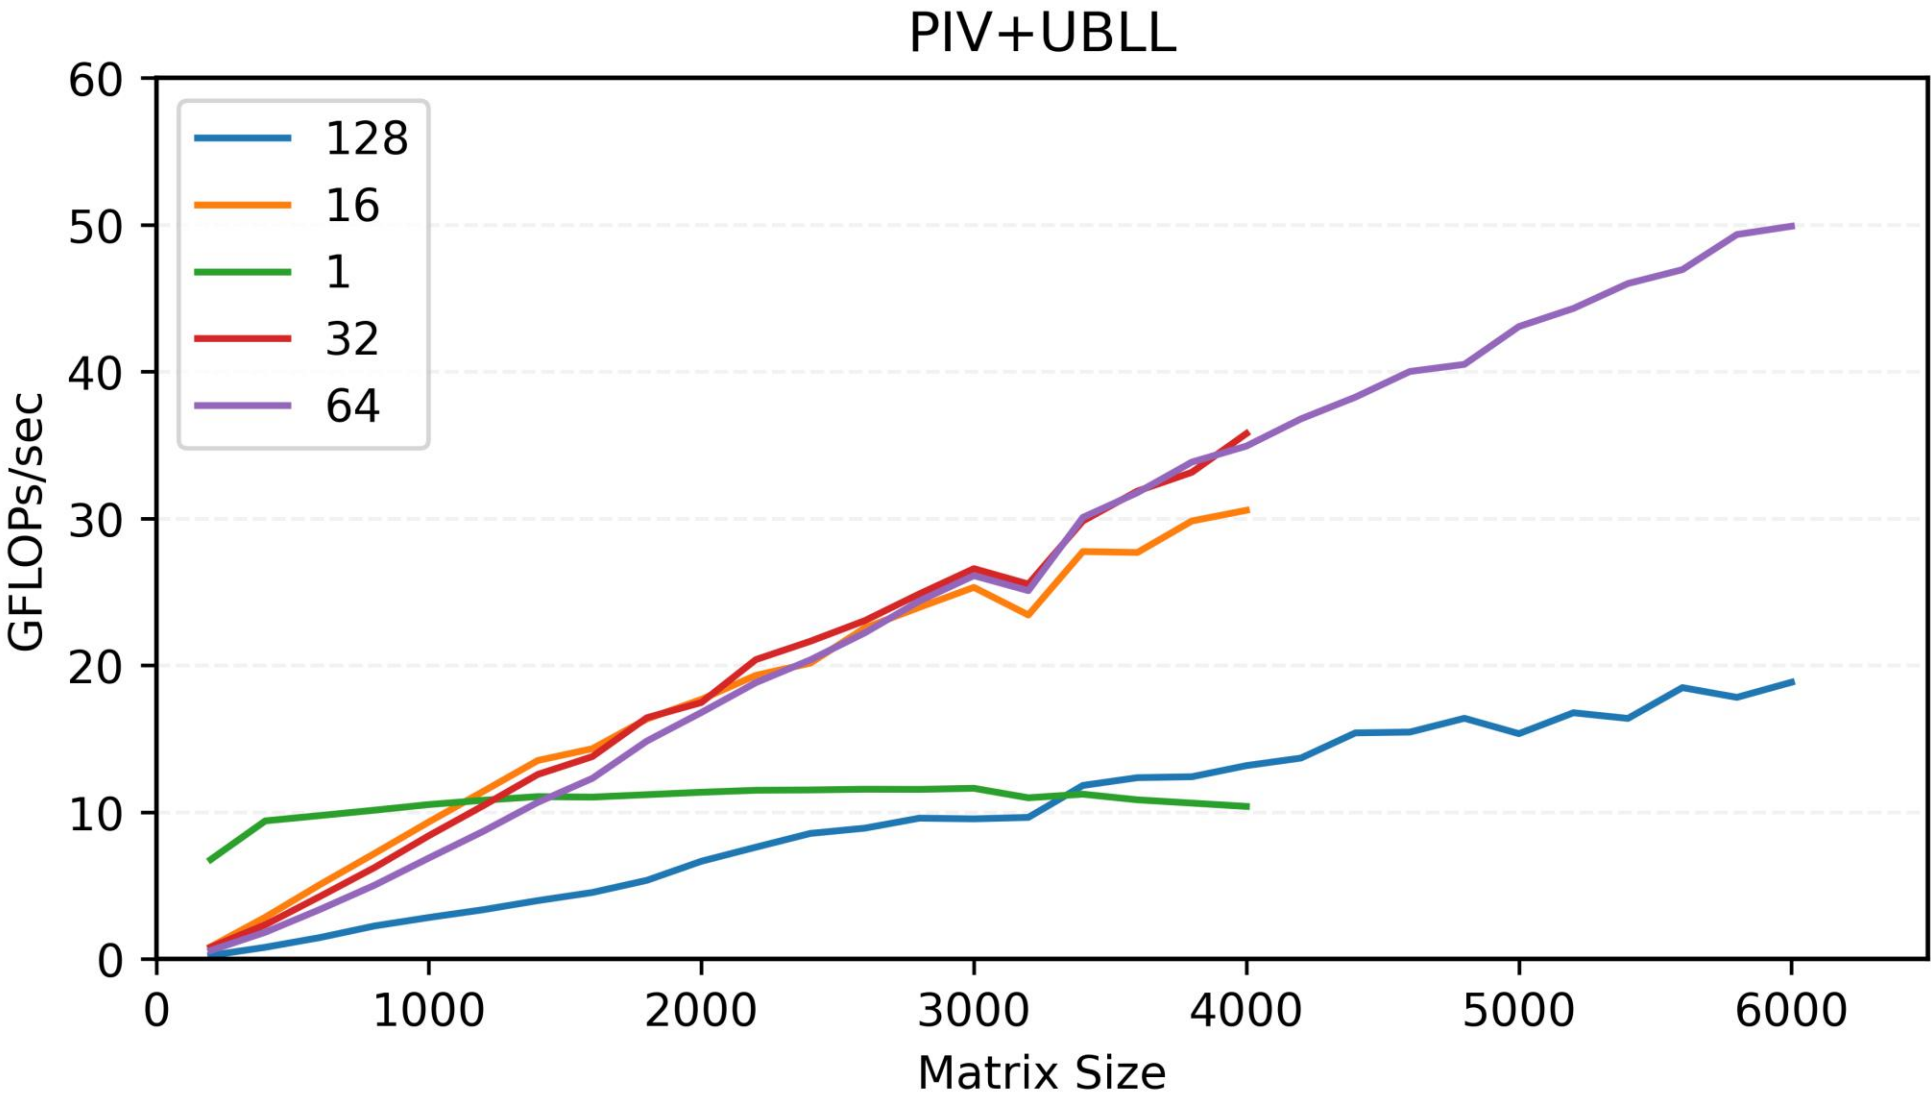
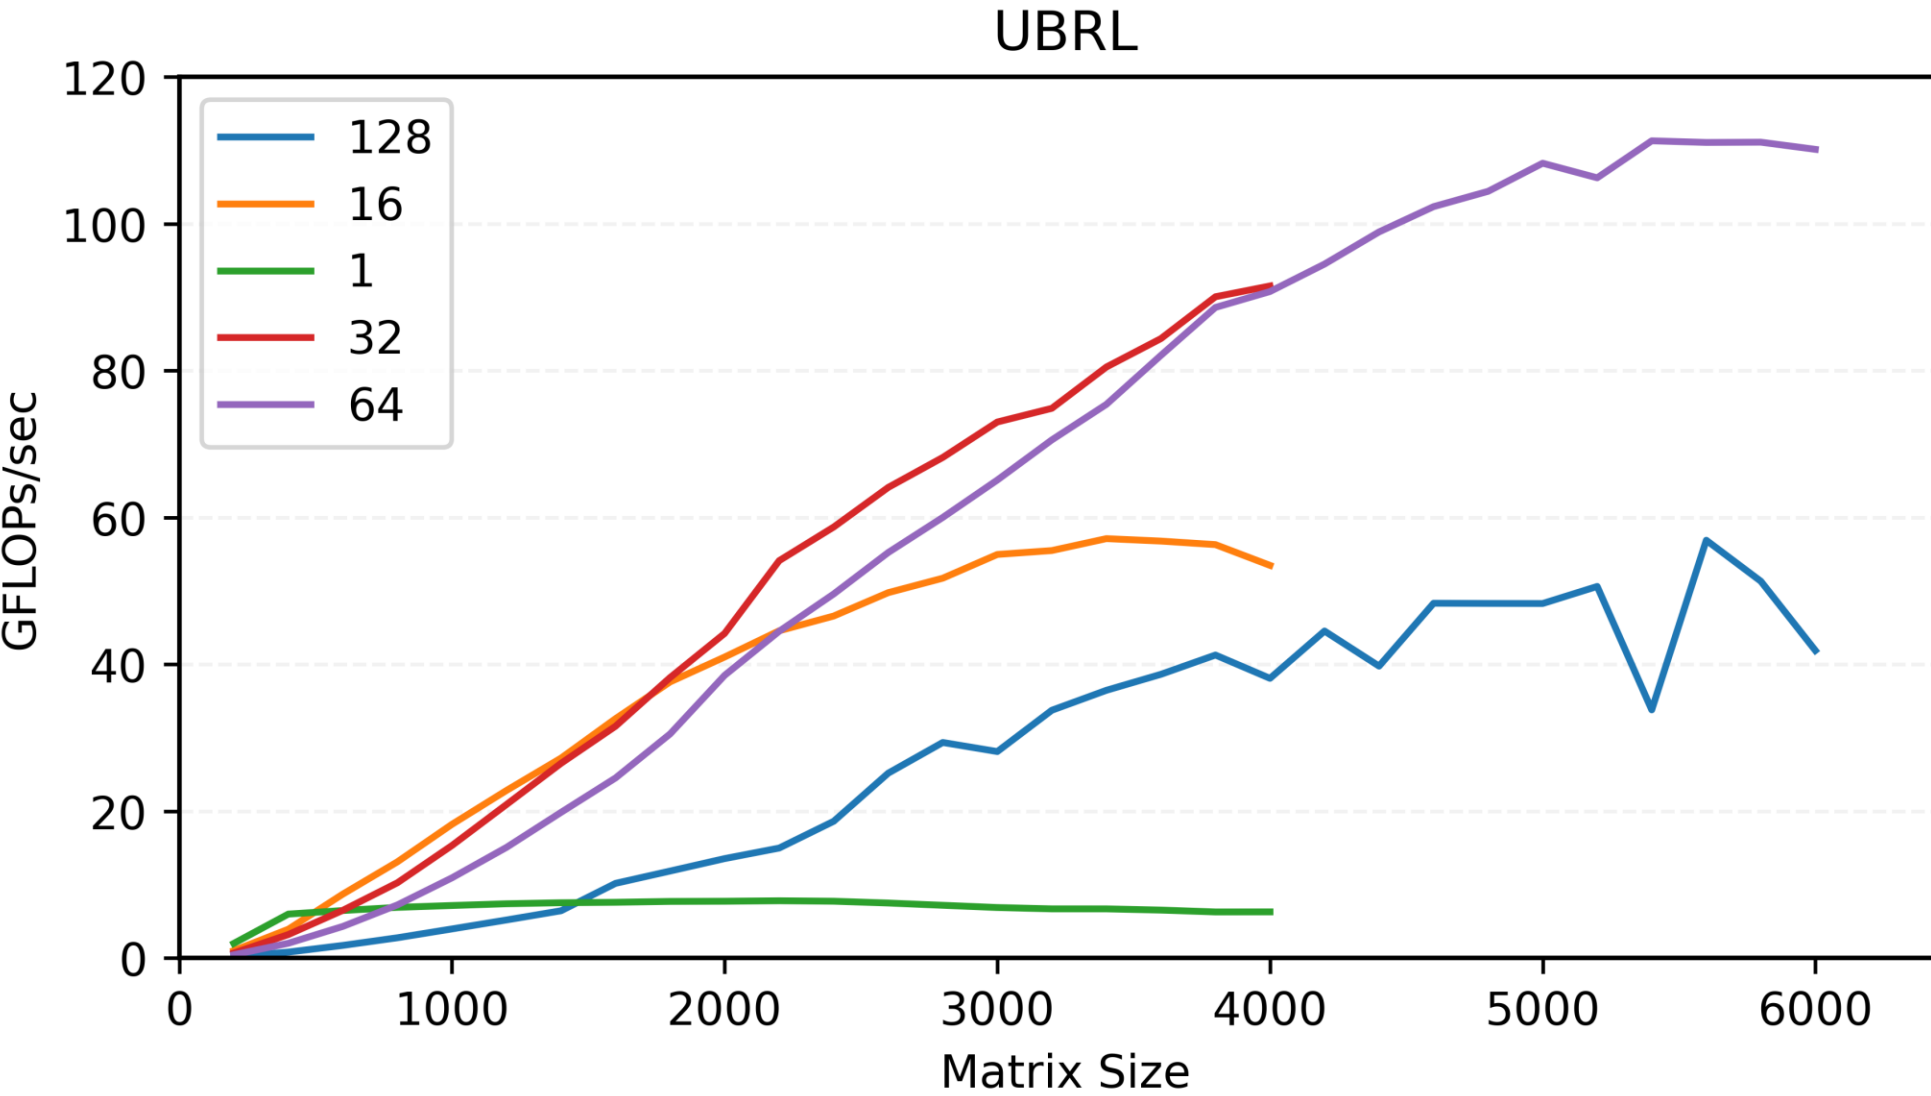

Result (Different block size)

PIV+ BRL + UBLL (SKTRI-GEMV, SKR2, SKTRI-GEMMT)

PIV+BRL+UBLL (32 threads)




Result (Unblocked Algorithms with multi-threading)



Comparison with other vendors (POTRF)

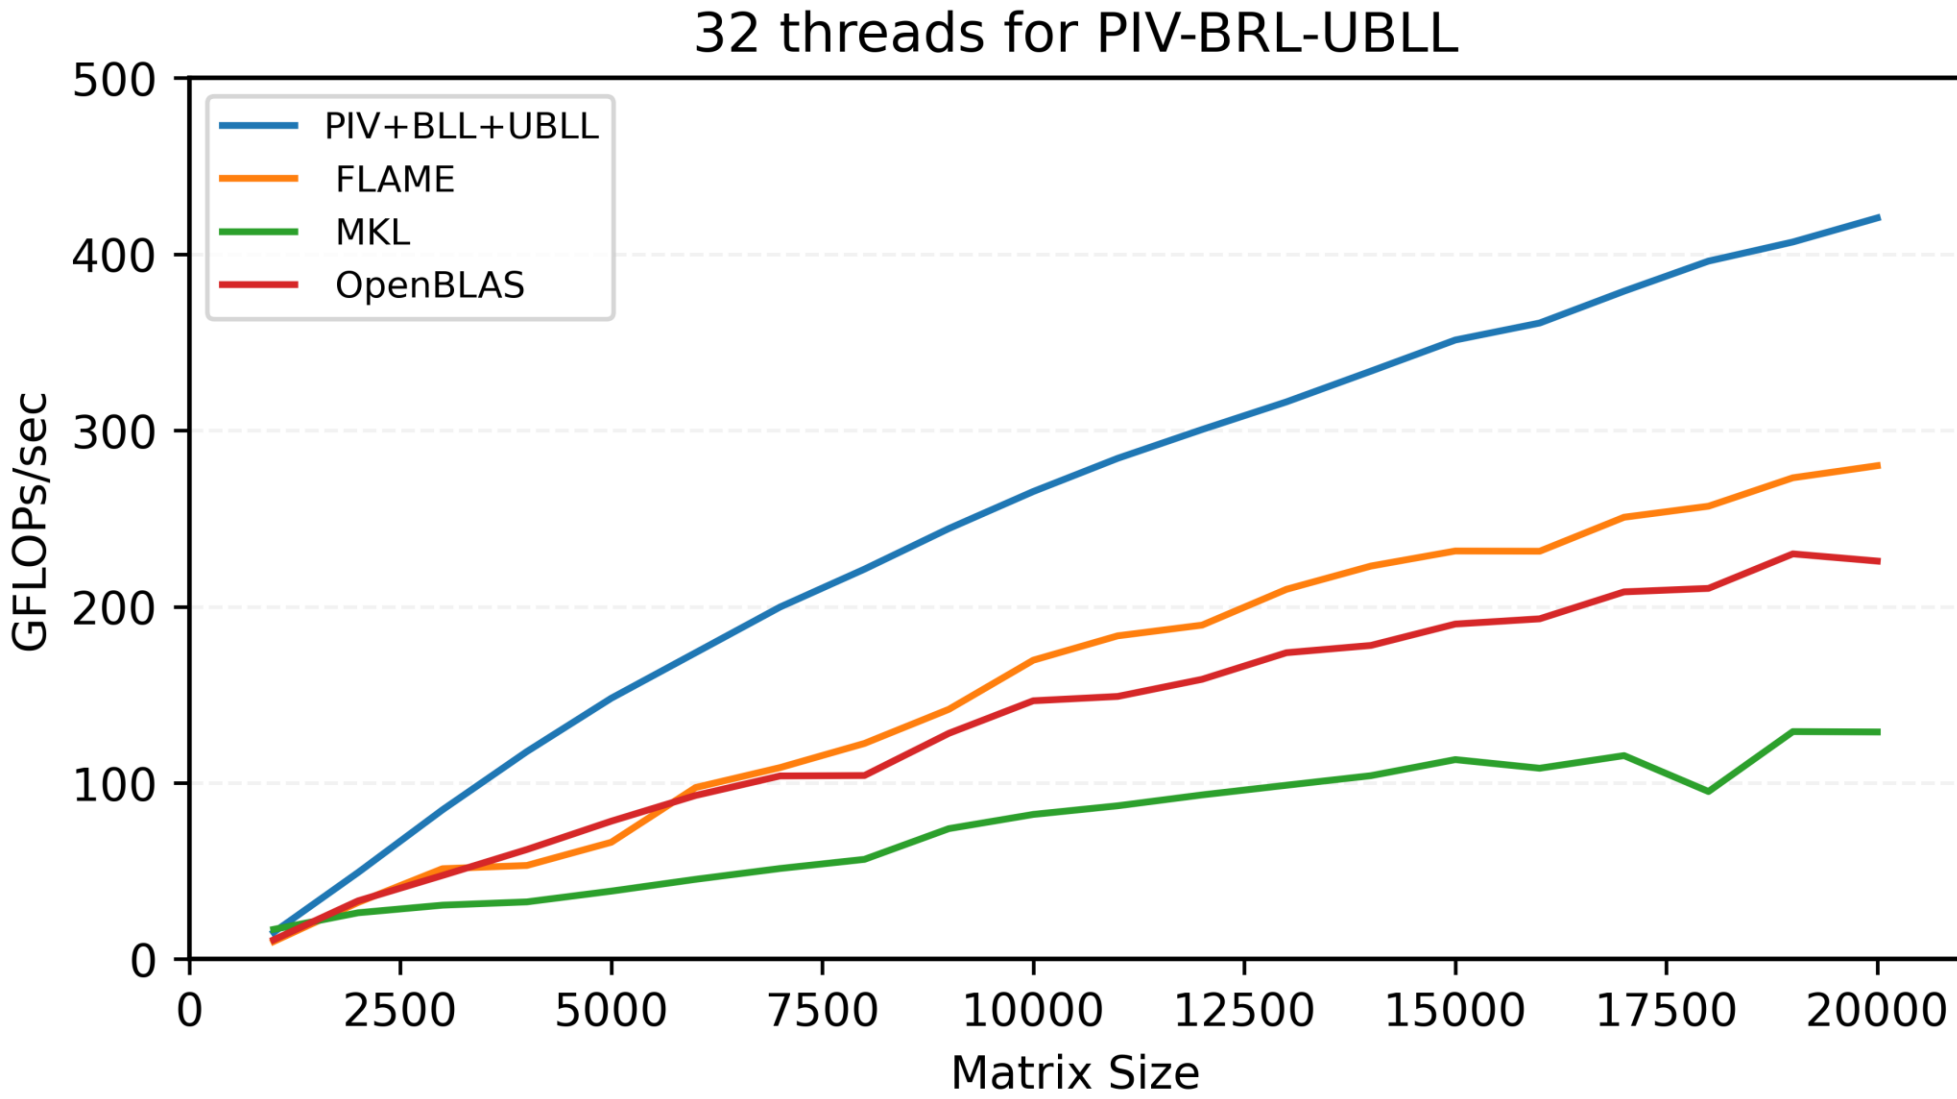

POTRF : computes the Cholesky factorization of a real symmetric positive definite matrix A

32 threads for blocked algorithms



Comparison with other vendors (PSTRF)

PSTRF : computes the Cholesky factorization with complete pivoting of a real symmetric positive semidefinite matrix

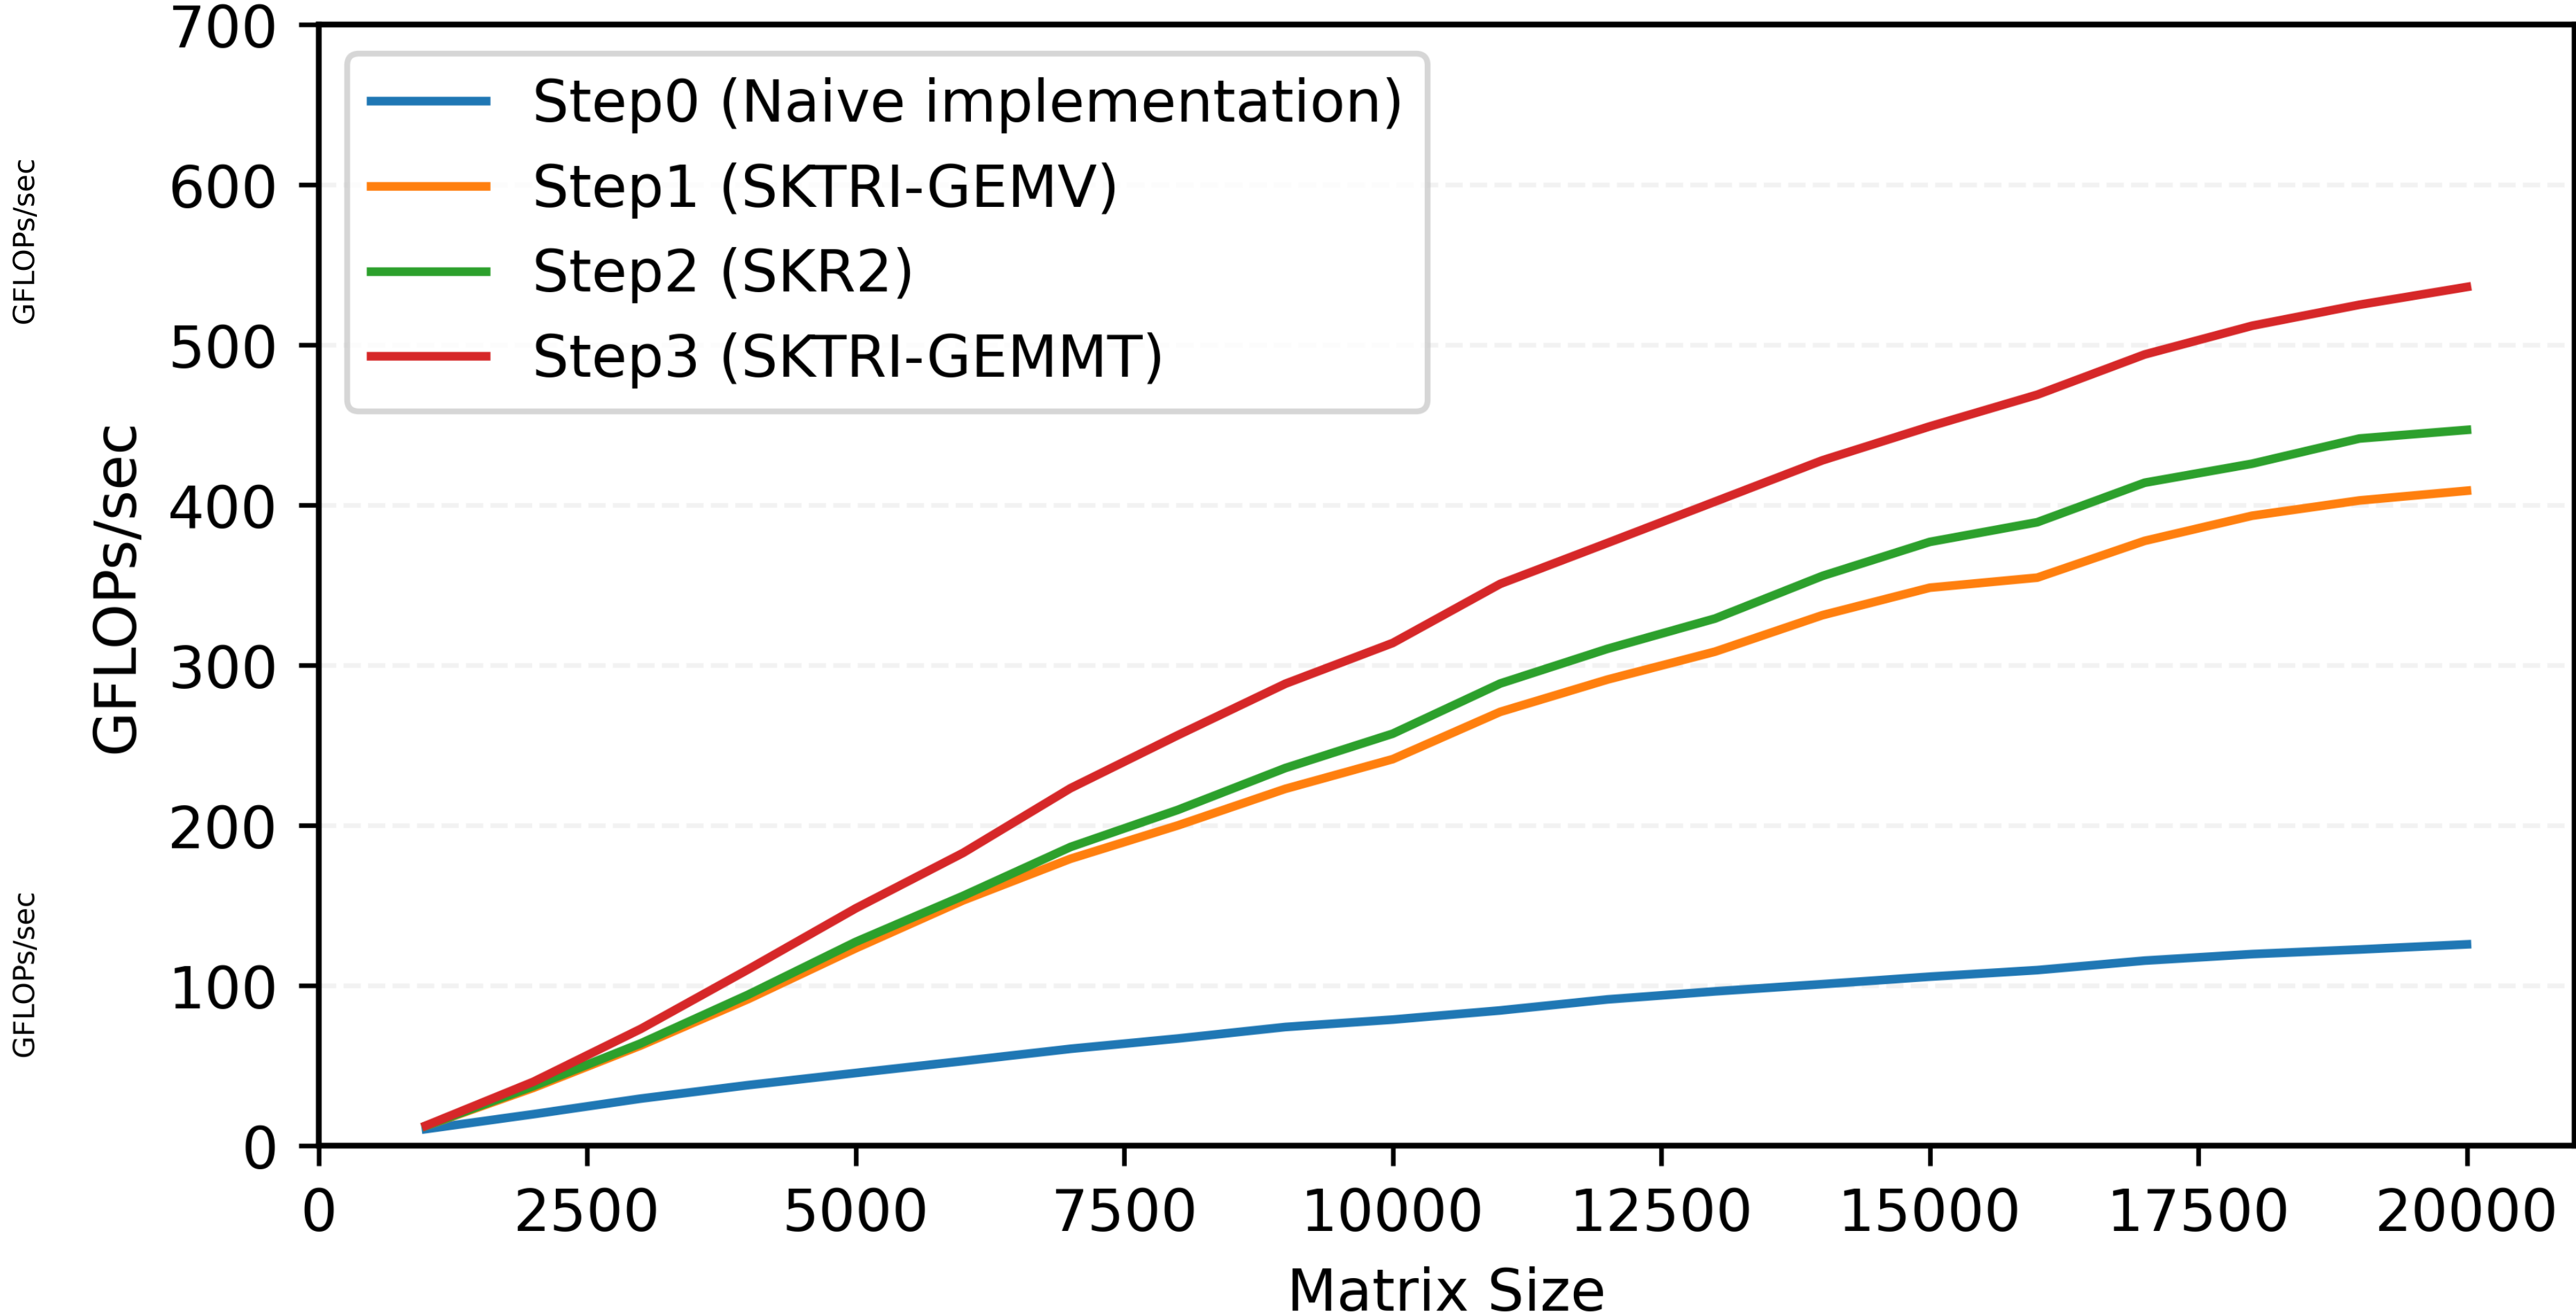


Optimization step by step

Optimizations for BRL + BLL step by step


- 1. SKTRI-GEMV
- 2. SKR2
- 3. SKTRI-GEMMT

BRL+UBLL (# of threads = 32, blocksize = 1024)



Summary

- By using FLAME methodology, we could systematically derive algorithms for LTLT decomposition of a Skew- symmetrical matrix
- Some new operations, which are not included in current BLAS, may probably needs to be optimized for those new algorithms to get high performance and BLIS could leverage those challenges.
- In this work, we optimize SKR2, GER2 and SKTRI-GEMV Level-2 operations, and two level-3 operations (SKTRI-GEMM, SKTRI-GEMMT)
- We profile the performance for 5 blocked algorithms and 4 unblocked algorithms including pivoting and unpivoting (BLL + UBLL is the best)
- BLIS-2.0 gives us opportunities to optimize those Level-2 operations.



I'm looking for a job!