Student Research Summary
Variations on Strassen-Winograd Algorithms for Rectangular Matrix-Multiplication

Sarah Loos
Department of Computer Science
Indiana University
150 S. Woodlawn, Room 215
Bloomington, IN
smloos@indiana.edu

ABSTRACT

The quadrant additions in the classical $O(n^{lg7})$, Strassen-Winograd matrix-multiplication algorithms are not ideal for preserving boundaries in properly rectangular multiplications. Cleaving an order-p vector into one of order $2^{\lceil \lg p \rceil - 1}$, and one of order $n - 2^{\lceil \lg p \rceil - 1}$, suggests a way to split an $m \times n$ rectangular matrix into quadrants. With $p = \max(m, n)$, the northwest is the only necessarily non-empty of the four quadrants. Adding it to other quadrants creates dense sums that, as factors, make more work for six of the seven recursive multiplications.

A very simple modification of the traditional algorithms, however, has only three recursions on factors built from that dense northwest block. The other four products are far more likely to have zero factors and be annihilated, accelerating the speed of multiplication.

Introduction

- The quadrant additions in the classical $O(n^{lg7})$, Strassen-Winograd matrix-multiplication algorithms are not ideal for preserving boundaries in properly rectangular multiplications. This is surprizing since these algorithms have been commonly used for over forty years.
- The difficulty was to modify Strassen's algorithm in a way that changed only which quadrants were added, and not the overall structure of the algorithm, to obtain a version which requires fewer computations on rectangular matrix-multiplications.
- There have been very few people who have attempted to modify Strassen's algorithm by changing which quadrants are added. None of them had recognized the positive effects of our modification when applied to rectangular matrices.

Approach

The new version of the C++ Matrix Template Library (MTL) includes tools for conveniently writing algorithms on dense matrices indexed as quadtrees and represented in hybrids of Morton order. I have used the MTL recursators to implement much more efficient matrix-multiplication.

Results

Experimental results plot total computations, time, L1 cache misses, L2 cache misses, and TLB misses for matrices of order 192 up to 4096. They compare among traditional Strassen-Winograd, the faster modified Strassen-Winograd, Douglas's benchmark implementation of Strassen-Winograd, classic recursion algorithm, and the BLAS3 dgemm from the manufacturer's ACML dgemm.

In spite of our simple and effective improvements to the $O(n^{lg7})$ algorithms, the race is still won by the classic $O(n^3)$ algorithm which uses the memory hierarchy much better.

Conclusion

Although the modified algorithms only win on certain properly rectangular matrix-multiplications, there is never a loss. So if a pure Strassen algorithm is to be used, it should be the modified version.

Opportunities for Undergraduate Research in Computer Science (OURCS)
Oct 2007, Pittsburgh, PA
.

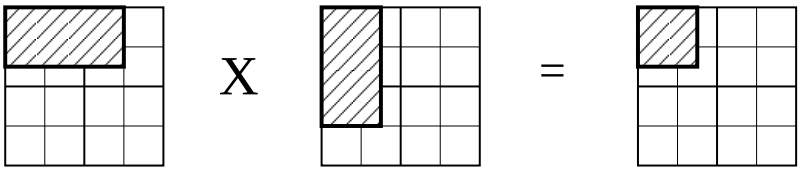

Step 1:

Modified Algo: 3 calls

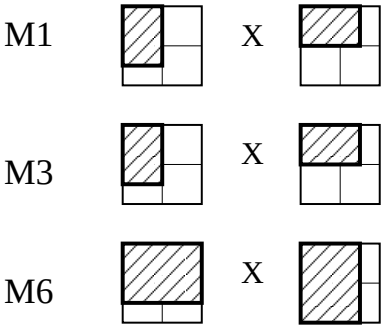

Original Algo: 3 calls

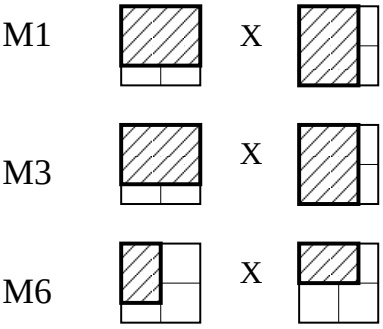

Step 2:

Modified Algo: 15 calls

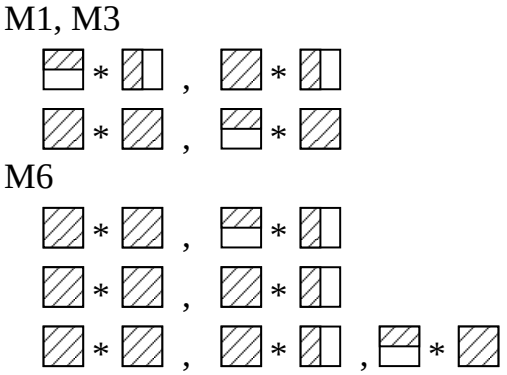

Original Algo: 18 calls

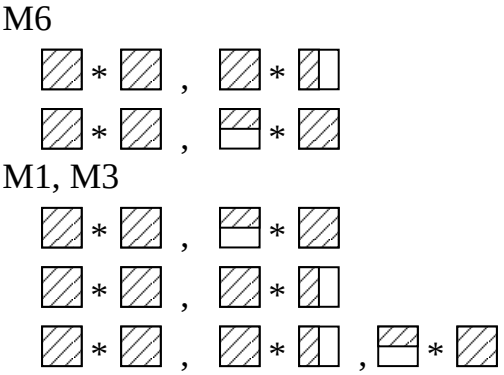

Step 3: (all full multiplies)

Modified Algo: 79 calls

Original Algo: 105 calls

Figure 1: Illustration of the recursive matrix multiply calls of both the modified and original Strassen algorithms for a rectangular matrix-multiplication.