# SINGLE-TRANSPOSE IMPLEMENTATION OF THE OUT-OF-ORDER 3D-FFT 

Alexander J. Yee

University of Illinois Urbana-Champaign

## The Problem

$\square$ FFTs are extremely memory-intensive.
$\square$ Completely bound by memory access.
$\square$ Memory bandwidth is always problem.

- Single-node shared memory: not enough bandwidth
- Multi-node: even worse
- Dominant factor in performance.
$\square$ Naïve implementations also bound by latency.
- Data-reordering can be many times slower than FFT computation itself!


## The Classic Approach to 3D-FFT

1. Perform x-dimension FFT.
2. In-memory transpose.
3. All-to-all communication.
4. Perform y-dimension FFT.
5. In-memory transpose.
6. All-to-all communication.
7. Perform z-dimension FFT.
8. In-memory transpose.
9. All-to-all communication.
$\square$ Exact order may differ.

- 3 all-to-all communication steps.
$\square$ An extra transpose may be needed at beginning to get data into order.



## What is an out-of-order FFT?

$\square \quad$ The Out-of-order FFT is mathematically the same as in-order FFT:

- Frequency domain is not in order.
$\square$ Forward Transform:
- Start from in-order time domain.
- End with out-of-order frequency domain.
- Use Decimation-in-Frequency algorithm.
$\square$ Inverse Transform:
- Start from out-of-order frequency domain.
- End with in-order time domain.
- Use Decimation-in-Time algorithm.
$\square$ Order of Frequency Domain:
- Bit-reversed is the most common.
- Other orders exist.
- Some algorithms are even faster - at the cost of further scrambling up the frequency domain.


Decimation-in-Frequency FFT
(Image taken from cnx.org)

## Why Out-of-Order?

$\square$ Many applications do not need an in-order frequency domain.

- Convolution
- Do not even need to look at Frequency Domain.
$\square$ Out-of-order FFT is faster:
$\square$ In-order FFTs require data-reordering -> bit-reversal
- Very poor memory access.
- Re-ordering is more expensive than FFT itself!
$\square$ In-order FFTs cannot be easily done in place.
- Requires double the memory of out-of-order FFT.
- Aggravates memory bottleneck.
$\square$ Out-of-order FFT can be several times faster!
$\square$ No need for final transpose for distributed FFTs over many nodes.


## Convolution via Out-of-order FFT


(lmages taken from cnx.org)

## Implementations of out-of-order FFT

$\square$ Prime95/MPrime - By George Woltman

- Used in GIMPs (Great Internet Mersenne Prime Search)
- World record holder for the largest prime number found. (August 2008)
- 9 of 10 largest known prime numbers found by GIMPS.
- Uses FFT for cyclic convolution.
- Fastest known out-of-order FFT. (x86-64 assembly for Windows + Linux)
$\square \quad$ y-cruncher Multi-threaded Pi Program - By Alexander J. Yee
- Fastest program to compute Pi and other constants.
$\square$ World Record holder for the most digits of Pi ever computed. (5 trillion digits - August 2010)
- Uses FFT and NTT for multiplying large numbers.
- Almost as fast as Prime95. (Standard C with Intel SSE Intrinsics - cross platform)
$\square \quad$ djbfft - By Daniel J. Bernstein
- One of the first implementations of out-of-order FFTs.
- Outperformed FFTW by factors $>3$ for convolution.
- Never widely-used, but motivated other out-of-order FFT projects.


## Our Approach to 3D-FFT

$\square$ Recognize that n-D FFT is same as 1D FFT.
$\square$ Different Twiddle Factors.

- Same Memory Access. Same Data-Flow.
$\square$ Implement a 1D-FFT with modified twiddle factors instead!
$\square$ All tricks for optimizing 1D-FFT are now available.
$\square$ Use Bailey's 4-step algorithm for 1D FFT.

1. First computation pass.
2. Matrix Transpose
3. Second computation pass.
4. Matrix Transpose data back to initial order.
$\square$ Out-of-order FFT -> No need for final transpose!
$\square$ Total: 1 transpose -> Only 1 all-to-all communication.

## Our Approach to 3D-FFT (cont.)

$\square$ To apply Bailey's 4-step method: Break the FFT into 2 passes.
$\square$ Easy way (slab decomposition):

- $x$ and $y$ into one pass.
- $z$ by itself in second pass.
- ( $y$ can go with $x$ or $z$ )
$\square$ Hard way (split dimension):
- Split the y dimension across the two passes.
- Overcomes scalability issue with $1^{\text {st }}$ method. (see next section)
$\square$ Both are being implemented.
$\square$ Frequency Domain will be Bit-reversed.



## Drawbacks

## Slab Decomposition

$\square \quad$ Can use standard FFT libraries.

- Optimal performance will still require custom sub-routines.
$\square$ Input data can be contiguous.
- Most common representation.
$\square$ \# of nodes must divide evenly into either x or z dimension.
- Scalability is limited to N nodes for $\mathrm{N}^{3}$ 3D-FFT
- Blue Waters will have more than 10,000 nodes...


## Splił Dimension

$\square$ Cannot use standard libraries.

- Everything must be written from scratch.
$\square$ Input data must be strided.
- Could imply extra transpose.
$\square$ \# of nodes must divide evenly into $x^{*} y$ or $x * z$.
- Scalability is limited to $\mathrm{N}^{3 / 2}$ nodes for $\mathrm{N}^{3}$ 3D-FFT.
- Not a problem on Blue Waters.


## Some Implementation Details

$\square$ All code written from scratch.

- No libraries.
- Everything is customized.


## Different Representations:

$$
\left\{r 0+i 0^{*} i, r 1+i 1^{*} i, r 2+i 2^{*} i, r 3+i 3^{*} i\right\}
$$

## Array of Structs (classic approach)



Complex Multiplication requires unpacking SSE3 addsubpd helps a little bit. But still slow.

Struct of Arrays


No unpacking needed.
When adjacent points need to operate: SSE3 Horizontal Instructions! (30\% faster)

## Benchmarks - Memory Bottleneck

## Complex Out-of-order 3D-FFT (10243) <br> Windows OpenMP - 16 GB needed <br> 2 x Intel Xeon X5482-64 GB DDR2



## Benchmarks - Shared Memory

Complex Out-of-order 3D-FFT (10243)
Slab Decomposition - 32 GB needed
2 x Intel Xeon X5482-64 GB DDR2


## Early Benchmarks - Distributed

## Complex Out-of-order 3D-FFT (10243)

Slab Decomposition - 32 GB needed
Accelerator Cluster - UIUC


## Analysis

$\square$ All-to-all communication steps reduced:

- Reduced from 3 to 1 for out-of-order FFT.
$\square$ In-order FFT doable by adding one transpose at end.
$\square$ Possibly communication optimal:
$\square$ No data is transferred more than once.
- Some data is never transferred at all.
- Hard to further reduce the \# of bytes transferred. (Is our current approach optimal?)
- Maybe possible to improve communication pattern instead?
$\square$ Lots of room for improvement within the node.
$\square$ FFT computation can be better optimized.
$\square$ Difficult to imagine more nodes than x or z dimension.
- Blue Waters: > 10,000 nodes
- 10,000 may be greater than one of the dimensions.
- May not be possible (or efficient) to use slab-decomposition.


## Next Steps

$\square$ Test current code on larger systems.
$\square$ Make sure the current implementation scales
$\square$ Port the code to PowerPC AltiVec.
$\square$ Currently implemented using x86-64 SSE3.
$\square$ Implement blocking and padding.
$\square$ Breaks cache associativity -> allows higher radix transforms.
$\square$ Overlapped communication and computation.
$\square$ Support for prime factors other than 2

- $3^{*} 2^{k}, 5^{*} 2^{k}$, and maybe $7^{*} 2^{k}$
$\square$ Real-input transforms.
$\square$ In-order FFT.


## Thanks for Listening

$\square$ Questions?

