

L10: Dense Linear Algebra on GPUs

CS6963



Administrative Issues

- Next assignment, linear algebra
 - Handed out by Friday
 - Due before spring break
 - handin cs6963 lab 3 <profile>”

CS6963

2
L10: Dense Linear Algebra

Outline

- Triangular solve assignment from last year (briefly)

Reading:

Paper: Volkov, V., and Demmel, J. W. 2008.
[Benchmarking GPUs to tune dense linear algebra](#), SC08,
 November 2008.

Paper link: <http://portal.acm.org/citation.cfm?id=1413402>

Talk link: <http://www.eecs.berkeley.edu/~volkov/volkov08-sc08talk.pdf>

Volkov code:

<http://forums.nvidia.com/index.php2/showtopic=47689&st=40&p=314014&#entry314014>

CS6963

3
L10: Dense Linear Algebra

Triangular Solve (STRSM)

```
for (j = 0; j < n; j++)
  for (k = 0; k < n; k++)
    if (B[j*n+k] != 0.0f) {
      for (i = k+1; i < n; i++)
        B[j*n+i] -= A[k*n+i] * B[j*n+k];
    }
```

Equivalent to:

```
cublasStrsm('l' /* left operator */, 'l' /* lower triangular */,
            'N' /* not transposed */, 'u' /* unit triangular */,
            N, N, alpha, d_A, N, d_B, N);
```

See: <http://www.netlib.org/blas/strsm.f>

CS6963

4
L10: Dense Linear Algebra

Last Year's Assignment

- Details:
 - Integrated with simpleCUBLAS test in SDK
 - Reference sequential version provided
- 1. Rewrite in CUDA
- 2. Compare performance with CUBLAS 2.0 library

CS6963

5
L10: Dense Linear Algebra

Symmetric Matrix Multiply (SSYMM)

```
float a[N][N], b[N][N], c[N][N];
float t1, t2;
for (j = 0; j < N; j++) {
  for (i = 0; i < N; i++) {
    t1 = b[i][i];
    t2 = 0;
    for (k = 0; k < i - 1; k++) {
      c[i][k] += t1 * a[i][k];
      t2 = t2 + b[k][i] * a[i][k];
    }
    c[i][i] += t1 * a[i][i] + t2;
  }
}
```

Equivalent to:
`cublasSsym('l' /* left operator */, 'u' /* upper triangular */, N, N, 1.0, d_A, N, d_B, N, 1.0, d_C, N);`

See: <http://www.netlib.org/blas/ssymm.f>



Performance Issues?

- + Abundant data reuse
- - Difficult edge cases
- - Different amounts of work for different $\langle j, k \rangle$ values
- - Complex mapping or load imbalance

CS6963

7
L10: Dense Linear Algebra

1. Project Proposal (due 3/7)

- Proposal Logistics:
 - Significant implementation, worth 55% of grade
 - Each person turns in the proposal (should be same as other team members)
- Proposal:
 - 3-4 page document (11pt, single-spaced)
 - Submit with handin program:
 - "handin cs6963 prop <pdf-file>"

CS6963

8
L10: Dense Linear Algebra

Decision Algorithm: Computational Decomposition

- Block Parallelism
 - for (i = 0; i < N; i++)
 - for (j = 0; j < N; j++)
 - a[i] = a[i] + c[j][i] * b[j];

No Dependence on i | Dependence on j

Block Candidate = i | Block Candidate = j

tile_by_index({"i"}, {T1}, {l1_control="ii", {"ii", "i", "j"}})

cudaize(block{"ii"}, thread{})

- Thread Parallelism
 - Thread Candidate = i
 - Thread Candidate = j

1D Block Setup(i) | 1D Thread Setup(i) | 2D Thread Setup(T_i, T_j)

Thread.X = i | Thread.X = j | Thread.Y = T_j

9
L10: Dense Linear Algebra
THE UNIVERSITY OF UTAH

Decision Algorithm: Data Staging

- Data Staging
 - Shared Memory
 - copy to shared memory
 - Data Reuse across threads
 - Registers
 - copy to registers
 - Data Reuse inside thread
- Final Loop Order
 - for (j = 0; j < N; j++)
 - for (i = 0; i < N; i++)
 - a[i] = a[i] + c[j][i] * b[j];
- Cudaize
 - Unrolling
 - cudaize(block{"ii"}, thread{"i"})
 - unroll to depth(1)

10
L10: Dense Linear Algebra
THE UNIVERSITY OF UTAH

CUDA-CHiLL Recipe

N = 1024
Ti = Tj = 32

```
tile_by_index({"i"}, {T1, Tj},
  {l1_control="ii",
  l2_control="k", {"ii", "jj", "i", "j"}})
normalize_index("i")
cudaize("mv_GPU", {a=N,
  b=N, c=N*N}, {block={"ii"},
  thread={"i"}})
```

copy_to_shared("b", "b", 1)
copy_to_registers("jj", "a")
unroll_to_depth(1)

11
L10: Dense Linear Algebra
THE UNIVERSITY OF UTAH

Matrix-Vector Multiply: GPU Code

Generated Code: with Computational decomposition only.

```
__global__ GPU_MV(float* a, float* b, float* c) {
  int bx = blockIdx.x; int tx = threadIdx.x;
  int i = 32*bx+tx;
  for (j = 0; j < N; j++)
    a[i] = a[i] + c[j][i] * b[j];
}
```

Final MV Generated Code: with Data staged in shared memory & registers.

```
__global__ GPU_MV(float* a, float* b, float* c) {
  int bx = blockIdx.x; int tx = threadIdx.x;
  __shared__ float bcopy[32];
  float acpy = a[tx + 32 * bx];
  for (jj = 0; jj < 32; jj++) {
    bcopy[tx] = b[32 * jj + tx];
    __syncthreads();
    //this loop is actually fully unrolled
    for (j = 32 * jj; j <= 32 * jj + 32; j++)
      acpy = acpy + c[j][32 * bx + tx] * bcopy[jj];
    __syncthreads();
  }
  a[tx + 32 * bx] = acpy;
}
```

Performance outperforms CUBLAS 3.2 (see later slide)

12
L10: Dense Linear Algebra
THE UNIVERSITY OF UTAH

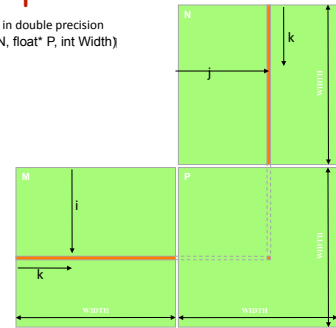
An added Complication

- What if input matrix *C* is transposed?
 $a[i] += c[i][j] * b[j];$
- What happens to global memory coalescing?
- What can be done?



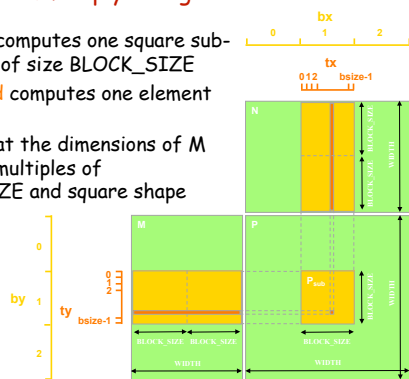
Reminder from L5: Matrix Multiplication in C

```
// Matrix multiplication on the (CPU) host in double precision
void MatrixMulOnHost(float* M, float* N, float* P, int Width)
{
  for (int i = 0; i < Width; ++i)
    for (int j = 0; j < Width; ++j) {
      double sum = 0;
      for (int k = 0; k < Width; ++k) {
        double a = M[i * width + k];
        double b = N[k * width + j];
        sum += a * b;
      }
      P[i * Width + j] = sum;
    }
}
```



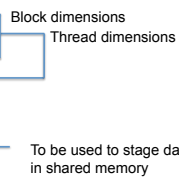
Tiled Matrix Multiply Using Thread Blocks

- One **block** computes one square sub-matrix P_{sub} of size `BLOCK_SIZE`
- One **thread** computes one element of P_{sub}
- Assume that the dimensions of *M* and *N* are multiples of `BLOCK_SIZE` and square shape



Strip-Mined Code

```
for (int ii = 0; ii < Width; ii+=TI)
  for (int i=ii; i<ii*TI-1; i++)
    for (int jj=0; jj<Width; jj+=TJ)
      for (int j = jj; j < jj*TJ-1; j++) {
        double sum = 0;
        for (int kk = 0; kk < Width; kk+=TK) {
          for (int k = kk; k < kk*TK-1; k++)
            sum += M[i][k] * N[k][j];
        }
        P[i][j] = sum;
      }
}
```



Final Code (from text, p. 87)

```

__global__ void MatrixMulKernel (float *Md, float *Nd, float *Pd, int Width) {
1.  __shared__ float Mds [TILE_WIDTH][TILE_WIDTH];
2.  __shared__ float Nds [TILE_WIDTH][TILE_WIDTH];
3 & 4.  int bx = blockDim.x; int by = blockDim.y; int tx = threadIdx.x; int ty = threadIdx.y;
//Identify the row and column of the Pd element to work on
5 & 6.  int Row = by * TILE_WIDTH + ty; int Col = bx * TILE_WIDTH + tx;
7.  float Pvalue = 0;
// Loop over the Md and Nd tiles required to compute the Pd element
8.  for (int m=0; m < Width / TILE_WIDTH; ++m) {
// Collaborative (parallel) loading of Md and Nd tiles into shared memory
9.  Mds [ty] [tx] = Md [Row*TILE_Width + (m*TILE_WIDTH + tx)];
10. Nds [ty] [tx] = Nd [(m*TILE_WIDTH + ty)*Width + Col];
11.  __syncthreads(); // make sure all threads have completed copy before calculation
12.  for (int k = 0; k < TILE_WIDTH; ++k) // Update Pvalue for TKxTK tiles in Mds and Nds
13.  Pvalue += Mds [ty] [k] * Nds [k] [tx];
14.  __syncthreads(); // make sure calculation complete before copying next tile
} // m loop
15. Pd [Row*Width + Col] = Pvalue;
}
    
```

LS: Memory Hierarchy, III

17



Preview: SGEMM (CUBLAS 2.x/3.x) on GPUs

- SGEMM result is not from algorithm of L5
- Why? Significant reuse can be managed within registers

	GPU Rule-of-Thumb	Lecture (for GTX 280)	Lecture (for Fermi C2050)
Threading	Generate lots of threads (up to 512/block) to hide memory latency	Only 64 threads/block provides 2 warps, sufficient to hide latency plus conserves registers	More cores/block, fewer registers/thread, so use 96 threads/block
Shared memory	Use to exploit reuse across threads	Communicate shared data across threads and coalesce global data	Communicate shared data across threads and coalesce global data
Registers	Use for temporary per-thread data	Exploit significant reuse within a thread	Exploit significant reuse within a thread
Texture memory	Not used	Not used	Increase bandwidth for global memory through parallel accesses

CS6963

18
L10: Dense Linear Algebra



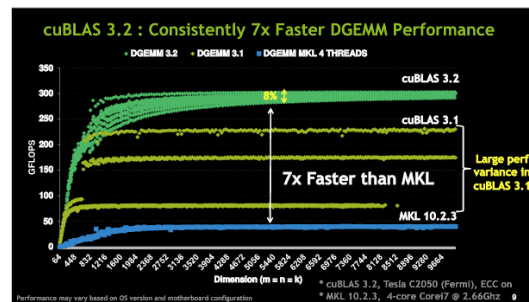
Volkov Slides 5-17, 24

CS6963

19
L10: Dense Linear Algebra



Comparison with MKL (Intel)



Slide source: <http://www.scribd.com/doc/47501296/CUDA-3-2-Math-Libraries-Performance>



CUDA-CHILL for Matrix Multiply (CUBLAS 2.x version)

```

init("mm.sp2", "MarkedLoop")
for (i = 0; i < n; i++)
  for (j = 0; j < n; j++)
    for (k = 0; k < n; k++)
      c[j][i] += a[k][i]*b[j][k];

tile_control({"i", "j"}, {TI, TJ},
  {l1_control="ii", l2_control="jj"},
  {"ii", "jj", "i", "j"})
tile_control({"k"}, {TK}, {l1_control="kk"},
  {"ii", "jj", "kk", "i", "j", "k"}, strided)
tile_control({"i"}, {TJ},
  {l1_control="ty", l1_tile="tx"},
  {"ii", "jj", "kk", "tx", "ty", "j", "k"})

--Assign loop levels to thread space and name the kernel
cudaize("mm_GPU",
  {a=N*N, b=N*N, c=N*N}, --array sizes for data copying
  {block={"ii", "jj"}, thread={"tx", "ty"}})
--Copy the "c" array usage to registers
copy_to_registers("kk", "c", {"tx", "ty"})
copy_to_shared("ty", "b")
--Unroll two innermost loop levels fully
unroll_to_depth(2)
    
```

Gabe Rudy Master's thesis

21
L10: Dense Linear Algebra

Final Result: Copy C to registers, B to shared memory and unroll

```

Steps 1 through 4 tile (for computation, data)
--Copy the "c" array usage to registers
5. copy_to_registers("kk", "c", {"tx", "ty"})
6. copy_to_shared("ty", "b")
--Unroll two innermost loop levels fully
7. unroll_to_depth(2)

float P1[16];
__shared__ float P2[16][17];
bx = blockDim.x, by = blockDim.y;
tx = threadIdx.x, ty = threadIdx.y;
P1[0:15] = c[16*by:16*by+15][tx+64*bx+16*ty];
for (t6 = 0; t10 <= 1008; t6+=16) {
  P2[tx][4*ty:4*ty+3] = b[16*by+4*ty:16*by+4*ty+3]
  [tx+t6];
  __syncthreads();
  P1[0:15] += a[t6][64*bx+16*ty+tx]*P2[0][0:15];
  P1[0:15] += a[t6+1][64*bx+16*ty+tx]*P2[1][0:15];
  ...
  P1[0:15] += a[t6+15][64*bx+16*ty+tx]*P2[15][0:15];
  __syncthreads();
}
c[16*by:16*by+15][tx+64*bx+16*ty] = P1[0:15];
    
```

B goes into shared memory

C goes into registers and is copied back at end

22
L10: Dense Linear Algebra

CUBLAS 3.x Example

```

1 tile_by_index({"i", "j"}, {TI, TJ},
  {l1_control="ii", l2_control="jj"},
  {"ii", "jj", "i", "j"})
2 tile_by_index({"k"}, {TK},
  {l1_control="kk"},
  {"ii", "jj", "kk", "i", "j", "k"},
  strided)
3 tile_by_index({"i"}, {TJ},
  {l1_control="tt", l1_tile="t"},
  {"ii", "jj", "kk", "t", "tt", "j", "k"})
4 cudaize("mm_GPU", {a=N*N, b=N*N, c=N*N},
  {block={"ii", "jj"}, threads={"t", "tt"}})
5 copy_to_registers("kk", "c")
6 copy_to_shared("tx", "b", -16)
7 copy_to_texture("b")
8 unroll_to_depth(2)

float P1[16];
__shared__ float P2[16][17];
tx = blockDim.x, ty = blockDim.y;
tx = threadIdx.x, ty = threadIdx.y;
P1[0:15] = c[16*ty:16*ty+15][tx+64*bx+16*ty];
for (t6 = 0; t10 <= 1008; t6+=16) {
  P2[tx][4*ty:4*ty+3] = tex1Dfetch(texRef_b, tx+16*ty+4*ty:16*ty+
  4*ty+3)+tx+t6);
  __syncthreads();
  P1[0:15] += a[t6][64*bx+16*ty+tx]*P2[0][0:15];
  P1[0:15] += a[t6+1][64*bx+16*ty+tx]*P2[1][0:15];
  ...
  P1[0:15] += a[t6+15][64*bx+16*ty+tx]*P2[15][0:15];
  __syncthreads();
}
c[16*ty:16*ty+15][tx+64*bx+16*ty] = P1[0:15];
    
```

THE UNIVERSITY OF UTAH

2D Convolution: CUDA-CHILL recipe and optimized code

```

Sequential Code
for(i=0; i<N; i++)
  for(j=0; j<N; j++)
    for(k=0; k<M; k++)
      c[i][j] = c[i][j] + a[k+i][j] * b[k][j];

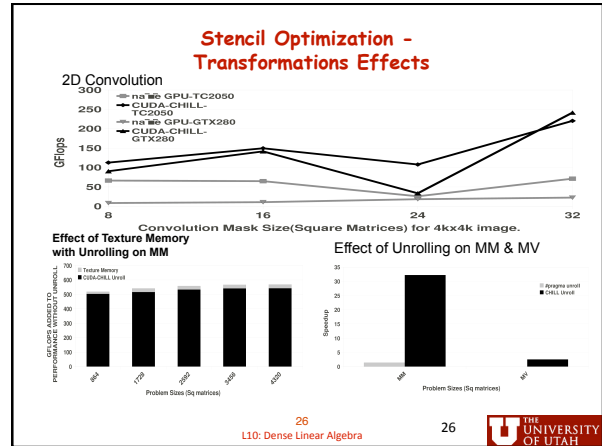
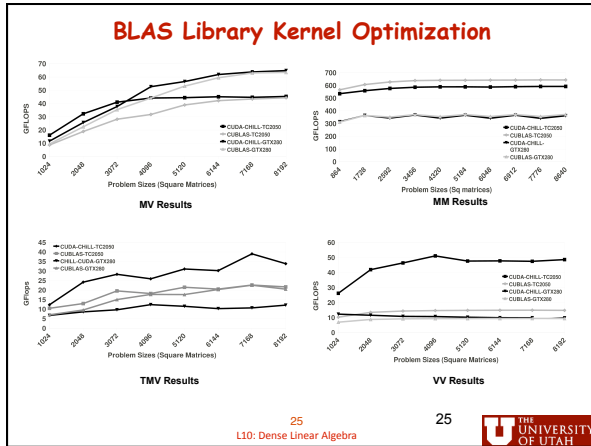
Complex bounds for shared memory copy loops

Optimized Code
__shared__ float _P1[47][31];
__shared__ float _P2[16][17]; float tmp3;
for (tmp = 16 * by + 3 * ty; tmp <= min(16 * by + 30, 16 *
  by + 3 * ty + 2); tmp++)
  for (tx1 = 2 * tx; tx1 <= min(2 * tx + 1, 46); tx1++)
    _P1[tx1][tmp - 16 * by] = a[tmp][32 * bx +
    tx1];
__syncthreads();
for (tx = 0; tmp <= 15; tmp++)
  for (tx1 = 2 * tx; tx1 <= 2 * tx + 1; tx1++)
    _P2[tx1][tmp] = b[tmp][tx1];
__syncthreads();
tmp3 = c[k + 16 * by][tx + 32 * bx];
for (l = 0; l <= 15; l++)
  tmp3 = tmp3 + _P1[l + tx][k + ty] * _P2[l][k];
c[k + 16 * by][tx + 32 * bx] = tmp3;
    
```

Complex bounds for shared memory copy loops

Data structures for shared memory

24
THE UNIVERSITY OF UTAH



Next Class

- Sparse linear algebra
- Appropriate data representations

27
L10: Dense Linear Algebra