Computer Architecture : A Programmer's Perspective

Abhishek Somani, Debdeep Mukhopadhyay

Mentor Graphics, IIT Kharagpur

September 9, 2016

Abhishek, Debdeep (IIT Kgp)

Comp. Architecture

Motivating Example

- 2 Memory Hierarchy
- 3 Parallelism in Single CPU
 - Dense Matrix Multiplication
 - The Problem
 - Analysis
 - Improvement
 - Better Cache utilization
- 5 Multicore Architectures
- 6 Appendix : Writing Efficient Serial Programs

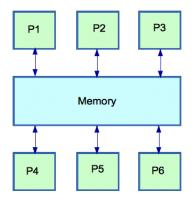
Outline

Motivating Example

- 2 Memory Hierarchy
- 3 Parallelism in Single CPU
- 4 Dense Matrix Multiplication
 - The Problem
 - Analysis
 - Improvement
 - Better Cache utilization
- 5 Multicore Architectures

6 Appendix : Writing Efficient Serial Programs

Communication Cost



- Communication cost in PRAM model : 1 unit per access
- Does it really hold in practice even within a single processor ?

Image: Image:

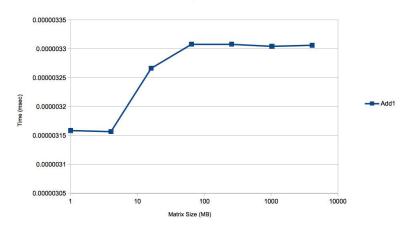
Add1

Add2

Abhishek, Debdeep (IIT Kgp)

æ

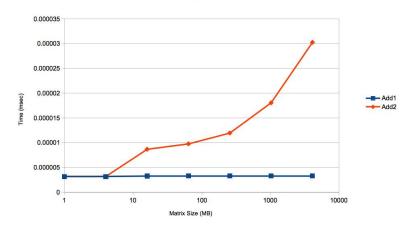
(日) (同) (三) (三)



Time per element

3

A B A B A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A



Time per element

∃ → (∃ →

э

A B > 4
 B > 4
 B

Outline

Motivating Example

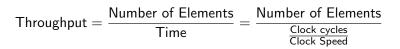
2 Memory Hierarchy

- 3 Parallelism in Single CPU
- 4 Dense Matrix Multiplication
 - The Problem
 - Analysis
 - Improvement
 - Better Cache utilization
- 5 Multicore Architectures

6 Appendix : Writing Efficient Serial Programs

Simple Addition

Image: A matrix and a matrix



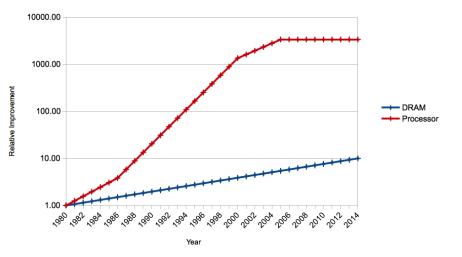
• For a fixed number of elements, how would stride impact throughput ?

• For a fixed stride, how would the number of elements impact throughput ?

Abhishek, Debdeep (IIT Kgp)

Performance Gap between Single Processor and DRAM

Relative Performance Improvement



- Clock Rate : 3.2 GHz
- Number of cores : 4
- Data Memory references per core per clock cycle : 2 64-bit references
- Peak Instruction Memory references per core per clock cycle : 1 128-bit reference
- Peak Memory bandwidth : 25.6 billion 64-bit data references + 12.8 billion 128-bit instruction references = 409.6 GB/s
- DRAM Peak bandwidth : 25 GB/s
- How is this gap managed ?

Memory Hierarchy

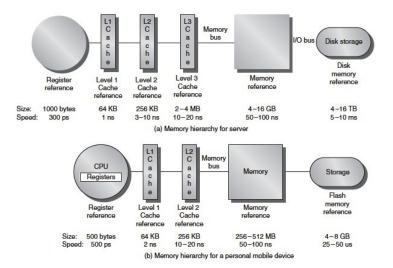


Figure : Courtesy of John L. Hennessey & David A. Patterson

Abhishek, Debdeep (IIT Kgp)

Comp. Architecture

September 9, 2016 13 / 96

3

・ロト ・ 日 ・ ・ ヨ ・ ・ ヨ ・

Memory Hierarchy in Intel Sandybridge

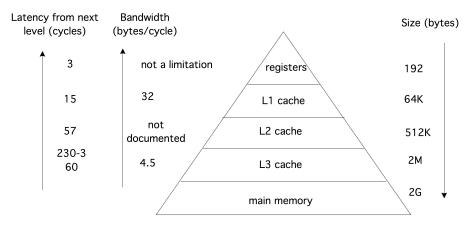


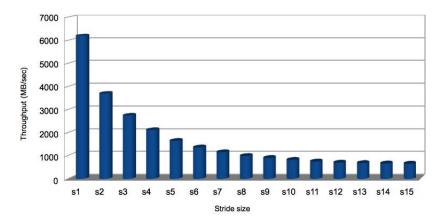
Figure : Courtesy of Victor Eijkhout

< 4 ► >

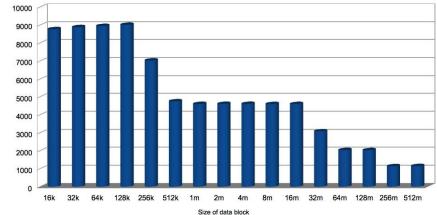
- Intel Xeon CPU E5-2697 v2
- Clock speed : 2.70GHz
- Number of processor cores : 24
- Caches :
 - L1D : 32 KB, L1I : 32 KB
 - Unified L2 : 256 KB
 - Unified L3 : 30720 KB
 - Line size : 64 Bytes
- 10.5.18.101, 10.5.18.102, 10.5.18.103, 10.5.18.104

Strided access

512 MB



Abhishek, Debdeep (IIT Kgp)



Regular access with stride 7

Throughput (MB/s)

- 一司

э

Outline

1 Motivating Example

2 Memory Hierarchy

Parallelism in Single CPU

- 4 Dense Matrix Multiplication
 - The Problem
 - Analysis
 - Improvement
 - Better Cache utilization
- 5 Multicore Architectures

O Appendix : Writing Efficient Serial Programs

- Factory Assembly Line analogy
- Fetch Decode Execute pipeline
- Improved throughput (instructions completed per unit time)
- Latency during initial "wind-up" phase
- Typical microprocessors have overall 10 35 pipeline stages
- Can the number of pipeline stages be increased indefinitely ?

- Pipeline depth : M
- Number of independent, subsequent operations : N

• Sequential time,
$$T_{seq} = MN$$

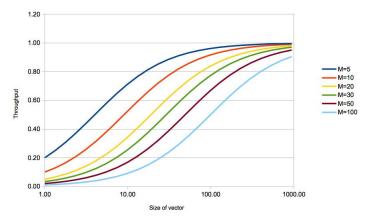
• Pipelined time,
$$T_{pipe} = M + N - 1$$

• Pipeline speedup,
$$\alpha = \frac{T_{seq}}{T_{pipe}} = \frac{MN}{M+N-1} = \frac{M}{1+\frac{M-1}{N}}$$

• Pipeline throughput,
$$p = \frac{N}{T_{pipe}} = \frac{N}{M+N-1} = \frac{1}{1+\frac{M-1}{N}}$$

Pipeline Throughput





æ

Pipeline Magic

Scale1

Scale2

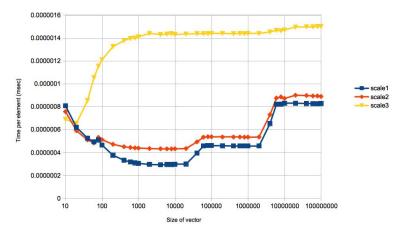
Scale3

3

<ロ> (日) (日) (日) (日) (日)

Influence of offset

scale1 : No offset, scale2 : +1 offset, scale3 : -1 offset



æ

∃ ► < ∃ ►</p>

A B A B A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A

Software Pipelining

• Pipelining can be effectively used for scale1 and scale2, but not scale3

- scale1 : Independent loop iterations
- scale2 : False dependency between loop iterations
- scale3 : Real dependency between loop iterations
- Software pipelining
 - Interleaving of instructions in different loop iterations
 - Usually done by the compiler
- Number of lines in assembly code generated by gcc under -O3 optimization
 - scale1 : 63
 - scale2 : 73
 - scale3 : 18

- Direct instruction-level parallelism
- Concurrent fetch and decode of multiple instructions
- Multiple floating-point pipelines can run in parallel
- Out-of-order execution and compiler optimization needed to properly exploit superscalarity
- Hard for compiler generated code to achieve more than 2-3 instructions per cycle
- Modern microprocessors are up to 6-way superscalar
- Very high performance may require assembly level programming

- Single Instruction Multiple Data
- Wide registers up to 512 bits
 - 16 integers
 - 16 floats
 - 8 doubles
- Intel : SSE, AMD : 3dNow!, etc.
- Advanced optimization options in recent compilers can generate relevant code to utilize SIMD
- Compiler intrinsics can be used to manually write SIMD code

Outline

1 Motivating Example

- 2 Memory Hierarchy
- 3 Parallelism in Single CPU
- Dense Matrix Multiplication
 - The Problem
 - Analysis
 - Improvement
 - Better Cache utilization
- 5 Multicore Architectures

O Appendix : Writing Efficient Serial Programs

Outline

1 Motivating Example

- 2 Memory Hierarchy
- 3 Parallelism in Single CPU
- Dense Matrix Multiplication
 - The Problem
 - Analysis
 - Improvement
 - Better Cache utilization
- 5 Multicore Architectures

6 Appendix : Writing Efficient Serial Programs

Why is matrix multiplication important?

RANK	SITE	SYSTEM	CORES	RMAX (TFLOP/S)	RPEAK (TFLOP/S)	POW (KW)
1	National Super Computer Center in Guangzhou China	Tianhe-2 [MilkyWay-2] - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 3151P NUDT	3,120,000	33,862.7	54,902.4	17,80
2	DOE/SC/Oak Ridge National Laboratory United States	Titan - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.	560,640	17,590.0	27,112.5	8,20
3	DOE/NNSA/LLNL United States	Sequoia - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1,572,864	17,173.2	20,132.7	7,89
4	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu	705,024	10,510.0	11,280.4	12,6
5	DOE/SC/Argonne National Laboratory United States	Mira - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	786,432	8,586.6	10,066.3	3,94
6	Swiss National Supercomputing Centre (CSCS) Switzerland	Piz Daint - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect , NVIDIA K20x Cray Inc.	115,984	6,271.0	7,788.9	2,32
7	King Abdullah University of Science and Technology Saudi Arabia	Shaheen II - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect Cray Inc.	196,608	5,537.0	7,235.2	2,83
8	Texas Advanced Computing Center/Univ. of Texas United States	Stampede - PowerEdge C8220, Xeon E5-2680 8C 2.700GHz, Infiniband FDR, Intel Xeon Phi SE10P Dell	462,462	5,168.1	8,520.1	4,51
9	Forschungszentrum Juelich (FZJ) Germany	JUQUEEN - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect IBM	458,752	5,008.9	5,872.0	2,30
hek, E	Debdeep (IIT Kgp)	Comp. Architecture		Se	ptember	9, 2

29 / 96

- Single array contains entire matrix
- Matrix arranged in row-major format
- m×n matrix contains m rows and n columns
- A(i,j) is the matrix entry at i^{th} row and j^{th} column of matrix **A**
- It is the $(i \times n + j)^{th}$ entry in the matrix array

Triple nested loop

```
void square_dgemm (int n, double* A, double* B, double* C)
{
   for (int i = 0; i < n; ++i)</pre>
    {
        const int iOffset = i*n;
        for (int j = 0; j < n; ++j)</pre>
        {
           double cij = 0.0;
            for( int k = 0; k < n; k++ )</pre>
                cii += A[iOffset+k] * B[k*n+j];
            C[iOffset+j] += cij;
        }
   }
}
```

• Total number of multiplications : n^3

Row-based data decomposition in matrix C

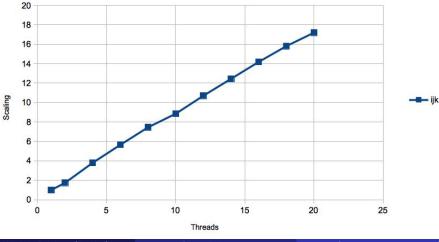
T1	1
Т2	10 11
	20 21
Т3	30
T4	31 40
Т5	41 50
Т6	51 60
Τ7	61
	70 71
T8	80

Abhishek, Debdeep (IIT Kgp)

```
void square_dgemm (int n, double* A, double* B, double* C)
Ł
#pragma omp parallel for schedule(static)
   for (int i = 0: i < n: ++i)</pre>
   ſ
       const int iOffset = i*n;
       for (int j = 0; j < n; ++j)</pre>
       ſ
           double cij = 0.0;
           for( int k = 0; k < n; k++ )</pre>
               cij += A[iOffset+k] * B[k*n+j];
           C[iOffset+j] += cij;
       }
   }
```

(Almost) Perfect Scaling for matrix of size 6000×6000

Matrix Multiplication Scaling



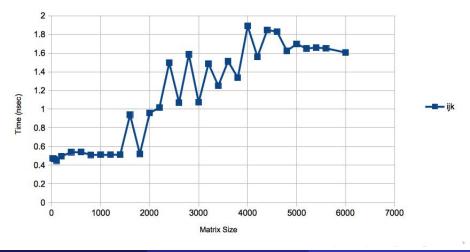
Abhishek, Debdeep (IIT Kgp)

Comp. Architecture

September 9, 2016 34 / 96

How good is the serial performance?

Time per scalar multiplication



- Normalized time becomes almost 4x when size of matrix grows from 1000 to 6000
- Experiments done on 3.2 GHz machine
- More than 5 clock cycles taken per double precision multiplication for 6000×6000 matrix !!!

Outline

1 Motivating Example

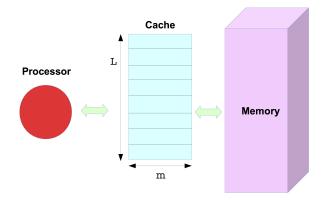
- 2 Memory Hierarchy
- 3 Parallelism in Single CPU

Dense Matrix Multiplication

- The Problem
- Analysis
- Improvement
- Better Cache utilization
- 5 Multicore Architectures

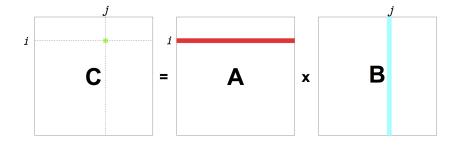
6 Appendix : Writing Efficient Serial Programs

Memory Hierarchy Model for analysis



- L lines of capacity m double precision numbers each
- Tall Cache assumption : L > m
- Replacement Policy : Least Recently Used
- No Hardware Prefetching

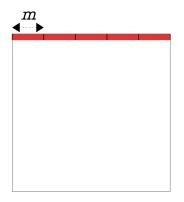
Memory Access Pattern during multiplication



• A, B and C are square matrices of size $n \times n$

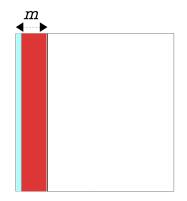
• **n** is large, i.e., n > L

Memory Access Pattern for A



• Sequential access : Accessing a row requires $\frac{n}{m}$ cache misses

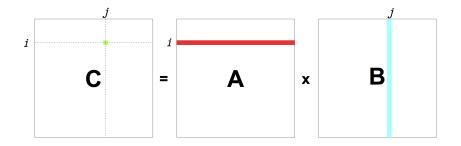
Memory Access Pattern for B



• Strided access : Accessing a column requires n cache misses

Abhishek, Debdeep (IIT Kgp)

Total cache misses



- For computing every C(i,j), the number of cache misses : $1 + \frac{n}{m} + n$
- If n < mL, total cache misses : $\frac{2n^2}{m} + n^3$
- If n > mL, total cache misses : $\frac{n^2}{m} + \frac{n^3}{m} + n^3$

- 64 bytes cache line size means m = 8
- 256 KB L2 cache means mL = 32768
- For practical problems n < 10 15k
- Thus n < mL and the total cache misses : $\frac{2n^2}{m} + n^3 = \Theta(n^3)$
- Can this be improved ?

Outline

1 Motivating Example

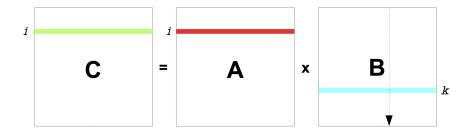
- 2 Memory Hierarchy
- 3 Parallelism in Single CPU

Dense Matrix Multiplication

- The Problem
- Analysis
- Improvement
- Better Cache utilization
- 5 Multicore Architectures

6 Appendix : Writing Efficient Serial Programs

Alternate Memory Access Pattern



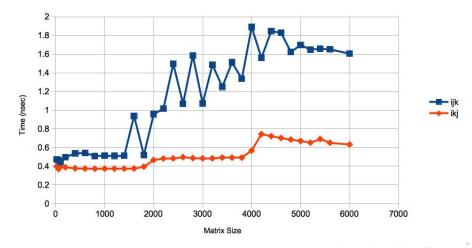
• For computing the row C(i,:), cache misses : $\frac{2n}{m} + \frac{n^2}{m}$

• Total cache misses :
$$\frac{2n^2}{m} + \frac{n^3}{m} = \Theta(\frac{n^3}{m})$$

```
void square_dgemm (int n, double* A, double* B, double* C)
ł
   for (int i = 0; i < n; ++i)</pre>
   {
       const int iOffset = i*n;
       for( int k = 0; k < n; k++ )
       {
           const int kOffset = k*n:
           for (int j = 0; j < n; ++j)</pre>
               C[iOffset+j] += A[iOffset+j] * B[kOffset+j];
       }
   }
}
```

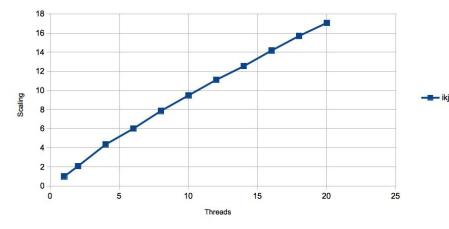
• Triple-nested loop with the order of (i, j, k) changed to (i, k, j)

Time per scalar multiplication



(Almost) Perfect Scaling for matrix of size 6000 imes 6000





Outline

1 Motivating Example

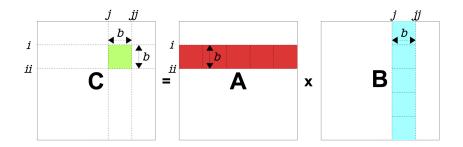
- 2 Memory Hierarchy
- 3 Parallelism in Single CPU

Dense Matrix Multiplication

- The Problem
- Analysis
- Improvement
- Better Cache utilization
- 5 Multicore Architectures

6 Appendix : Writing Efficient Serial Programs

Blocking / Tiling



- Assumptions for analysis : n%b = 0 and b%m = 0
- Cache misses in loading a block : $\frac{b^2}{m}$
- Cache misses in finding a block of C : $\frac{b^2}{m} + \frac{b^2}{m} \frac{2n}{b} = \frac{b^2}{m} + \frac{2nb}{m}$
- Total cache misses : $\frac{n^2}{b^2}(\frac{b^2}{m} + \frac{2nb}{m}) = \frac{n^2}{m} + \frac{2n^3}{mb} = \Theta(\frac{n^3}{mb})$

• The 3 blocks for A, B and C should just fit in the cache

•
$$3b^2 = mL$$
, i.e., $b = \sqrt{\frac{mL}{3}}$

- For L1 cache of capacity 32KB, mL = 4096 and b = 36.95
- A good value for b is 32
- Total cache misses : $\frac{2n^3}{mb} = \frac{2\sqrt{3}n^3}{m\sqrt{mL}} = \Theta(\frac{n^3}{m\sqrt{\text{Cache Size}}})$

Tiled Multiply

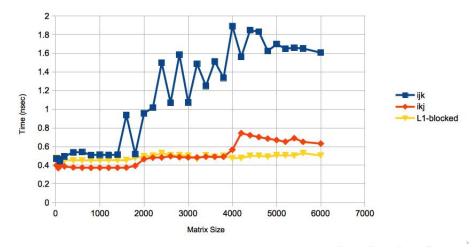
```
void square_dgemm (int n, double* A, double* B, double* C)
{
   for (int i = 0; i < n; i += BLOCK_SIZE)</pre>
   {
       const int iOffset = i * n;
       for (int j = 0; j < n; j += BLOCK_SIZE)
           for (int k = 0; k < n; k += BLOCK_SIZE)
           {
               /* Correct block dimensions if block "goes off
                   edge of " the matrix */
               int M = min (BLOCK_SIZE, n-i);
               int N = min (BLOCK_SIZE, n-j);
               int K = min (BLOCK_SIZE, n-k);
               /* Perform individual block dgemm */
              do_block(n, M, N, K, A + iOffset + k,
                      B + k*n + j, C + iOffset + j);
           }
```

}

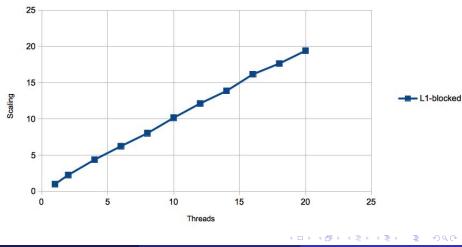
```
static void do_block (int n, int M, int N, int K, double* A,
    double* B, double* C)
ł
   for (int i = 0; i < M; ++i)</pre>
   ſ
       const int iOffset = i*n;
       for (int j = 0; j < N; ++j)</pre>
       ſ
           double cij = 0.0;
           for (int k = 0; k < K; ++k)
               cij += A[iOffset+k] * B[k*n+j];
           C[iOffset+j] += cij;
       }
   }
```

Tiled versus Normal

Time per scalar multiplication



Multi-threaded Scaling



Abhishek, Debdeep (IIT Kgp)

Comp. Architecture

September 9, 2016 55 / 96

- Given that we have made the data being worked upon available in the cache closest to the processor, we could use some ILP
- ILP kicks in when there is significant amount of independent work available in a single block of code
- Loop unrolling can help us achieve that
- Compilers also unroll loop but in this case there are too many nesting levels for the compiler to do the correct thing

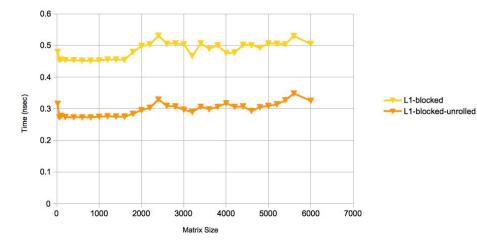
Tiled Multiply with unrolling

```
for (int k = 0; k < K; ++k)
    cij += A[iOffset+k] * B[k*n+j];</pre>
```

```
for (int k = 0; k < K; k += 8)
{
   const double d0 = A[iOffset+k] * B[k*n+j];
   const double d1 = A[iOffset+k+1] * B[(k+1)*n+j];
   const double d2 = A[iOffset+k+2] * B[(k+2)*n+j];
   const double d3 = A[iOffset+k+3] * B[(k+3)*n+j];
   const double d4 = A[iOffset+k+4] * B[(k+4)*n+j];
   const double d5 = A[iOffset+k+5] * B[(k+5)*n+j];
   const double d6 = A[iOffset+k+6] * B[(k+6)*n+j];
   const double d7 = A[iOffset+k+7] * B[(k+7)*n+j];
   cij += (d0 + d1 + d2 + d3 + d4 + d5 + d6 + d7);
}
```

Tiled Multiply with unrolling...

Time per scalar multiplication



- In addition to L1, blocking can be done for the L2 cache also \Rightarrow 2-level tiled code
- Next programming assignment

- Manual optimization and tuning is tedious and error-prone
- Entire process needs to be redone in full for any new architecture
- Multi-threaded optimization adds further complexity
- Code generated automatically by parameterized code generators
 - Automatically Tuned Linear Algebra Software
 - Portable High Performance ANSI C
- Essentially a search problem

Outline

1 Motivating Example

- 2 Memory Hierarchy
- 3 Parallelism in Single CPU
- 4 Dense Matrix Multiplication
 - The Problem
 - Analysis
 - Improvement
 - Better Cache utilization

5 Multicore Architectures

Appendix : Writing Efficient Serial Programs

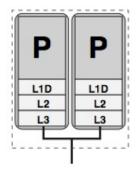


Figure : Courtesy of G. Hager & G. Wellein

- Each core has it's own cache for all levels
- eg., Intel Montecito

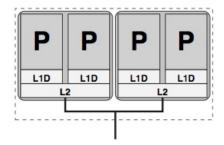


Figure : Courtesy of G. Hager & G. Wellein

- Separate L1, Shared L2 (2 dual-core L2 groups)
- Shared cache enables inter-core communication without going to the main memory
- Reduced latency and improved bandwidth
- eg., Intel Harpertown

Hexa Core

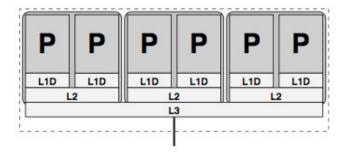


Figure : Courtesy of G. Hager & G. Wellein

- 6 single-core L1 groups, 3 dual-core L2 groups
- L3 shared for all cores
- Cache bandwidth shared across number of cores connected
- eg., Intel Dunnington

Uniform Memory Access (UMA)

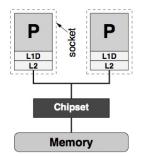


Figure : Courtesy of G. Hager & G. Wellein

- 2 single-core CPUs share a common FrontSide bus (FSB)
- Arbitration protocols built into the CPUs
- \bullet Chipset connects to memory and other I/O systems
- $\bullet\,$ Data can be transfered to/from only one CPU at a time

Uniform Memory Access...

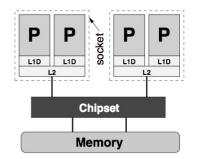


Figure : Courtesy of G. Hager & G. Wellein

- FSB not shared by sockets
- Role of chipset becomes more important
- Anisotropic system Cores on same socket are "closer" than those on other sockets

Abhishek, Debdeep (IIT Kgp)

Integrated Memory Controller

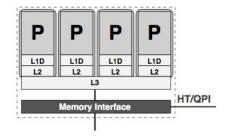


Figure : Courtesy of G. Hager & G. Wellein

- Integrated memory controller allows direct connection to memory and/or other sockets
- Intel QuickPath (QPI), AMD HyperTransport (HT)
- eg. Intel Nehalem, AMD Shanghai

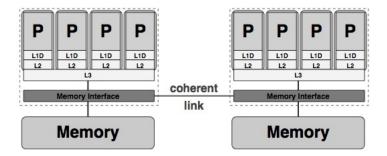


Figure : Courtesy of G. Hager & G. Wellein

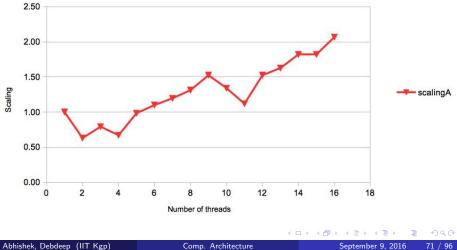
- cache-coherent Non Uniform Memory Access
- Every UMA building block is a Locality Domain (LD)
- Provides scalable bandwidth for large number of processors

- Explicit logic required to maintain cache coherence
- MESI protocol
 - Modified : Cache line modified in this cache, resides in no other cache
 - Exclusive : Read from memory, not modified yet, resides in no other cache
 - Shared : Read from memory, not modified yet, may reside in other caches
 - Invalid : Data in cache line is garbage
- Cache coherence traffic can hurt application performance if same cache line is modified frequently by different locality domains (false sharing).

Back to π

```
const double deltaX = 1.0/(double)numPoints:
   double pi = 0.0;
   omp_set_num_threads(numThreads);
   double components[numThreads];
   for(int i = 0; i < numThreads; ++i)</pre>
       components[i] = 0.0;
#pragma omp parallel shared(components)
    Ł
       const int nt = omp_get_num_threads();
       const int pointsPerThread = numPoints/nt:
       const int threadId = omp_get_thread_num();
       double xi = (0.5 + pointsPerThread * threadId) * deltaX;
       for(int i = 0; i < pointsPerThread; ++i)</pre>
       Ł
           components[threadId] += 4.0/(1 + xi * xi);
           xi += deltaX:
       }
   3
   for(int i = 0; i < numThreads; ++i)</pre>
       pi += components[i];
   pi *= deltaX:
```

Multi-threaded scaling



Abhishek, Debdeep (IIT Kgp)

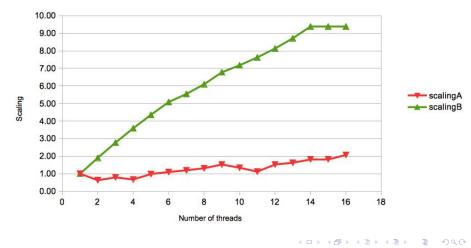
Comp. Architecture

September 9, 2016

Back to π ...

```
const double deltaX = 1.0/(double)numPoints:
   double pi = 0.0;
   omp_set_num_threads(numThreads);
   double components[numThreads];
#pragma omp parallel shared(components)
       const int nt = omp_get_num_threads();
       const int pointsPerThread = numPoints/nt;
       double myComponent = 0.0;
       const int threadId = omp_get_thread_num();
       double xi = (0.5 + pointsPerThread * threadId) * deltaX;
       for(int i = 0; i < pointsPerThread; ++i)</pre>
       ſ
           myComponent += 4.0/(1 + xi * xi);
           xi += deltaX:
       3
       components[threadId] = myComponent;
   3
   for(int i = 0; i < numThreads; ++i)</pre>
       pi += components[i];
   pi *= deltaX:
```

Multi-threaded scaling



numactl --hardware

available: 2 nodes (0-1) node 0 cpus: 0 2 4 6 8 10 12 14 16 18 20 22 24 26 28 30 32 34 36 38 40 42 44 46 node 0 size: 131026 MB node 0 free: 124290 MB node 1 cpus: 1 3 5 7 9 11 13 15 17 19 21 23 25 27 29 31 33 35 37 39 41 43 45 47 node 1 size: 131072 MB node 1 free: 126752 MB node distances: node 0 1 0: 10 20 1: 20 10

Highly scalable ccNUMA

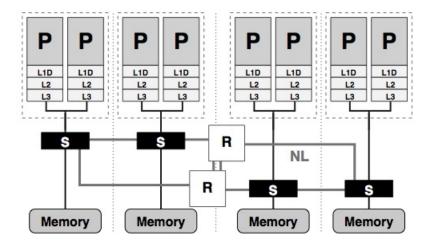


Figure : Courtesy of G. Hager & G. Wellein

Abhishek, Debdeep (IIT Kgp)

э

< 🗇 🕨

Outline

1 Motivating Example

- 2 Memory Hierarchy
- 3 Parallelism in Single CPU
- 4 Dense Matrix Multiplication
 - The Problem
 - Analysis
 - Improvement
 - Better Cache utilization
- 5 Multicore Architectures

6 Appendix : Writing Efficient Serial Programs

- Compiler modifies each function call to log the number of calls, its callers and the time taken
- Best suited when each function call takes significant time
- Overhead significant if many functions with short runtime
- eg., gprof from GNU binutils package

- Program is sampled at regular intervals and program counter and current call stack are recorded
- Program needs to run long enough for results to be accurate
- Possible to get profiling information down to the source and assembly level
- eg., gperftools from Google, Vtune Amplifier from Intel

- Special on-chip registers which get incremented every time a certain event occurs
- Example events
 - Bus transactions
 - Mis-predicted branches
 - Cache Misses at various levels
 - Pipeline stalls
 - Number of loads and stores
 - Number of instructions executed
- eg., Vtune Amplifier from Intel, oprofile, PAPI

- Sequential access of data in arrays
- If arr[i][j] is in the cache, arr[i][j + 1] is likely to be in the cache also. However, arr[i + 1][j] is NOT likely to be in the cache
- Avoid using nested containers like vector of vectors for storing matrices
- Redesign data-structures for locality of access

Optimize Memory Access...

```
const int size = 10000;
int a[size], b[size];
for(int i = 0; i < size; ++i)
{
    b[i] = func(a[i]);
}
```

```
typedef struct { int a; int b;} myPair;
myPair ab[size];
for(int i = 0; i < size; ++i)
{
    ab[i].b = func(ab[i].a);
}
```

- Use inline functions for short functions.
- Replace long if...else if...else if... chains by switch statements.
 - Such chains may lead to frequent branch mis-prediction.
 - Pipelined architectures incur severe cost (15-20 cycles) for every mis-predicted branch.
 - Compiler may optimize switch into a table lookup requiring a single jump.
 - If converting to switch is not possible, put the most common clauses at the beginning of the if chain.
- Where applicable, replace deeply recursive functions by iterative ones. eg., BFS, DFS.

- Most modern servers have 4-way superscalar cores.
- Blocks of code (eg., in the body of a loop) should have enough independent instructions.
- Unrolling of loops may help in achieving this.
- Inlined functions (small ones) also help, better register optimization being the other benefit.

Loop Unrolling Example

```
for(int i = 0; i < 100; ++i)
{
    if(i % 2 == 0)
        func1(i);
    else
        func2(i);
    func3(i);
}</pre>
```

```
for(int i = 0; i < 100; i += 2)
{
    func1(i);
    func3(i);
    func2(i+1);
    func3(i+1);
}</pre>
```

э

• • • • • • • •

- First and foremost job : Correct and reliable mapping of high-level source code to machine code
- Major code transformation areas
 - Function inlining
 - Constant folding
 - Constant propagation
 - Common subexpression elimination
 - Register variables
 - Branch analysis
 - Loop analysis
 - Algebraic Reduction

Function Inlining

```
double square (double a)
{
    return a * a;
}
double parabola (double b)
{
    return square(b) + 1.0;
}
```

```
double parabola (double b)
{
    return b * b + 1.0;
}
```

double x, y, z; y = x * (17.0/19.0); z = x * 17.0 / 19.0;

double x, y, z; y = x * 0.89473684210526316374; z = x * 17.0 / 19.0;

イロト イポト イヨト イヨト 二日

```
double parabola (double b)
{
    return b * b + 1.0;
}
double x, y;
x = parabola( 13.5 );
y = x * 2.3;
```

double x, y; x = 183.25; y = 421.475;

3

```
double a, b, c, d;
b = (a + 6.0);
c = (a + b) * (a + b);
d = (a + b) / a;
```

```
double a, b, c, d, temp;
temp = a + a + 6.0;
c = temp * temp;
d = temp / a;
```

Branch Analysis : Join identical branches

```
double x, y, z;
bool b;
if(b)
{
    y = parabola(x);
    z = y + 4.0;
}
else
{
    y = square(x);
    z = y + 4.0;
}
...
```

if(b)
{
 y = parabola(x);
}
else
{
 y = square(x);
}
z = y + 4.0;
...

3

< ロ > < 同 > < 三 > < 三

Branch Analysis : Eliminate jumps

```
int foo (int a, bool b)
{
    if(b)
        a = a * 4;
    else
        a = a * 5;
    return a;
}
```

```
int foo (int a, bool b)
{
    if(b)
    {
        a = a * 4;
        return a;
    }
    else
    {
        a = a * 5;
        return a;
    }
}
```

3

< ロ > < 同 > < 回 > < 回 > < 回 > < 回

- Loop unrolling
- Loop invariant code motion
- Instructions reordering and scheduling
- Pointer elimination

• ...

Option	Details	Number of optimization flags
-00	Default, Fast compilation and	
	low memory usage during compilation	0
-01	Quick and light transformations	
	that preserve execution ordering	39
-02	More optimizations with instruction	
	reordering and inlining	83
-03	Heavy duty optimizations with a	
	lot of transformations	93
-Ofast	O3 with fast, standards incompliant	
	floating point calculations	94
-Os	Optimize for size of executable	66

- Compilers can not optimize across modules
- Declare objects and fixed size arrays (not very large) inside functions that need them; Avoid dynamic memory allocation
- Write programs to access data in arrays sequentially; Compilers can not do this transformation
- Use <u>__restrict</u> keyword for pointers when the program logic rules out pointer aliasing
- ALWAYS PROFILE AFTER MAKING A SIGNIFICANT CHANGE

Pointer Aliasing

```
void foo (int * a, int * p)
{
   for( int i = 0; i < 1000; ++i)</pre>
       a[i] = *p + 2;
}
void func1 ()
ł
    int arr[1000];
   foo(arr, &arr[10]);
}
void func2 ()
{
    int arr[1000], b = 20;
   foo(arr, &b);
}
```

< 67 ▶

- Introduction to High Performance Computing for Scientists and Engineers Hager, Wellein : Chapter 1, 2, 3
- http://agner.org/optimize/optimizing_cpp.pdf