

Program Optimization

Spring 2012

Instructors:

Aykut & Erkut Erdem

Acknowledgement: The course slides are adapted from the slides prepared by R.E. Bryant, D.R. O'Hallaron, G. Kesden and Markus Püschel of Carnegie-Mellon Univ.

Today

- **Overview**
- **Program optimization**
 - Code motion/precomputation
 - Strength reduction
 - Sharing of common subexpressions
 - Optimization blocker: Procedure calls
 - Optimization blocker: Memory aliasing
- **Exploiting Instruction-Level Parallelism**

Harsh Reality

- *There's more to runtime performance than asymptotic complexity*
- *One can easily loose 10x, 100x in runtime or even more*
- **What matters:**
 - Constants ($100n$ and $5n$ is both $O(n)$, but)
 - Coding style (unnecessary procedure calls, unrolling, reordering, ...)
 - Algorithm structure (locality, instruction level parallelism, ...)
 - Data representation (complicated structs or simple arrays)

Harsh Reality

■ Must optimize at multiple levels:

- Algorithm
- Data representations
- Procedures
- Loops

■ Must understand system to optimize performance

- How programs are compiled and executed
 - Execution units, memory hierarchy
- How to measure program performance and identify bottlenecks
- How to improve performance without destroying code modularity and generality

Optimizing Compilers

-O

- Use optimization flags, **default is no optimization (-O0)**!
- Good choices for gcc: -O2, -O3, -march=xxx, -m64
- Try different flags and maybe different compilers

Example

```
double a[4][4];
double b[4][4];
double c[4][4]; # set to zero

/* Multiply 4 x 4 matrices a and b */
void mmm(double *a, double *b, double *c, int n) {
    int i, j, k;
    for (i = 0; i < 4; i++)
        for (j = 0; j < 4; j++)
            for (k = 0; k < 4; k++)
                c[i*4+j] += a[i*4 + k]*b[k*4 + j];
}
```

- **Compiled without flags:**
~1300 cycles
- **Compiled with `-O3 -m64 -march=... -fno-tree-vectorize`**
~150 cycles
- **Core 2 Duo, 2.66 GHz**

Optimizing Compilers

- Compilers are **good** at: mapping program to machine instructions
 - register allocation
 - code selection and ordering (scheduling)
 - dead code elimination
 - eliminating minor inefficiencies
- Compilers are **not good** at: improving asymptotic efficiency
 - up to programmer to select best overall algorithm
 - big-O savings are (often) more important than constant factors
 - but constant factors also matter
- Compilers are **not good** at: overcoming “optimization blockers”
 - potential memory aliasing
 - potential procedure side-effects

Limitations of Optimizing Compilers

- *If in doubt, the compiler is conservative*
- **Operate under fundamental constraints**
 - Must not change program behavior under any possible condition
 - Often prevents it from making optimizations when would only affect behavior under pathological conditions.
- **Behavior that may be obvious to the programmer can be obfuscated by languages and coding styles**
 - e.g., data ranges may be more limited than variable types suggest
- **Most analysis is performed only within procedures**
 - Whole-program analysis is too expensive in most cases
- **Most analysis is based only on *static* information**
 - Compiler has difficulty anticipating run-time inputs

Today

- Overview
- **Program optimization**
 - **Code motion/precomputation**
 - Strength reduction
 - Sharing of common subexpressions
 - Optimization blocker: Procedure calls
 - Optimization blocker: Memory aliasing
- Exploiting Instruction-Level Parallelism

Generally Useful Optimizations

- Optimizations that you or the compiler should do regardless of processor / compiler
- Code Motion
 - Reduce frequency with which computation performed
 - If it will always produce same result
 - Especially moving code out of loop

```
void set_row(double *a, double *b,  
            long i, long n)  
{  
    long j;  
    for (j = 0; j < n; j++)  
        a[n*i+j] = b[j];  
}
```

```
long j;  
int ni = n*i;  
for (j = 0; j < n; j++)  
    a[ni+j] = b[j];
```

Compiler-Generated Code Motion

```

void set_row(double *a, double *b,
            long i, long n)
{
    long j;
    for (j = 0; j < n; j++)
        a[n*i+j] = b[j];
}
  
```

```

long j;
long ni = n*i;
double *rowp = a+ni;
for (j = 0; j < n; j++)
    *rowp++ = b[j];
  
```

Where are the FP operations?

```

set_row:
    testq    %rcx, %rcx                # Test n
    jle     .L4                        # If 0, goto done
    movq    %rcx, %rax                # rax = n
    imulq   %rdx, %rax                # rax *= i
    leaq    (%rdi,%rax,8), %rdx        # rowp = A + n*i*8
    movl    $0, %r8d                  # j = 0
.L3:                                     # loop:
    movq    (%rsi,%r8,8), %rax         # t = b[j]
    movq    %rax, (%rdx)               # *rowp = t
    addq    $1, %r8                    # j++
    addq    $8, %rdx                   # rowp++
    cmpq    %r8, %rcx                 # Compare n:j
    jg     .L3                         # If >, goto loop
.L4:                                     # done:
    rep ; ret
  
```

Today

- Overview
- **Program optimization**
 - Code motion/precomputation
 - **Strength reduction**
 - Sharing of common subexpressions
 - Optimization blocker: Procedure calls
 - Optimization blocker: Memory aliasing
- Exploiting Instruction-Level Parallelism

Reduction in Strength

- Replace costly operation with simpler one
- Shift, add instead of multiply or divide
 - $16 * x \quad \rightarrow \quad x \ll 4$
 - Utility machine dependent
 - Depends on cost of multiply or divide instruction
 - On Intel Nehalem, integer multiply requires 3 CPU cycles
- Recognize sequence of products

```

for (i = 0; i < n; i++)
  for (j = 0; j < n; j++)
    a[n*i + j] = b[j];
  
```



```

int ni = 0;
for (i = 0; i < n; i++) {
  for (j = 0; j < n; j++)
    a[ni + j] = b[j];
  ni += n;
}
  
```

Today

- Overview
- **Program optimization**
 - Code motion/precomputation
 - Strength reduction
 - **Sharing of common subexpressions**
 - Optimization blocker: Procedure calls
 - Optimization blocker: Memory aliasing
- Exploiting Instruction-Level Parallelism

Share Common Subexpressions

- Reuse portions of expressions
- Compilers often not very sophisticated in exploiting arithmetic properties

```

/* Sum neighbors of i,j */
up =    val[(i-1)*n + j  ];
down =  val[(i+1)*n + j  ];
left =  val[i*n          + j-1];
right = val[i*n          + j+1];
sum = up + down + left + right;

```

3 multiplications: $i*n$, $(i-1)*n$, $(i+1)*n$

```

leaq    1(%rsi), %rax    # i+1
leaq   -1(%rsi), %r8    # i-1
imulq   %rcx, %rsi     # i*n
imulq   %rcx, %rax     # (i+1)*n
imulq   %rcx, %r8     # (i-1)*n
addq    %rdx, %rsi     # i*n+j
addq    %rdx, %rax     # (i+1)*n+j
addq    %rdx, %r8     # (i-1)*n+j

```

```

long inj = i*n + j;
up =    val[inj - n];
down =  val[inj + n];
left =  val[inj - 1];
right = val[inj + 1];
sum = up + down + left + right;

```

1 multiplication: $i*n$

```

imulq   %rcx, %rsi    # i*n
addq    %rdx, %rsi    # i*n+j
movq    %rsi, %rax    # i*n+j
subq    %rcx, %rax    # i*n+j-n
leaq    (%rsi,%rcx), %rcx # i*n+j+n

```

Today

- Overview
- **Program optimization**
 - Code motion/precomputation
 - Strength reduction
 - Sharing of common subexpressions
 - **Optimization blocker: Procedure calls**
 - Optimization blocker: Memory aliasing
- Exploiting Instruction-Level Parallelism


Optimization Blocker #1: Procedure Calls

■ Procedure to Convert String to Lower Case

```
void lower(char *s)
{
    int i;
    for (i = 0; i < strlen(s); i++)
        if (s[i] >= 'A' && s[i] <= 'Z')
            s[i] -= ('A' - 'a');
}
```

Lower Case Conversion Performance

- Time quadruples when double string length
- Quadratic performance



Convert Loop To Goto Form

```
void lower(char *s)
{
    int i = 0;
    if (i >= strlen(s))
        goto done;
loop:
    if (s[i] >= 'A' && s[i] <= 'Z')
        s[i] -= ('A' - 'a');
    i++;
    if (i < strlen(s))
        goto loop;
done:
}
```

- `strlen` executed every iteration

Calling Strlen

```
/* My version of strlen */
size_t strlen(const char *s)
{
    size_t length = 0;
    while (*s != '\0') {
        s++;
        length++;
    }
    return length;
}
```

■ Strlen performance

- Only way to determine length of string is to scan its entire length, looking for null character.

■ Overall performance, string of length N

- N calls to strlen
- Require times N, N-1, N-2, ..., 1
- Overall $O(N^2)$ performance


Improving Performance

```
void lower(char *s)
{
    int i;
    int len = strlen(s);
    for (i = 0; i < len; i++)
        if (s[i] >= 'A' && s[i] <= 'Z')
            s[i] -= ('A' - 'a');
}
```

- Move call to `strlen` outside of loop
- Since result does not change from one iteration to another
- Form of code motion

Lower Case Conversion Performance

- Time doubles when double string length
- Linear performance of lower2



Optimization Blocker: Procedure Calls

■ *Why couldn't compiler move strlen out of inner loop?*

- Procedure may have side effects
 - Alters global state each time called
- Function may not return same value for given arguments
 - Depends on other parts of global state
 - Procedure `lower` could interact with `strlen`

■ **Warning:**

- Compiler treats procedure call as a black box
- Weak optimizations near them

■ **Remedies:**

- Use of `inline` functions
 - GCC does this with `-O2`
 - See web aside ASM:OPT
- Do your own code motion

```
int lencnt = 0;
size_t strlen(const char *s)
{
    size_t length = 0;
    while (*s != '\0') {
        s++; length++;
    }
    lencnt += length;
    return length;
}
```

Today

- Overview
- **Program optimization**
 - Code motion/precomputation
 - Strength reduction
 - Sharing of common subexpressions
 - Optimization blocker: Procedure calls
 - **Optimization blocker: Memory aliasing**
- Exploiting Instruction-Level Parallelism

Memory Matters

```

/* Sum rows is of n X n matrix a
   and store in vector b */
void sum_rows1(double *a, double *b, long n) {
    long i, j;
    for (i = 0; i < n; i++) {
        b[i] = 0;
        for (j = 0; j < n; j++)
            b[i] += a[i*n + j];
    }
}

```

```

# sum_rows1 inner loop
.L53:
    addsd    (%rcx), %xmm0           # FP add
    addq    $8, %rcx
    decq    %rax
    movsd   %xmm0, (%rsi,%r8,8)     # FP store
    jne     .L53

```

- Code updates `b[i]` on every iteration
- Why couldn't compiler optimize this away?

Memory Aliasing (Simple Example)

```
void twiddle1(int *xp, int *yp) {  
    *xp += *yp;  
    *xp += *yp;  
}
```

```
void twiddle2(int *xp, int *yp) {  
    *xp += 2 * *yp;  
}
```

$i=2, j=2;$

$xp = \&i; yp = \&j;$

$*xp += *yp; \quad // *xp = 2+2 = 4$

$*xp += *yp \quad // *xp = 4+2 = 6$

$*xp += 2 * (*yp); \quad // *xp = 2 + 2*2 = 6$

What if xp and yp point to the same address?

$int i=2;$

$xp = yp = \&i;$

twiddle1:

$*xp += *yp; \quad // *xp = 2 + 2 = 4$

$*xp += *yp; \quad // *xp = 4 + 4 = 8$

twiddle2:

$*xp += 2 * (*yp); \quad // *xp = 2 + 2*2 = 6$

Memory Aliasing

```

/* Sum rows is of n X n matrix a
   and store in vector b */
void sum_rows1(double *a, double *b, long n) {
    long i, j;
    for (i = 0; i < n; i++) {
        b[i] = 0;
        for (j = 0; j < n; j++)
            b[i] += a[i*n + j];
    }
}

```

```

double A[9] =
    { 0, 1, 2,
      4, 8, 16},
    { 32, 64, 128};

double B[3] = A+3;

sum_rows1(A, B, 3);

```

Value of B:

init: [4, 8, 16]

i = 0: [3, 8, 16]

i = 1: [3, 22, 16]

i = 2: [3, 22, 224]

- Code updates `b[i]` on every iteration
- Must consider possibility that these updates will affect program behavior

Removing Aliasing

```

/* Sum rows is of n X n matrix a
   and store in vector b */
void sum_rows2(double *a, double *b, long n) {
    long i, j;
    for (i = 0; i < n; i++) {
        double val = 0;
        for (j = 0; j < n; j++)
            val += a[i*n + j];
        b[i] = val;
    }
}

```

```

# sum_rows2 inner loop
.L66:
    addsd    (%rcx), %xmm0    # FP Add
    addq    $8, %rcx
    decq    %rax
    jne     .L66

```

- No need to store intermediate results

Optimization Blocker: Memory Aliasing

- **Memory aliasing: Two different memory references write to the same location**
- **Easy to have happen in C**
 - Since allowed to do address arithmetic
 - Direct access to storage structures
- **Hard to analyze = compiler cannot figure it out**
 - Hence is conservative
- **Solution: Scalar replacement in innermost loop**
 - Copy memory variables **that are reused** into local variables
 - Basic scheme:
 - **Load:** $t1 = a[i], t2 = b[i+1], \dots$
 - **Compute:** $t4 = t1 * t2; \dots$
 - **Store:** $a[i] = t12, b[i+1] = t7, \dots$

Today

- Overview
- Program optimization
 - Code motion/precomputation
 - Strength reduction
 - Sharing of common subexpressions
 - Optimization blocker: Procedure calls
 - Optimization blocker: Memory aliasing
- **Exploiting Instruction-Level Parallelism**


Exploiting Instruction-Level Parallelism

- **Need general understanding of modern processor design**
 - Hardware can execute multiple instructions in parallel
- **Performance limited by data dependencies**
- **Simple transformations can have dramatic performance improvement**
 - Compilers often cannot make these transformations
 - Lack of associativity and distributivity in floating-point arithmetic

Benchmark Example: Data Type for Vectors

```

/* data structure for vectors */
typedef struct{
    int len;
    double *data;
} vec;
  
```



```

/* retrieve vector element and store at val */
double get_vec_element(*vec, idx, double *val)
{
    if (idx < 0 || idx >= v->len)
        return 0;
    *val = v->data[idx];
    return 1;
}
  
```


Benchmark Computation

```

void combine1(vec_ptr v, data_t *dest)
{
    long int i;
    *dest = IDENT;
    for (i = 0; i < vec_length(v); i++) {
        data_t val;
        get_vec_element(v, i, &val);
        *dest = *dest OP val;
    }
}
  
```

Compute sum or product of vector elements

■ Data Types


- Use different declarations for `data_t`
- `int`
- `float`
- `double`

■ Operations

- Use different definitions of `OP` and `IDENT`
- `+` / `0`
- `*` / `1`

Cycles Per Element (CPE)

- Convenient way to express performance of program that operates on vectors or lists
- Length = n
- In our case: **CPE = cycles per OP**
- $T = CPE * n + \text{Overhead}$
 - CPE is slope of line



Benchmark Performance

```

void combine1(vec_ptr v, data_t *dest)
{
    long int i;
    *dest = IDENT;
    for (i = 0; i < vec_length(v); i++) {
        data_t val;
        get_vec_element(v, i, &val);
        *dest = *dest OP val;
    }
}

```

Compute sum or product of vector elements

Method	Integer		Double FP	
	Add	Mult	Add	Mult
Combine1 unoptimized	29.0	29.2	27.4	27.9
Combine1 -O1	12.0	12.0	12.0	13.0

Basic Optimizations

```
void combine4(vec_ptr v, data_t *dest)
{
    int i;
    int length = vec_length(v);
    data_t *d = get_vec_start(v);
    data_t t = IDENT;
    for (i = 0; i < length; i++)
        t = t OP d[i];
    *dest = t;
}
```

- Move `vec_length` out of loop
- Avoid bounds check on each cycle
- Accumulate in temporary

Effect of Basic Optimizations

```

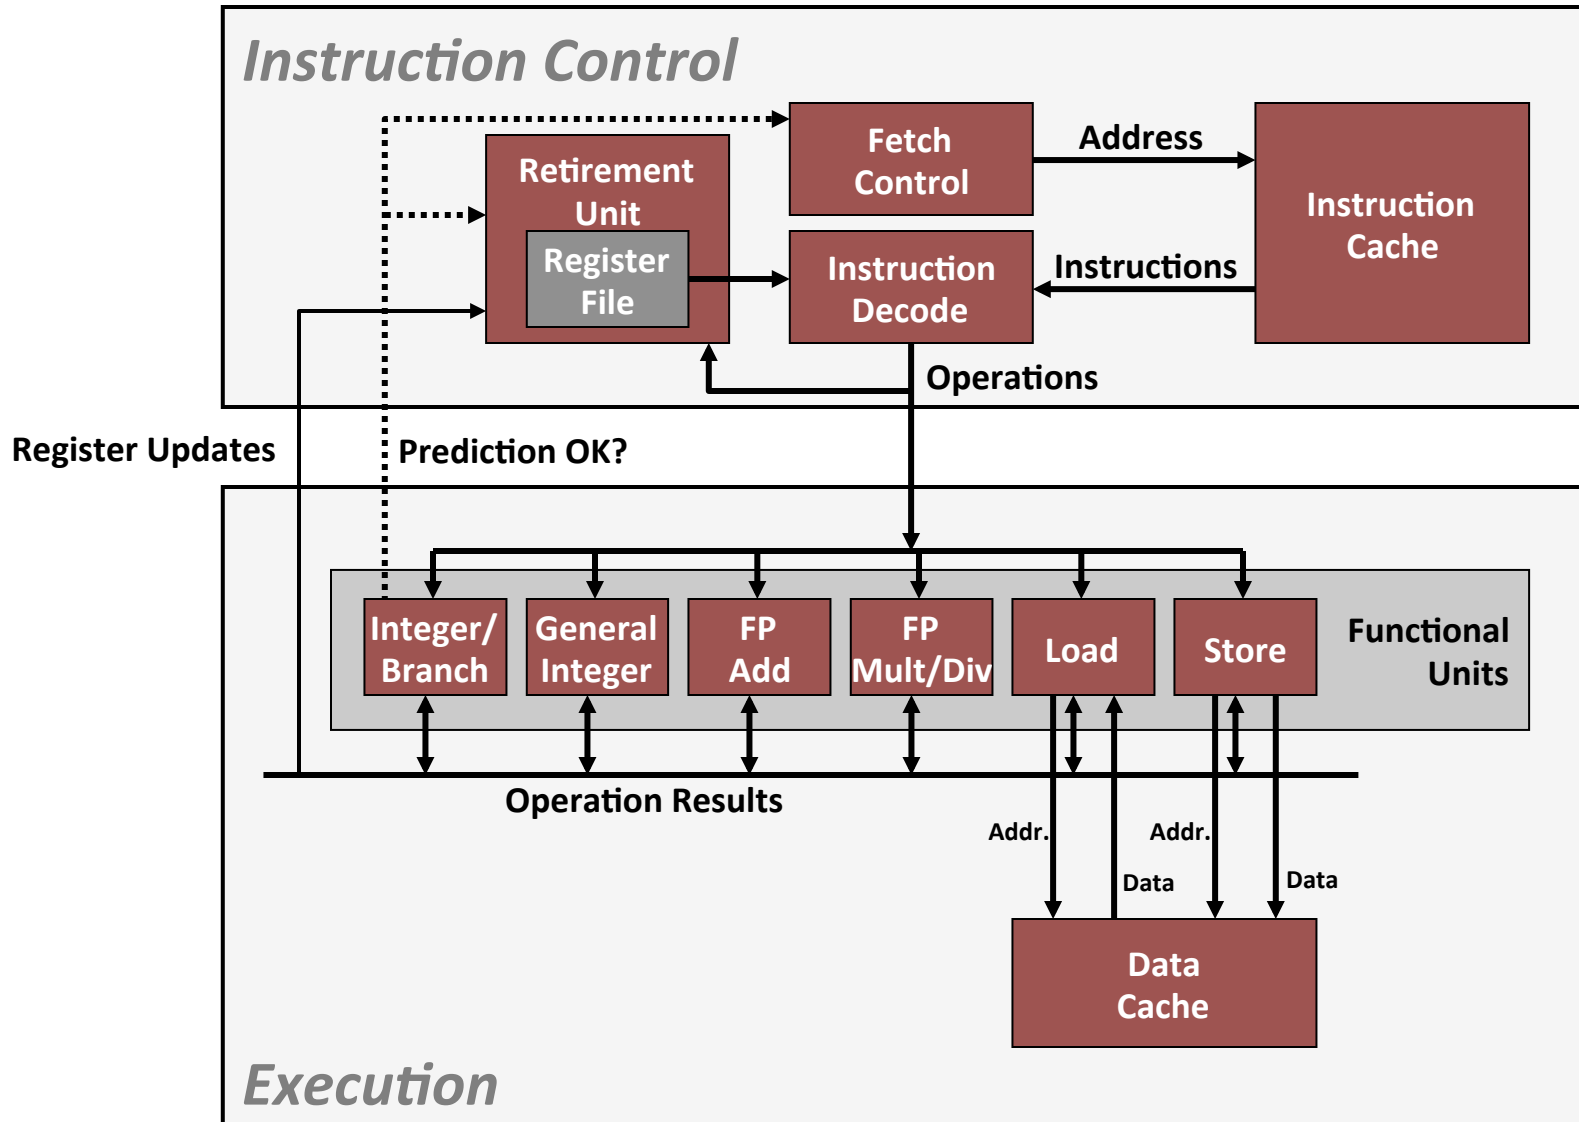
void combine4(vec_ptr v, data_t *dest)
{
    int i;
    int length = vec_length(v);
    data_t *d = get_vec_start(v);
    data_t t = IDENT;
    for (i = 0; i < length; i++)
        t = t OP d[i];
    *dest = t;
}

```

Method	Integer		Double FP	
	Add	Mult	Add	Mult
Combine1 -O1	12.0	12.0	12.0	13.0
Combine4	2.0	3.0	3.0	5.0


- Eliminates sources of overhead in loop

Modern CPU Design



Latency versus Throughput

■ Example:	<i>latency</i>	<i>cycles/issue</i>
Integer Multiply	10	1



■ Consequence:

- How fast can 10 independent int mults be executed?

$t1 = t2 * t3; t4 = t5 * t6; \dots$

- How fast can 10 sequentially dependent int mults be executed?

$t1 = t2 * t3; t4 = t5 * t1; t6 = t7 * t4; \dots$

- Major problem for fast execution: **Keep pipelines filled**

Superscalar Processor

- **Definition:** A superscalar processor can issue and execute *multiple instructions in one cycle*. The instructions are retrieved from a sequential instruction stream and are usually scheduled dynamically.
- **Benefit:** without programming effort, superscalar processor can take advantage of the *instruction level parallelism* that most programs have
- Most CPUs since about 1998 are superscalar.
- Intel: since Pentium Pro

Nehalem CPU

■ Multiple instructions can execute in parallel

- 1 load, with address computation
- 1 store, with address computation
- 2 simple integer (one may be branch)
- 1 complex integer (multiply/divide)
- 1 FP Multiply
- 1 FP Add

■ Some instructions take > 1 cycle, but can be pipelined

<i>Instruction</i>	<i>Latency</i>	<i>Cycles/Issue</i>
Load / Store	4	1
Integer Multiply	3	1
Integer/Long Divide	11--21	11--21
Single/Double FP Multiply	4/5	1
Single/Double FP Add	3	1
Single/Double FP Divide	10--23	10--23

x86-64 Compilation of Combine4


■ Inner Loop (Case: Integer Multiply)

```

.L519:                                # Loop:
    imull  (%rax,%rdx,4), %ecx        # t = t * d[i]
    addq   $1, %rdx                  # i++
    cmpq   %rdx, %rbp                # Compare length:i
    jg     .L519                      # If >, goto Loop
    
```

Method	Integer		Double FP	
	Add	Mult	Add	Mult
Combine4	2.0	3.0	3.0	5.0
Latency Bound	1.0	3.0	3.0	5.0

Combine4 = Serial Computation (OP = *)



- **Computation (length=8)**

`(((((1 * d[0]) * d[1]) * d[2]) * d[3]) * d[4]) * d[5]) * d[6]) * d[7])`

- **Sequential dependence**

- Performance: determined by latency of OP

Loop Unrolling

```
void unroll2a_combine(vec_ptr v, data_t *dest)
{
    int length = vec_length(v);
    int limit = length-1;
    data_t *d = get_vec_start(v);
    data_t x = IDENT;
    int i;
    /* Combine 2 elements at a time */
    for (i = 0; i < limit; i+=2) {
        x = (x OP d[i]) OP d[i+1];
    }
    /* Finish any remaining elements */
    for (; i < length; i++) {
        x = x OP d[i];
    }
    *dest = x;
}
```

- Perform 2x more useful work per iteration

Effect of Loop Unrolling

Method	Integer		Double FP	
Operation	Add	Mult	Add	Mult
Combine4	2.0	3.0	3.0	5.0
Unroll 2x	2.0	1.5	3.0	5.0
Latency Bound	1.0	3.0	3.0	5.0

- **Helps integer multiply**
 - below latency bound
 - Compiler does clever optimization
- **Others don't improve. *Why?***
 - Still sequential dependency

```
x = (x OP d[i]) OP d[i+1];
```

Loop Unrolling with Reassociation

```

void unroll2aa_combine(vec_ptr v, data_t *dest)
{
    int length = vec_length(v);
    int limit = length-1;
    data_t *d = get_vec_start(v);
    data_t x = IDENT;
    int i;
    /* Combine 2 elements at a time */
    for (i = 0; i < limit; i+=2) {
        x = x OP (d[i] OP d[i+1]);
    }
    /* Finish any remaining elements */
    for (; i < length; i++) {
        x = x OP d[i];
    }
    *dest = x;
}

```

Compare to before

$x = (x \text{ OP } d[i]) \text{ OP } d[i+1];$

- Can this change the result of the computation?
- Yes, for FP. *Why?*

Effect of Reassociation

Method	Integer		Double FP	
Operation	Add	Mult	Add	Mult
Combine4	2.0	3.0	3.0	5.0
Unroll 2x	2.0	1.5	3.0	5.0
Unroll 2x, reassociate	2.0	1.5	1.5	3.0
Latency Bound	1.0	3.0	3.0	5.0
Throughput Bound	1.0	1.0	1.0	1.0

■ Nearly 2x speedup for Int *, FP +, FP *


- Reason: Breaks sequential dependency

```
x = x OP (d[i] OP d[i+1]);
```

- Why is that? (next slide)

Reassociated Computation

```
x = x OP (d[i] OP d[i+1]);
```



■ What changed:

- Ops in the next iteration can be started early (no dependency)

■ Overall Performance

- N elements, D cycles latency/op
- Should be $(N/2+1)*D$ cycles:
CPE = D/2
- Measured CPE slightly worse for FP mult

Loop Unrolling with Separate Accumulators

```
void unroll2a_combine(vec_ptr v, data_t *dest)
{
    int length = vec_length(v);
    int limit = length-1;
    data_t *d = get_vec_start(v);
    data_t x0 = IDENT;
    data_t x1 = IDENT;
    int i;
    /* Combine 2 elements at a time */
    for (i = 0; i < limit; i+=2) {
        x0 = x0 OP d[i];
        x1 = x1 OP d[i+1];
    }
    /* Finish any remaining elements */
    for (; i < length; i++) {
        x0 = x0 OP d[i];
    }
    *dest = x0 OP x1;
}
```

- Different form of reassociation

Effect of Separate Accumulators

Method	Integer		Double FP	
	Add	Mult	Add	Mult
Combine4	2.0	3.0	3.0	5.0
Unroll 2x	2.0	1.5	3.0	5.0
Unroll 2x, reassociate	2.0	1.5	1.5	3.0
Unroll 2x Parallel 2x	1.5	1.5	1.5	2.5
Latency Bound	1.0	3.0	3.0	5.0
Throughput Bound	1.0	1.0	1.0	1.0

- **2x speedup (over Combine4) for Int *, FP +, FP ***
 - Breaks sequential dependency in a “cleaner,” more obvious way

```


x0 = x0 OP d[i];
x1 = x1 OP d[i+1];

```

Separate Accumulators

```

x0 = x0 OP d[i];
x1 = x1 OP d[i+1];
    
```



■ What changed:

- Two independent “streams” of operations

■ Overall Performance

- N elements, D cycles latency/op
- Should be $(N/2+1)*D$ cycles:
CPE = D/2
- CPE matches prediction!

What Now?

Unrolling & Accumulating

■ Idea

- Can unroll to any degree L
- Can accumulate K results in parallel
- L must be multiple of K

■ Limitations

- Diminishing returns
 - Cannot go beyond throughput limitations of execution units
- Large overhead for short lengths
 - Finish off iterations sequentially

Unrolling & Accumulating: Double *

■ Case

- Intel Nehalem
- Double FP Multiplication
- Latency bound: 5.00. Throughput bound: 1.00

	FP *	Unrolling Factor L							
	K	1	2	3	4	6	8	10	12
<i>Accumulators</i>	1	5.00	5.00	5.00	5.00	5.00	5.00		
	2		2.50		2.50		2.50		
	3			1.67					
	4				1.25		1.25		
	6					1.00			1.19
	8						1.02		
	10							1.01	
	12								1.00

Unrolling & Accumulating: Int +

■ Case

- Intel Nehalem
- Integer addition
- Latency bound: 1.00. Throughput bound: 1.00

	FP *	Unrolling Factor L							
	K	1	2	3	4	6	8	10	12
<i>Accumulators</i>	1	2.00	2.00	1.00	1.01	1.02	1.03		
	2		1.50		1.26		1.03		
	3			1.00					
	4				1.00		1.24		
	6					1.00			1.02
	8						1.03		
	10							1.01	
	12								1.09

Achievable Performance

Method	Integer		Double FP	
Operation	Add	Mult	Add	Mult
Scalar Optimum	1.00	1.00	1.00	1.00
Latency Bound	1.00	3.00	3.00	5.00
Throughput Bound	1.00	1.00	1.00	1.00

- Limited only by throughput of functional units
- Up to 29X improvement over original, unoptimized code

Using Vector Instructions

Method	Integer		Double FP	
Operation	Add	Mult	Add	Mult
Scalar Optimum	1.00	1.00	1.00	1.00
Vector Optimum	0.25	0.53	0.53	0.57
Latency Bound	1.00	3.00	3.00	5.00
Throughput Bound	1.00	1.00	1.00	1.00
Vec Throughput Bound	0.25	0.50	0.50	0.50

■ Make use of SSE Instructions

- Parallel operations on multiple data elements
- See Web Aside OPT:SIMD on CS:APP web page

Reading Assignment

- **Section 5.14 – Identifying and Eliminating Performance Bottlenecks**
 - Program Profiling
 - Using a Profiler to Guide Optimization
 - Amdahl's Law

Getting High Performance

- **Good compiler and flags**
- **Don't do anything stupid**
 - Watch out for hidden algorithmic inefficiencies
 - Write compiler-friendly code
 - Watch out for optimization blockers:
procedure calls & memory references
 - Look carefully at innermost loops (where most work is done)
- **Tune code for machine**
 - Exploit instruction-level parallelism
 - Avoid unpredictable branches (not covered)
 - Make code cache friendly (covered in course before)