

Simple
measures,
large impact

Vipin Vasu

Elimination of
Common
Subexpressions

Avoiding
Branches

Using SIMD
instruction
sets

Conclusion

Simple measures, large impact

Vipin Vasu

Simple
measures,
large impact

Vipin Vasu

Elimination of
Common
Subexpressions

Avoiding
Branches

Using SIMD
instruction
sets

Conclusion

Outline

- 1 Elimination of Common Subexpressions
- 2 Avoiding Branches
- 3 Using SIMD instruction sets
- 4 Conclusion

Elimination of Common Subexpressions

- Common subexpression elimination is an optimization that is often considered a task for compilers
- Basically one tries to save time by precalculating parts of complex expressions and assigning them to temporary variables before a code construct starts that uses those parts multiple times.

Elimination of Common Subexpressions

```
1 ! inefficient
2 do i=1,N
3   A(i)=A(i)+s+r*sin(x)
4 enddo
```

→

```
tmp=s+r*sin(x)
do i=1,N
  A(i)=A(i)+tmp
enddo
```

This optimization is also called loop invariant code motion

Avoiding Branches

- “Tight” loops, i.e., loops that have few operations in them, are typical candidates for software pipelining , loop unrolling, and other optimization techniques .
- Compiler optimization fails or is inefficient, performance will suffer if the loop body contains conditional branches.

Simple
measures,
large impact

Vipin Vasu

Elimination of
Common
Subexpressions

Avoiding
Branches

Using SIMD
instruction
sets

Conclusion

Avoiding Branches

```
1 do j=1,N
2   do i=1,N
3     if(i.ge.j) then
4       sign=1.d0
5     else if(i.lt.j) then
6       sign=-1.d0
7     else
8       sign=0.d0
9     endif
10    C(j) = C(j) + sign * A(i,j) * B(i)
11  enddo
12 enddo
```

```
1 do j=1,N
2   do i=j+1,N
3     C(j) = C(j) + A(i,j) * B(i)
4   enddo
5 enddo
6 do j=1,N
7   do i=1,j-1
8     C(j) = C(j) - A(i,j) * B(i)
9   enddo
10 enddo
```

Using SIMD instruction sets

- The use of SIMD in microprocessors is often termed “vectorization,”
- Generally speaking, a “vectorizable” loop in this context will run faster if more operations can be performed with a single instruction.
- Preferring SIMD instructions over scalar ones is no guarantee for a performance improvement.
- If the code is strongly limited by memory bandwidth, no SIMD technique can bridge this gap.

Simple
measures,
large impact

Vipin Vasu

Elimination of
Common
Subexpressions

Avoiding
Branches

Using SIMD
instruction
sets

Conclusion

Using SIMD instruction sets

```
1 real, dimension(1:N) :: r, x, y
2 do i=1, N
3   r(i) = x(i) + y(i)
4 enddo
```

Using SIMD instruction sets

```
1  ! vectorized part
2  rest = mod(N,4)
3  do i=1,N-rest,4
4      load R1 = [x(i),x(i+1),x(i+2),x(i+3)]
5      load R2 = [y(i),y(i+1),y(i+2),y(i+3)]
6      ! "packed" addition (4 SP flops)
7      R3 = ADD(R1,R2)
8      store [r(i),r(i+1),r(i+2),r(i+3)] = R3
9  enddo
10 ! remainder loop
11 do i=N-rest+1,N
12     r(i) = x(i) + y(i)
13 enddo
```

Using SIMD instruction sets

- Some SIMD instruction sets distinguish between aligned and unaligned data.
- In cases where the compiler knows nothing about the alignment of arrays used in a vectorized loop and cannot otherwise influence it, unaligned loads and stores must be used, incurring some performance penalty.
- A loop with a true dependency cannot be SIMD vectorized.

```
1 do i=2,N
2   A(i)=s*A(i-1)
3 enddo
```

Using SIMD instruction Sets

- There are no fixed guidelines for when a loop qualifies as vectorized.
- Load and store instructions could still be scalar; compilers tend to report such loops as “vectorized” .
- On x86 processors with SSE(Streaming SIMD Extensions) support, the lower and higher 64 bits of a register can be moved independently.

Using SIMD instruction Sets

```
1 rest = mod(N,2)
2 do i=1,N-rest,2
3     ! scalar loads
4     load R1.low = x(i)
5     load R1.high = x(i+1)
6     load R2.low = y(i)
7     load R2.high = y(i+1)
8     ! "packed" addition (2 DP flops)
9     R3 = ADD(R1,R2)
10    ! scalar stores
11    store r(i) = R3.low
12    store r(i+1) = R3.high
13 enddo
14 ! remainder "loop"
15 if(rest.eq.1) r(N) = x(N) + y(N)
```

Using SIMD instruction Sets

- If the compiler cannot be convinced to properly vectorize a loop even with additional command line options or source code directives, before using assembly language altogether is to employ compiler intrinsics.
- Intrinsics are constructs that resemble assembly instructions so closely that they can usually be translated 1:1 by the compiler.

Simple
measures,
large impact

Vipin Vasu

Elimination of
Common
Subexpressions

Avoiding
Branches

Using SIMD
instruction
sets

Conclusion

The End