



C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Compiler Optimization

Mirko Köster

Seminar
Effiziente Programmierung in C
Fachbereich Informatik
Universität Hamburg

2012-11-29



Overview

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Key Aspects

- What is the compiler capable of?
- What are its weaknesses?
- How can you make use of it?

Content

- Automatic Optimization
- Profile Guided Optimization
- Aiding Optimizations
- 'Safe' / 'Unsafe' Optimizations
- OpenMP



Preface

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Compiler

- Some examples are from the GNU C compiler
- There are lots of other good compilers available
- But I'll just give you an overview of the concepts
- Refer to the manual of your compiler for specific optimizations

Architecture

- In this presentation I'll focus on the x86 architecture
- If you are developing for another architecture get familiar with it (but the basic concepts will work there as well)



Automatic Optimization

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

- Changes that don't affect the result
- May optimize
 - Execution speed
 - File size of the executable
 - or even power consumption
- activated by compiler options / flags

How does it work?

- 1 Analyse source code
- 2 Assume stricter rules as the c-language
- 3 Prove assumptions
- 4 Apply optimization(s)



Automatic Optimization

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

How to use it

- activated by `-O[level]`
- or manually by the specific flag

-O1

- `-fauto-inc-dec`
- `-fcompare-elim`
- `-fcprop-registers`
- `-fdce`
- `-fdefer-pop`
- `-fdelayed-branch`
- `-fdse`
- `-fguess-branch-probability`
- `-fif-conversion2`
- `-fif-conversion`

-O1

- `-fipa-pure-const`
- `-fipa-profile`
- `-fipa-reference`
- `-fmerge-constants`
- `-fsplit-wide-types`
- `-ftree-bit-ccp`
- `-ftree-builtin-call-dce`
- `-ftree-ccp`
- `-ftree-ch`
- `-ftree-copyrename`
- `-ftree-dce`

-O1

- `-ftree-dominator-opts`
- `-ftree-dse`
- `-ftree-forwprop`
- `-ftree-fre`
- `-ftree-phirop`
- `-ftree-slsr`
- `-ftree-sra`
- `-ftree-pta`
- `-ftree-ter`
- `-funit-at-a-time`



Automatic Optimization

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

-o2 (includes all from -O1)

- -fthread-jumps
- -falign-functions -falign-jumps
- -falign-loops -falign-labels
- -fcaller-saves
- -fcrossjumping
- -fcse-follow-jumps -fcse-skip-blocks
- -fdelete-null-pointer-checks
- -fdevirtualize
- -fexpensive-optimizations
- -fgcse -fgcse-lm
- -fhoist-adjacent-loads
- -finline-small-functions
- -findirect-inlining

-o2 (includes all from -O1)

- -fipa-sra
- -foptimize-sibling-calls
- -fpartial-inlining
- -fpeephole2
- -fregmove
- -freorder-blocks -freorder-functions
- -frerun-cse-after-loop
- -fsched-interblock -fsched-spec
- -fschedule-insns -fschedule-insns2
- -fstrict-aliasing -fstrict-overflow
- -ftree-switch-conversion -ftree-tail-merge
- -ftree-pre
- -ftree-vrp



Automatic Optimization

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

-o3 (includes all from -O2)

- -finline-functions
- -funswitch-loops
- -fpredictive-commoning
- -fgcse-after-reload
- -ftree-vectorize
- -fvect-cost-model
- -ftree-partial-pre
- -fipa-cp-clone

-O0 (default)

Reduce compilation time and make debugging
produce the expected results

-Os (Optimize for size)

disables

- -falign-functions
- -falign-jumps
- -falign-loops
- -falign-labels
- -freorder-blocks
- -freorder-blocks-and-partition
- -fprefetch-loop-arrays
- -ftree-vect-loop-version



Automatic Optimization

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Some optimizations are very time-consuming

- Some problems are np hard
- Some problems are even undecidable
- Tradeoff: in those cases the compiler won't give the optimal result but a good result (to save time/space during compilation)



Architecture Independent Optimizations

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

- Do not rely upon knowledge of the underlying architecture
- Can be applied under any circumstances after the assumptions have been proven



Loop Invariant Code Motion

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

Moves code out of a loop if it is invariant of the loop variable

unoptimized

```
1 int sum=0, x;  
  for(int i = 0; i < n; i++) {  
3     sum += i;  
     x = 5;  
5  }
```

optimized

```
1 int sum=0, x = 5;  
  for(int i = 0; i < n; i++) {  
3     sum += i;  
  }
```



Const Propagation (with Loop Optimization)

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

Evaluation of expressions with known values at compile time

unoptimized

```
1 int N = 10, sum = 0;
2 for(int i = 0; i < N; i++)
3     sum += i;
4 printf("sum = %d\n", sum);
```

optimized

```
1 printf("sum = %d\n", 45);
```



Dead Code Elimination

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

Removes code that is unnecessary or never executed

unoptimized

```
1 unsigned int x = foobar();  
   if(x < 0) {  
3     printf("never executed\n");  
   } else {  
5     printf("x: %u\n", x);  
   }
```

optimized

```
printf("x: %u\n", foobar());
```



Common Subexpression Elimination

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

Reduces occurrences of multiple common subexpressions

unoptimized

```
1 void foo(int *a, int n) {  
    for(int i = 0; i < n; i++)  
3     a[i] += a[i]/n + a[i]*n;  
}
```

optimized

```
void foo(int *a, int n) {  
2     int temp;  
    for(int i = 0; i < n; i++)  
4         temp = a[i]  
        a[i] += temp/n + temp*n;  
6 }
```



Interprocedural Optimization

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

looks at multiple functions and how they work together

Arguments in Registers

- passing arguments in registers instead of pushing/popping them to/from stack
- reduces call/return overhead
- requires modification of caller and callee



Definition

- For small functions the overhead of calling may be larger in relation to the body.
- Inlining replaces the call to the function with the body.

unoptimized

```
1 int foo(int a) {  
2     return a * (a+1);  
3 }  
4 ...  
5 int a[5];  
6 for(int i = 0; i < 5; i++)  
    a[i] = foo(i);
```

optimized

```
1 int a[5];  
2  
3 a[0] = 0 * 1;  
4 a[1] = 1 * 2;  
5 a[2] = 2 * 3;  
6 a[3] = 3 * 4;  
7 a[4] = 4 * 5;
```



Interprocedural Constant Propagation

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

Evaluation of expressions with known values at compile time taking multiple functions into account

unoptimized

```
1 static int square(int x) {  
    return x*x;  
3 }  
  
5 printf("5^2=%d\n",square(5));
```

optimized

```
1 static int square(int x) {  
    return x*x;  
3 }  
  
5 printf("5^2 = %d\n", 25);
```




Architecture Dependent Optimizations

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

- Target-specific optimizations
- The compiler has to know the target architecture
- caution: the executable may not run on older machines

What makes a target architecture?

- Instruction set (e.g. x86)
- Number of (special purpose) registers
- Cache size & type
- possibly some instruction set extensions (MMX, SSE...)



Overview History

- 1985 x86 32bit
- 1989 x87 FPU (Co-Processor)
- 1993 MMX
- 1997 SSE, 3DNow!
- 2000 SSE2
- 2003 x86-64 64bit
- 2004 SSE3
- 2007 SSE4a
- 2011 SSE5/AVX
- 2013 AVX2, FMA3



-mtune & -march

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

-mtune

This option optimizes for the given architecture, making the code faster on those machines. But it will still run on other architectures.

-march

This option will make the most of the given architecture. May not run on other architectures.

example options

- i386
- pentium
- corei7
- amdfam10



gcc: 32 vs 64 Bit

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Advantage of compiling for 64bit machines

- The compiler can make use of
 - at least MMX, SSE and SSE2, since every x86-64 machine supports these.
 - 16 registers (64 bit) instead of 8 registers (32 bit)
 - larger virtual address space (at least 48 bit = 256 TiB)



Automatic Vectorization

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

The compiler makes use of SIMD

source

```
1 float a[128];  
   ...  
3 for(int i=0; i < 128; i++)  
   a[i] *= 2.5f;
```

Example Optimization using AVX

- Width of SIMD registers: 256bit
- Float uses 32bit
- -> 8 calculations in parallel
- 16 * 8 simultaneous multiplications instead of 128 in sequence



Profile Guided Optimization

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

The execution of the program is profiled, so the compiler can learn from the 'behaviour' of the code

Steps

- compile and link it with profiling enabled
- run the program - make sure all the time-critical parts are executed
- profiling data will be written to disk
- recompile making use of the profiling data



Function Ordering

Definition

Re-orders functions to improve instruction cache hit rate

unoptimized

```
1 int foo() {  
2   ... //several lines of code  
3 }  
4 float someFunction() {  
5   ... //several lines of code  
6 }  
7 ... //more functions  
8 int bar() {  
9   ... //several lines of code  
10 }
```

optimized

```
1 int foo() {  
2   ... //several lines of code  
3 }  
4 int bar() {  
5   ... //several lines of code  
6 }  
7 float someFunction() {  
8   ... //several lines of code  
9 }  
10 ... //more functions
```

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources



Basic Block Ordering

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

- Similar to function ordering
- Same goal: improve instruction cache hit rate
- Re-orders blocks



Switch Statement Optimization

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

Sorts the cases in a switch statement by frequency of execution

unoptimized

```
switch(expression)
2 {
3     case constant1:
4         statements; break;
5     case constant2:
6         statements; break;
7     case constant3:
8         statements; break;
9     default:
10        statements;
11 }
```

optimized

```
1 switch(expression)
2 {
3     case constant3:
4         statements; break;
5     case constant1:
6         statements; break;
7     case constant2:
8         statements; break;
9     default:
10        statements;
11 }
```



Improved Register Allocation

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

Keeps the locally most frequently used variables in registers

note

The problem of register allocation is np-hard without profiling



Aiding Optimizations

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Why this is useful

- The compiler 'enforces rules of the C-Standard' to ensure correct programs
- Often the compiler has to make conservative assumptions
- If it had more knowledge about the code, it could optimize more aggressively
- The programmer can help the compiler



Data Layout

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

A good data layout uses memory space and cache more efficiently

unoptimized

```
1 struct foo {  
    char a;  
3   float x[8];  
    char b;  
5   float y[8];  
    char c;  
7   float z[8];  
};
```

optimized

```
struct foo {  
2   float x[8];  
    float y[8];  
4   float z[8];  
    char a;  
6   char b;  
    char c;  
8   };
```



Pragma Vector Aligned

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Definition

- Communicates data layout information to the compiler
- Some architectures contain instructions that execute faster if the data is guaranteed to be aligned on specific memory boundaries

source

```
float a[128];  
2 ...  
   #pragma vector aligned  
4 for(int i=0; i < 128; i++)  
   a[i] *= 2.5f;
```

options

- aligned
- unaligned
- always



'Safe' / 'Unsafe' Optimizations

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

'normal' behaviour

- Most optimizations won't change the result of computations
- especially not the `-o[level]` options
- the compiler is conservative

more optimizations

- compiler options that might change the results
- but the computations may be faster
- caution: only use them if you don't need the precision
- e.g. `-ffast-math`



Definition

Shared-Memory Multithreading Programming Interface

unoptimized

```
1 void foobar(int *a, int n) {  
    for (int i = 0; i < n; i++)  
2     a[i] = 2 * i;  
3 }
```

optimized

```
void foobar(int *a, int n) {  
2   #pragma omp parallel for  
    for (int i = 0; i < n; i++)  
4     a[i] = 2 * i;  
}
```



What we've learned today

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Remember

- Which optimizations the compiler can do by himself -> **readability over manual optimization**
- Tell the compiler details about the destination architecture
- Where the compiler needs some help (aided optimization)
- Optimize manually, where the compiler can't help - but only if you can expect a real performance impact. if not -> **readability over manual optimization**
- or: "Premature Optimization is the root of all evil"



Questions?

C-Compiler

Mirko Köster

Automatic
Optimization

Architecture
Independent

Inter-
Procedural

Architecture
Dependent

Profile Guided
Optimization

Aiding
Optimizations

'Safe' /
'Unsafe'
Optimizations

OpenMP

Conclusion

Questions?

Sources

Thank you for your attention

Questions?



Resources I used to prepare this presentation

- http://en.wikipedia.org/wiki/Optimizing_compiler
- <http://gcc.gnu.org/onlinedocs/gcc-4.7.2/gcc/Optimize-Options.html>
- http://gcc.gnu.org/onlinedocs/gcc/i386-and-x86_002d64-Options.html
- <http://www.embedded.com/design/mcus-processors-and-socs/4008892/Tuning-C-C--compilers-for-optimal-parallel->