

Optimization

Compiler Construction

2025-01-14



Prof. Dr. Michael Kuhn

michael.kuhn@ovgu.de

Parallel Computing and I/O

Institute for Intelligent Cooperating Systems

Faculty of Computer Science

Otto von Guericke University Magdeburg

<https://parcio.ovgu.de>

Outline

Optimization

Overview

Optimization

Summary

- Optimizations are important to improve performance
 - Straightforward translations are often not very fast
- Optimizations can be applied at different levels
 - Local within a basic block, global within a full function, interprocedural across functions
- Optimizations cannot fix everything
 - $O(n^2)$ algorithms will not be simply replaced by $O(n \cdot \log(n))$ ones

Optimization

Overview

Optimization

Summary

- Can the following expressions be optimized?

```
1 i: integer;  
2 j: integer;  
3 k: float;  
4 l: float;  
5  
6 i = i * 9;  
7 i = i * j + i * j;  
8 k = k * 9;  
9 k = k * l + k * l;
```

- Can the following expressions be optimized?

```
1 i: integer;
2 j: integer;
3 k: float;
4 l: float;
5
6 i = i * 9;
7 i = i * j + i * j;
8 k = k * 9;
9 k = k * l + k * l;
```

```
1 i = (i << 3) + i;
2 i = (i * j) << 1;
3 k = ???;
4 k = ???;
```

- Perform loop unrolling for the following snippet. Discuss advantages/disadvantages.

```
1 int a[];
2 int b[];
3 int c[];
4
5 for (int i = 0; i < N; i++) {
6     c[i] = a[i] + b[i];
7 }
```

- Perform loop unrolling for the following snippet. Discuss advantages/disadvantages.

```
1  int i;
2  for (i = 0; i < N - 4; i += 4) {
3      c[i] = a[i] + b[i];
4      c[i + 1] = a[i + 1] + b[i + 1];
5      c[i + 2] = a[i + 2] + b[i + 2];
6      c[i + 3] = a[i + 3] + b[i + 3];
7  }
8  for (; i < N; i++) {
9      c[i] = a[i] + b[i];
10 }
```

- Code size increases
- Can provide opportunities for vectorization

- Perform function inlining for the following code snippet.

```
1 int foo (int a, int b) {  
2     return a + b;  
3 }  
4  
5 int a = 0;  
6 for (int i = 0; i < N; i++) {  
7     a = foo(a, i) * foo(a, i + 1);  
8 }
```

- Perform function inlining for the following code snippet.

```
1 int a = 0;
2 for (int i = 0; i < N; i++) {
3     a = (a + i) * (a + i + 1);
4 }
```

- Code size increases

- Perform function inlining for the following code snippet.

```
1 int foo (int a, int b) {  
2     b = a;  
3     return a + b;  
4 }  
5  
6 int a = 0;  
7 for (int i = 0; i < N; i++) {  
8     a = foo(a, i) * foo(a, i + 1);  
9 }
```

- Perform function inlining for the following code snippet.

```
1 int a = 0;
2 for (int i = 0; i < N; i++) {
3     int foo1, foo2;
4     i = a;
5     foo1 = a + i;
6     i + 1 = a;
7     foo2 = a + i + 1;
8     a = foo1 * foo2;
9 }
```

- Side effects have to be considered

- Perform function inlining for the following code snippet.

```
1 int a = 0;
2 for (int i = 0; i < N; i++) {
3     a = (a + a) * (a + a);
4 }
```

- Perform code hoisting for the following code snippet.

```
1 while (term_iteration > 0) {
2     for (i = 1; i < N; i++) {
3         for (j = 1; j < N; j++) {
4             star = 0.25 * (Matrix[m2][i - 1][j] + Matrix[m2][i][j - 1] +
5                 ↪ Matrix[m2][i][j + 1] + Matrix[m2][i + 1][j]);
6             star += 0.25 * (2 * M_PI * M_PI) * h * h * sin(M_PI * h *
7                 ↪ (double)i) * sin(M_PI * h * (double)j);
8
9             Matrix[m1][i][j] = star;
10        }
11    }
12 }
```

- Perform code hoisting for the following code snippet.

```
1 double pih = M_PI * h;
2 double fpisin = 0.25 * (2 * M_PI * M_PI) * h * h;
3 while (term_iteration > 0) {
4     for (i = 1; i < N; i++) {
5         double fpisin_i = fpisin * sin(pih * (double)i);
6         for (j = 1; j < N; j++) {
7             star = 0.25 * (Matrix[m2][i - 1][j] + Matrix[m2][i][j - 1] +
8                 ↪ Matrix[m2][i][j + 1] + Matrix[m2][i + 1][j]);
9             star += fpisin_i * sin(pih * (double)j);
10
11             Matrix[m1][i][j] = star;
12         }
13     }
```

- Perform dead code elimination for the following code snippet.

```
1 int foo (int* bar) {  
2     bool b = false;  
3     if (b && bar == NULL) {  
4         ...  
5         return 1;  
6     } else {  
7         ...  
8         return 2;  
9     }  
10    return 3;  
11 }
```


- Perform dead code elimination for the following code snippet.

```
1 int foo (int* bar) {  
2     bool b = false;  
3     if (b && bar == NULL) {  
4         ...  
5         return 1;  
6     } else {  
7         ...  
8         return 2;  
9     }  
10 }
```

- Perform dead code elimination for the following code snippet.

```
1 int foo (int* bar) {  
2     ...  
3     return 2;  
4 }
```

- Perform dead code elimination for the following code snippet.

```
1 int foo (int* bar) {  
2     int v = *bar;  
3     if (bar == NULL) {  
4         return NULL;  
5     }  
6     return *bar;  
7 }
```

- Perform dead code elimination for the following code snippet.

```
1 int foo (int* bar) {  
2     if (bar == NULL) {  
3         return NULL;  
4     }  
5     return *bar;  
6 }
```

- Perform dead code elimination for the following code snippet.

```
1 int foo (int* bar) {  
2     return *bar;  
3 }
```

- Might lead to undefined behavior

- Real-world compiler optimizations using partdiff and GCC 14

- Real-world compiler optimizations using partdiff and GCC 14
- -O0: 1,447 lines of assembly

- Real-world compiler optimizations using partdiff and GCC 14
- -O0: 1,447 lines of assembly
- -Og: 923 lines of assembly

- Real-world compiler optimizations using partdiff and GCC 14
- -O0: 1,447 lines of assembly
- -Og: 923 lines of assembly
- -O1: 787 lines of assembly

- Real-world compiler optimizations using partdiff and GCC 14
- -O0: 1,447 lines of assembly
- -Og: 923 lines of assembly
- -O1: 787 lines of assembly
- -O2: 830 lines of assembly

- Real-world compiler optimizations using partdiff and GCC 14
- -O0: 1,447 lines of assembly
- -Og: 923 lines of assembly
- -O1: 787 lines of assembly
- -O2: 830 lines of assembly
- -O3: 1,016 lines of assembly

- Real-world compiler optimizations using partdiff and GCC 14
- -O0: 1,447 lines of assembly
- -Og: 923 lines of assembly
- -O1: 787 lines of assembly
- -O2: 830 lines of assembly
- -O3: 1,016 lines of assembly
- -Ofast: 1,126 lines of assembly

- Real-world compiler optimizations using partdiff and GCC 14
- -O0: 1,447 lines of assembly
- -Og: 923 lines of assembly
- -O1: 787 lines of assembly
- -O2: 830 lines of assembly
- -O3: 1,016 lines of assembly
- -Ofast: 1,126 lines of assembly
- -Os: 679 lines of assembly

Outline

Optimization

Overview

Optimization

Summary

- Be careful about correctness
 - For example, order of floating point operations
- Be careful about external side effects
 - For example, system calls like `write`
- Be careful about changes to debugging
 - For example, breakpoints in inlined functions
- Interactions can be positive or negative
 - For example, constant propagation before reachability analysis

References

[Thain, 2020] Thain, D. (2020). *Introduction to Compilers and Language Design: Second Edition*. <http://compilerbook.org/>.