

# Compilers 101

Optimizations – Part 2

# Previously...

Preprocessing

Lexical analysis

Syntax analysis

Semantic analysis

IR Generation

IR Optimization

Optimization

Code generation

Frontend

Middle-end

Backend

# Previously... (2)

- LLVM IR
- LLVM IR passes
- LLVM IR optimizations

# Today

- opt tool showcase
- Loop optimizations

# Running optimizations

- Generate LLVM IR
- Use opt tool
- Guide: <https://llvm.org/docs/CommandGuide/opt.html>

# opt tool

- LLVM optimizer
  - Input: LLVM IR
  - Output: LLVM IR
- Passes list: <https://llvm.org/docs/Passes.html>
- Options: <https://llvm.org/docs/CommandGuide/opt.html>
- Command line examples:

```
opt -print-passes
```

```
opt -O2 input.ll -o optimized.bc
```

```
opt -passes='dce' input.bc
```

```
opt -dce input.bc
```

```
opt -load MyPass.so -my-custom-pass input.ll -o  
output.ll
```

# Compiler Explorer



Is an interactive compiler exploration website. Edit code in C, C++, C#, F#, Rust, Go, D, Haskell, Swift, Pascal, ispc, Python, Java, or any of the other 30+ supported languages components, and see how that code looks after being compiled in real time.

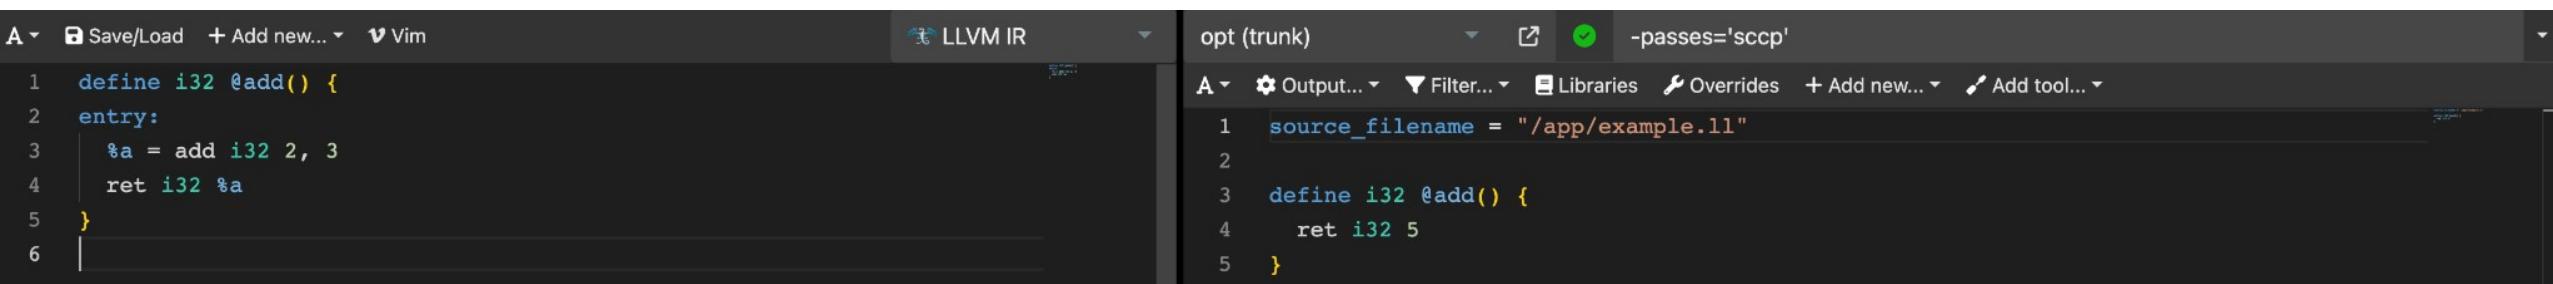
Official website: <https://godbolt.org/>

GitHub repository: <https://github.com/compiler-explorer/compiler-explorer>

# Running LLVM IR optimizations

Learn by example (1):

- SCCP:



The screenshot shows the LLVM DevTools interface. On the left, there is a code editor window titled "LLVM IR" containing the following LLVM IR code:

```
1 define i32 @add() {
2 entry:
3     %a = add i32 2, 3
4     ret i32 %a
5 }
```

On the right, there is a terminal window titled "opt (trunk)" with the command "-passes='sccp'". The output of the command is:

```
1 source_filename = "/app/example.ll"
2
3 define i32 @add() {
4     ret i32 5
5 }
```

# CFG visualization

```
opt -dot-cfg cfg_example.ll  
dot -Tpng cfg_example.dot
```

The screenshot shows the LLVM DevTools interface with three main windows:

- LLVM IR source #1:** Displays the LLVM IR code for a function named `test`. The code includes an entry block, a `check` block, and two `if` blocks (`if_true` and `if_false`) that both lead to a `ret i32` instruction.
- LLVM IR:** An LLVM IR editor window showing the same code as the source file.
- CFG opt (trunk) (Editor #1, Compiler #1):** A Control Flow Graph (CFG) visualization. The graph consists of four nodes:
  - A root node labeled `test` with a single outgoing edge to the `check` node.
  - The `check` node contains the LLVM IR for the `check` block, including the `%cmp = icmp eq i32 %x, 0` comparison and the `br il %cmp, label %if_true, label %if_false` branch.
  - The `if_true` node contains the LLVM IR for the `if_true` block, which simply `ret i32 1`.
  - The `if_false` node contains the LLVM IR for the `if_false` block, which simply `ret i32 0`.Arrows indicate the flow from the `test` node to the `check` node, and from the `check` node to both the `if_true` and `if_false` nodes.

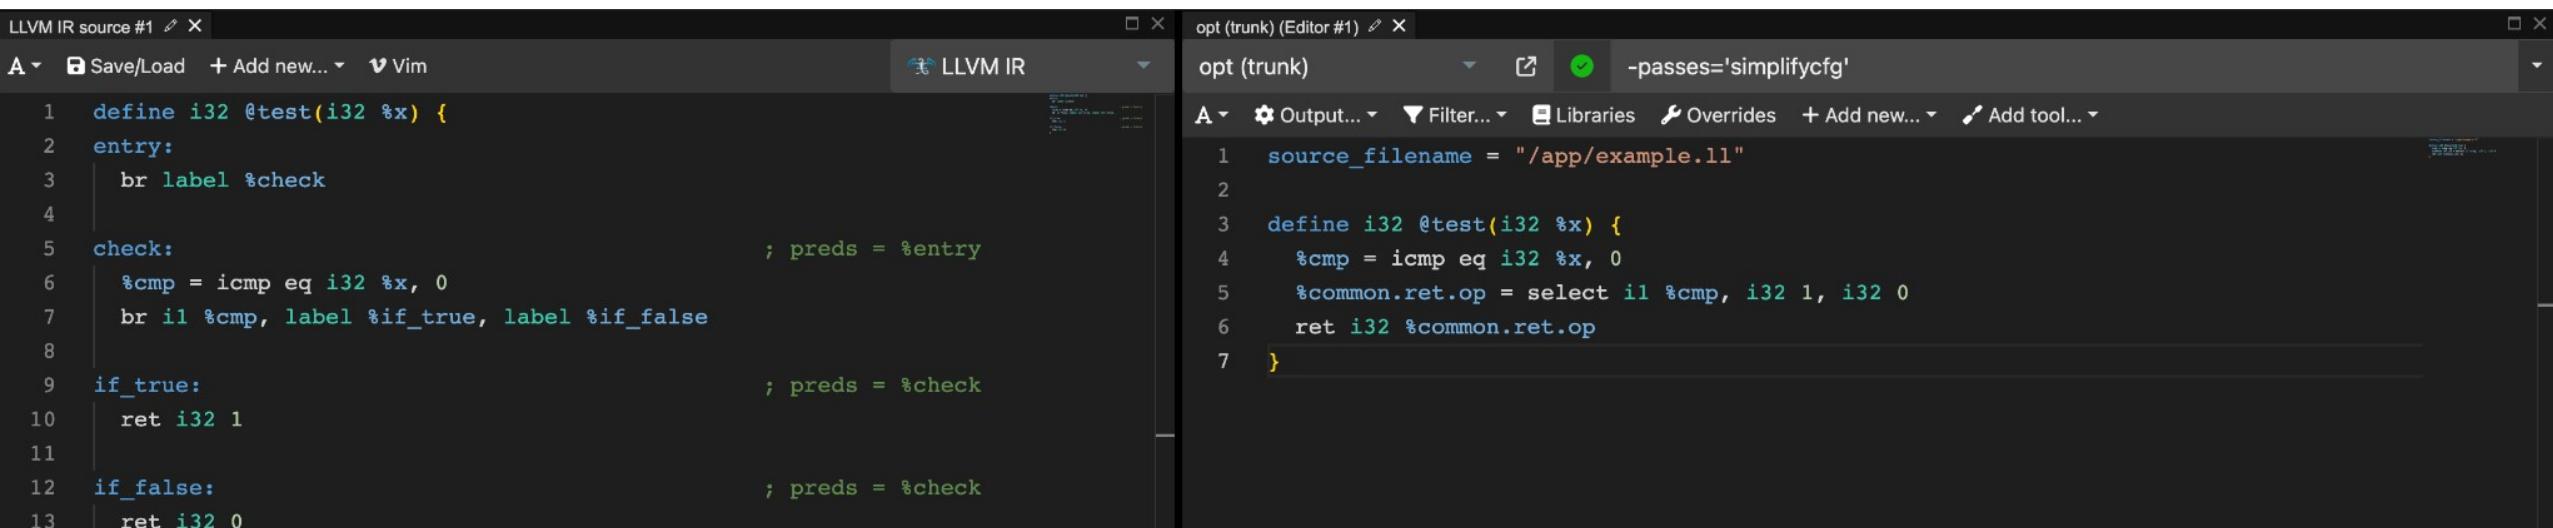
At the bottom of the interface, there is a status bar with the text "Layout time: 1ms" and "Basic blocks: 4".

<https://godbolt.org/z/T6bWarjjK>

# Running LLVM IR optimizations

Learn by example (2):

- CFG simplification:



The screenshot shows two code editors side-by-side. The left editor is titled 'LLVM IR source #1' and contains the following LLVM IR code:

```
define i32 @test(i32 %x) {
entry:
  br label %check

check:                                ; preds = %entry
  %cmp = icmp eq i32 %x, 0
  br i1 %cmp, label %if_true, label %if_false

if_true:                                 ; preds = %check
  ret i32 1

if_false:                               ; preds = %check
  ret i32 0
```

The right editor is titled 'opt (trunk) (Editor #1)' and contains the following command-line output:

```
opt (trunk)
      -passes='simplifycfg'

source_filename = "/app/example.ll"
define i32 @test(i32 %x) {
  %cmp = icmp eq i32 %x, 0
  %common.ret.op = select i1 %cmp, i32 1, i32 0
  ret i32 %common.ret.op}
```

# Loop optimizations

Why are they important?

- Pareto principle
  - 20 percent of effort gives 80 percent of result
- Most computationally intensive algorithms involve loops

# Terminology

- Trip count is a minimum number of times a loop executes
- Induction variable is a variable that gets increased or decreased on each loop iteration
- Loop invariant is the part of a computation inside a loop that remains unchanged (invariant) in each iteration

# Detecting induction variables

- $i$  is a basic induction variable in a loop  $L$  if the only definitions of  $i$  within  $L$  are of the form  $i := i +/ - c$  where  $c$  is loop invariant
- $k$  is a derived induction variable in loop  $L$  if
  - There's only 1 definition of  $k$  within  $L$  in the form of  $k := j */+ c$  where  $j$  is an induction variable
  - If  $j$  is an induction variable in the family of  $i$  then
    - The only definition of  $j$  that reaches  $k$  is the one in the loop
    - There's no definition of  $i$  on any path between the definition of  $j$  and the definition of  $k$

# Structured control flow

Structured control flow is a programming paradigm characterized by the clear, hierarchical organization of control paths through a program.

- It relies exclusively on structured control constructs such as conditionals (if, else), loops (for, while), and function calls, which lead to well-defined entry and exit points.
- The key attribute of structured CF is that it avoids the use of arbitrary jumps or goto statements, making the control flow predictable and easier to follow.

# Structured control flow

## Edgar Dijkstra: Go To Statement Considered Harmful

### Go To Statement Considered Harmful

Key Words and Phrases: go to statement, jump instruction, branch instruction, conditional clause, alternative clause, repetitive clause, program intelligibility, program sequencing

CR Categories: 4.22, 5.23, 5.24

#### EDITOR:

For a number of years I have been familiar with the observation that the quality of programmers is a decreasing function of the density of **go to** statements in the programs they produce. More recently I discovered why the use of the **go to** statement has such disastrous effects, and I became convinced that the **go to** statement should be abolished from all "higher level" programming languages (i.e. everything except, perhaps, plain machine code). At that time I did not attach too much importance to this discovery; I now submit my considerations for publication because in very recent discussions in which the subject turned up, I have been urged to do so.

My first remark is that, although the programmer's activity ends when he has constructed a correct program, the process taking place under control of his program is the true subject

dynamic progress is only characterized when we also give to which call of the procedure we refer. With the inclusion of procedures we can characterize the progress of the process via a sequence of textual indices, the length of this sequence being equal to the dynamic depth of procedure calling.

Let us now consider repetition clauses (like, **while B repeat A** or **repeat A until B**). Logically speaking, such clauses are now superfluous, because we can express repetition with the aid of recursive procedures. For reasons of realism I don't wish to exclude them: on the one hand, repetition clauses can be implemented quite comfortably with present day finite equipment; on the other hand, the reasoning pattern known as "induction" makes us well equipped to retain our intellectual grasp on the processes generated by repetition clauses. With the inclusion of the repetition clauses textual indices are no longer sufficient to describe the dynamic progress of the process. With each entry into a repetition clause, however, we can associate a so-called "dynamic index," inexorably counting the ordinal number of the corresponding current repetition. As repetition clauses (just as procedure calls) may be applied nestedly, we find that now the progress of the process can always be uniquely characterized by a

# Structured control flow

Is this program structured?

```
1 void foo(char *a, int N) {  
2     for (int i = 0; i < N; i++) {  
3         if (a[i] == 42)  
4             break;  
5     }  
6 }
```

# Loop-invariant code motion

```
1 void baz();  
2 void bar(int i);  
3  
4 void foo(char *a, int N) {  
5     for (int i = 0; i < N; i++) {  
6         bar(i);  
7         baz();  
8         a[0] = 42;  
9     }  
10 }
```

```
1 void baz();  
2 void bar(int i);  
3  
4 void foo(char *a, int N) {  
5     a[0] = 42;  
6     for (int i = 0; i < N; i++) {  
7         bar(i);  
8         baz();  
9     }  
10 }
```

# Conditions for safe hoisting

- hoisting - moving to outer scope
- sinking - moving to inner scope
  
- An invariant assignment  $d: x := y \text{ op } z$  is safe to hoist if
  - $d$  dominates all loop exits at which  $x$  is live and
  - there is only one definition of  $x$  in the loop, and
  - $x$  is not live at the entry point for the loop

# Loop unroll

```
1 void bar(int i);
2
3 void foo(char *a, int N) {
4     for (int i = 0; i < N / 4; i += 4) {
5         bar(i + 0);
6         bar(i + 1);
7         bar(i + 2);
8         bar(i + 3);
9     }
10    for (int i = N - N % 2; i < N; i++) {
11        bar(i);
12    }
13}
```

# When it works and when it fails?

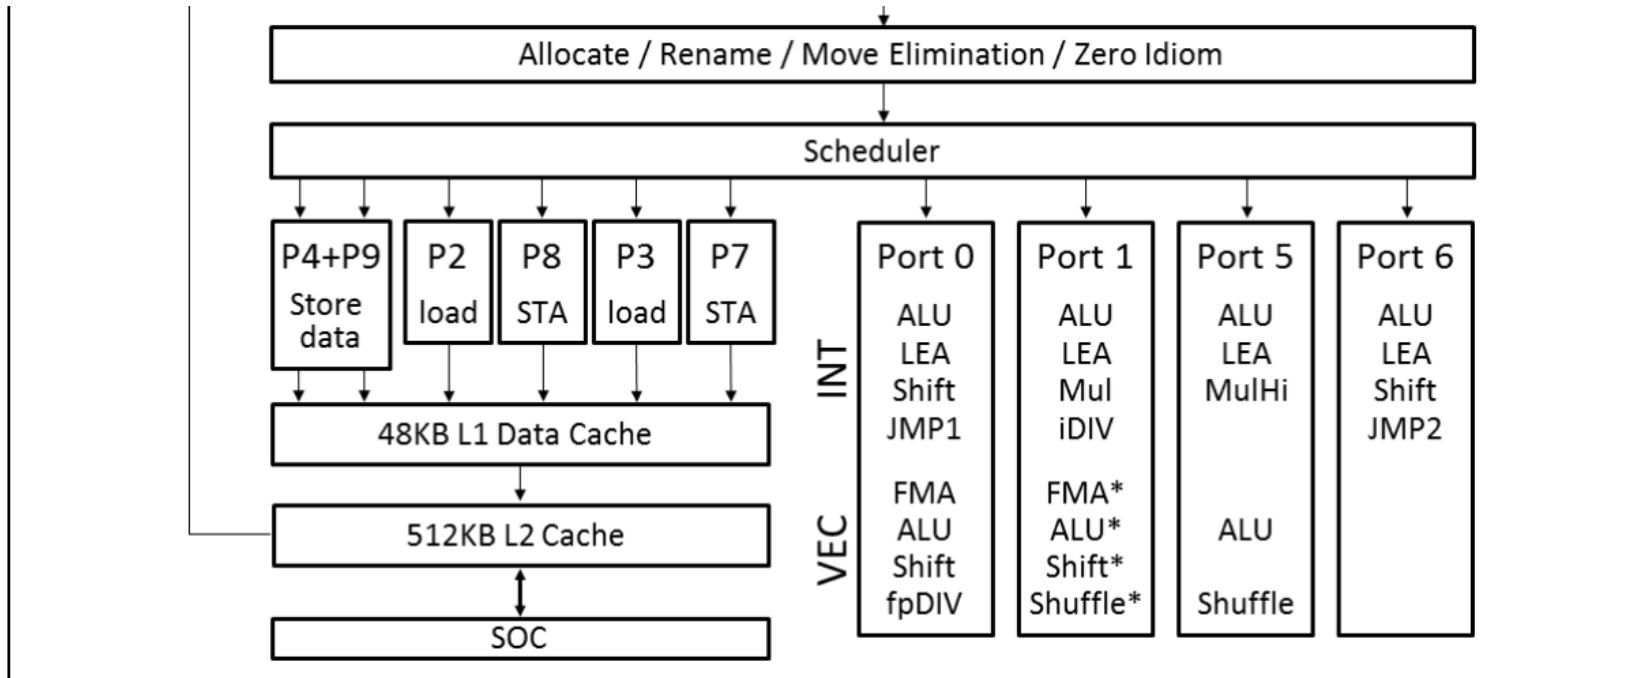


Figure 2-1. Processor Core Pipeline Functionality of the Ice Lake Client Microarchitecture<sup>1</sup>

<https://www.intel.com/content/www/us/en/developer/articles/technical/intel-sdm.html>

# Loop fusion

```
1 void bar(int i);
2
3 void foo(char *a, int N) {
4     for (int i = 0; i < N; i++) {
5         a[i] = 42;
6     }
7     for (int i = N; i >= 0; i--) {
8         bar(i);
9     }
10 }
```

```
1 void bar(int i);
2
3 void foo(char *a, int N) {
4     for (int i = 0; i < N; i++) {
5         a[i] = 42;
6         bar(N - i);
7     }
8 }
```

# Loop peeling

```
1 void foo(int *a, int *b, int N) {  
2     int p = 10;  
3     for (int i = 0; i < N; i++) {  
4         a[i] = b[i] + b[p];  
5         p = i;  
6     }  
7 }
```

```
1 void foo(int *a, int *b, int N) {  
2     a[0] = b[0] + b[10];  
3     for (int i = 1; i < N; i++) {  
4         a[i] = b[i] + b[i - 1];  
5     }  
6 }
```

# Loop unswitch

```
1 void foo(int *a, int *b, bool c, int N) {  
2     for (int i = 0; i < N; i++) {  
3         a[i] = i;  
4         if (c) {  
5             b[i] = 42;  
6         }  
7     }  
8 }
```

```
1 void foo(int *a, int *b, bool c, int N) {  
2     if (c) {  
3         for (int i = 0; i < N; i++) {  
4             a[i] = i;  
5             b[i] = 42;  
6         }  
7     } else {  
8         for (int i = 0; i < N; i++) {  
9             a[i] = i;  
10        }  
11    }  
12 }
```

# Control flow sink

```
1 void bazz();
2 void bazz(int i);
3 void bar(int i);
4
5 void foo(char *a, int N) {
6     a[0] = 42;
7     for (int i = 0; i < N; i++) {
8         bar(i);
9         bazz();
10        bazz(a[0]);
11    }
12}
```

```
1 void bazz();
2 void bazz(int i);
3 void bar(int i);
4
5 void foo(char *a, int N) {
6     for (int i = 0; i < N; i++) {
7         bar(i);
8         bazz();
9         a[0] = 42;
10        bazz(a[0]);
11    }
12}
```

# Loop reroll

```
1 void bar(int i);
2
3 void foo(char *a, int N) {
4     for (int i = 0; i < N / 4; i += 4) {
5         bar(i + 0);
6         bar(i + 1);
7         bar(i + 2);
8         bar(i + 3);
9     }
10    for (int i = N - N % 2; i < N; i++) {
11        bar(i);
12    }
13 }
```

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

# Loop fission

```
1 void bar(int i);
2
3 void foo(char *a, int N) {
4     for (int i = 0; i < N; i++) {
5         a[i] = 42;
6         bar(N - i);
7     }
8 }
```

```
1 void bar(int i);
2
3 void foo(char *a, int N) {
4     for (int i = 0; i < N; i++) {
5         a[i] = 42;
6     }
7     for (int i = N; i >= 0; i--) {
8         bar(i);
9     }
10 }
```

# Loop tiling

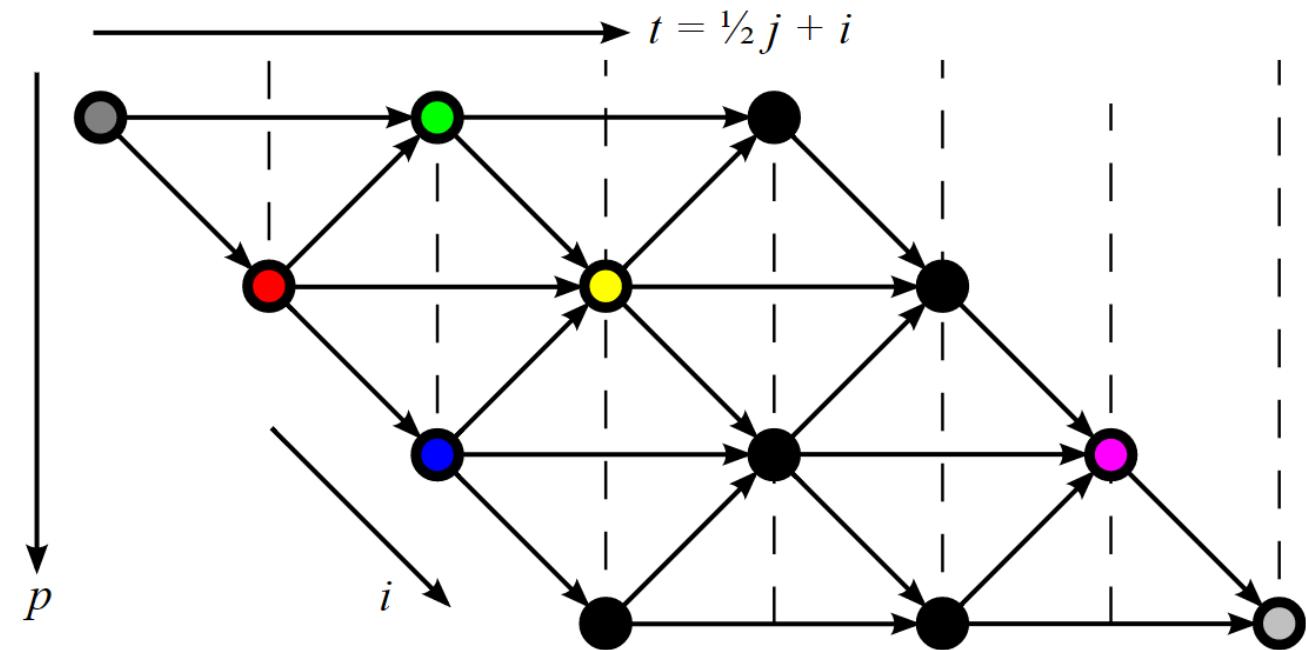
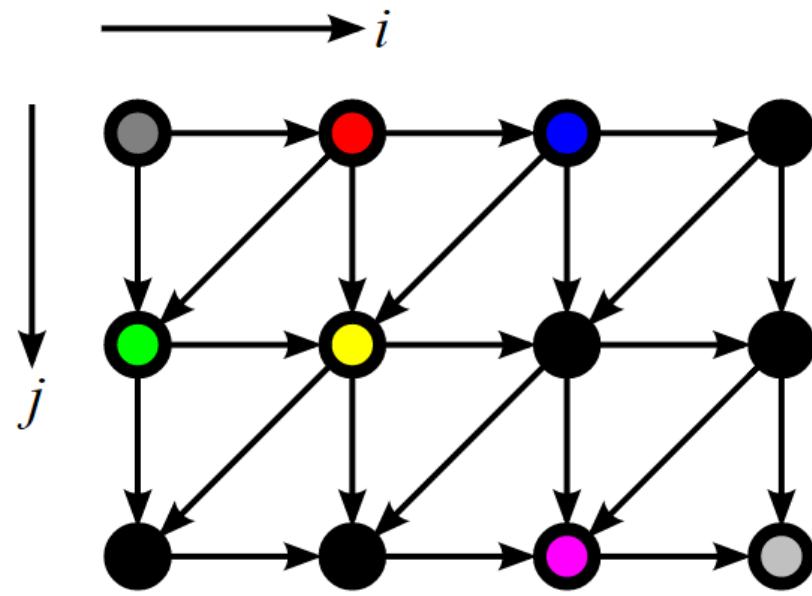
```
1 void foo(int N) {  
2     for (int i = 0; i < N; i++) {  
3         // ...  
4     }  
5 }  
  
1 #include <algorithm>  
2  
3 void foo(int N) {  
4     int B = 10;  
5     for (int i = 0; i < N; i += B) {  
6         for (int j = i; i < std::min(N, j + B); j++) {  
7             // ...  
8         }  
9     }  
}
```

# Loop rotation

```
1 void foo(int *a, int *b, int n) {  
2     for (int i = 0; i < n; i++)  
3         a[i] = a[i] + b[i];  
4 }
```

```
1 void foo(int *a, int *b, int n) {  
2     int i = 0;  
3     if (n > 0) {  
4         do {  
5             a[i] = a[i] + b[i];  
6             i++;  
7         } while (i < n);  
8     }  
9 }
```

# Polytope model



[https://en.wikipedia.org/wiki/Polytope\\_model](https://en.wikipedia.org/wiki/Polytope_model)

# Affine transformations

```
#define ERR(x, y) (dst[x][y] - src[x][y])

void dither(unsigned char** src, unsigned char** dst, int w, int h)
{
    int i, j;
    for (j = 0; j < h; ++j) {
        for (i = 0; i < w; ++i) {
            int v = src[i][j];
            if (i > 0)
                v -= ERR(i - 1, j) / 2;
            if (j > 0) {
                v -= ERR(i, j - 1) / 4;
                if (i < w - 1)
                    v -= ERR(i + 1, j - 1) / 4;
            }
            dst[i][j] = (v < 128) ? 0 : 255;
            src[i][j] = (v < 0) ? 0 : (v < 255) ? v : 255;
        }
    }
}
```

```
void dither_skewed(unsigned char **src, unsigned char **dst, int w, int h)
{
    int t, p;
    for (t = 0; t < (w + (2 * h)); ++t) {
        int pmin = max(t % 2, t - (2 * h) + 2);
        int pmax = min(t, w - 1);
        for (p = pmin; p <= pmax; p += 2) {
            int i = p;
            int j = (t - p) / 2;
            int v = src[i][j];
            if (i > 0)
                v -= ERR(i - 1, j) / 2;
            if (j > 0)
                v -= ERR(i, j - 1) / 4;
            if (j > 0 && i < w - 1)
                v -= ERR(i + 1, j - 1) / 4;
            dst[i][j] = (v < 128) ? 0 : 255;
            src[i][j] = (v < 0) ? 0 : (v < 255) ? v : 255;
        }
    }
}
```

# Polyhedral transformations in LLVM

Polly is a high-performance loop and data-locality optimizer embedded framework within the LLVM. It utilizes mathematical models to analyze and optimize memory access patterns, significantly enhancing code efficiency, particularly in loops and matrix multiplications

Docs & info: <https://polly.llvm.org/>

# Next time...

- Vectorization

# Lab assignment #2

- Write a plugin with simple LLVM pass
- Task list in Google Docs
- Deadline: April, 9
- Where to seek help
  - <https://www.llvm.org/docs/WritingAnLLVMNewPMPass.html>
  - <https://github.com/llvm/llvm-project/tree/main/llvm/examples/IRTransforms>
  - contact teachers

# Test

<https://forms.gle/ywfMNSijV2XqDiVF8>

Submission time: 10 minutes

Почему loop unroll оптимизация применима и эффективна не всегда?

Your answer

---

Приведите примеры оптимизационных пассов для циклов. Объясните, зачем каждая может быть нужна.

Note: точное название пасса и флага не требуется

Your answer

---



# Extra materials

- Loop Optimizations in LLVM: The Good, The Bad, and The Ugly -  
<https://llvm.org/devmtg/2018-10/slides/Kruse-LoopTransforms.pdf>
- EuroLLVM Developers' Meeting: J. Doerfert & T. Grosser "Analyzing and Optimizing ... Polly" -  
[https://www.youtube.com/watch?v=mXve\\_W4XU2g](https://www.youtube.com/watch?v=mXve_W4XU2g)
- LLVM Developers' Meeting: T. Grosser "Polly - Polyhedral optimizations in LLVM" -  
<https://www.youtube.com/watch?v=WwfZkQEuwEE>

