

Loop unrolling in details

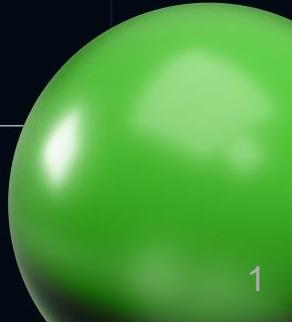
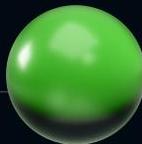


Иван
Афанасьев

Unipro



C++ Russia
2023



Bio



Иван
Афанасьев

- ММФ НГУ. к.ф.м.н.
- 12 лет C++
- 2 года в разработке
компиляторов 

Outline of the talk

- Loop from compiler perspective
- Loop unrolling basics
- Loop unrolling overhead
- GCC and CLANG unroll details
- New optimization opportunities after unrolling
- Example

Outline of the talk

- Loop from compiler perspective
- Loop unrolling basics
- Loop unrolling overhead
- GCC and CLANG unroll details
- New optimization opportunities after unrolling
- Example

Loop from compiler perspective

```
for (i = s; i < b; i++)  
    statements(i);
```

Loop from compiler perspective

```
for (i = s; i < b; i++)  
    statements(i);
```

```
if (s < b) {  
    i = s;  
    do {  
        statements(i);  
        i++;  
    } while (i < b);  
}
```

Loop from compiler perspective

```
for (i = s; i < b; i++)  
    statements(i);
```

```
bool cond = s < b;  
if (!cond) goto LBB_EXIT;  
i = s;
```

LBB_HEAD:

```
statements(i);
```

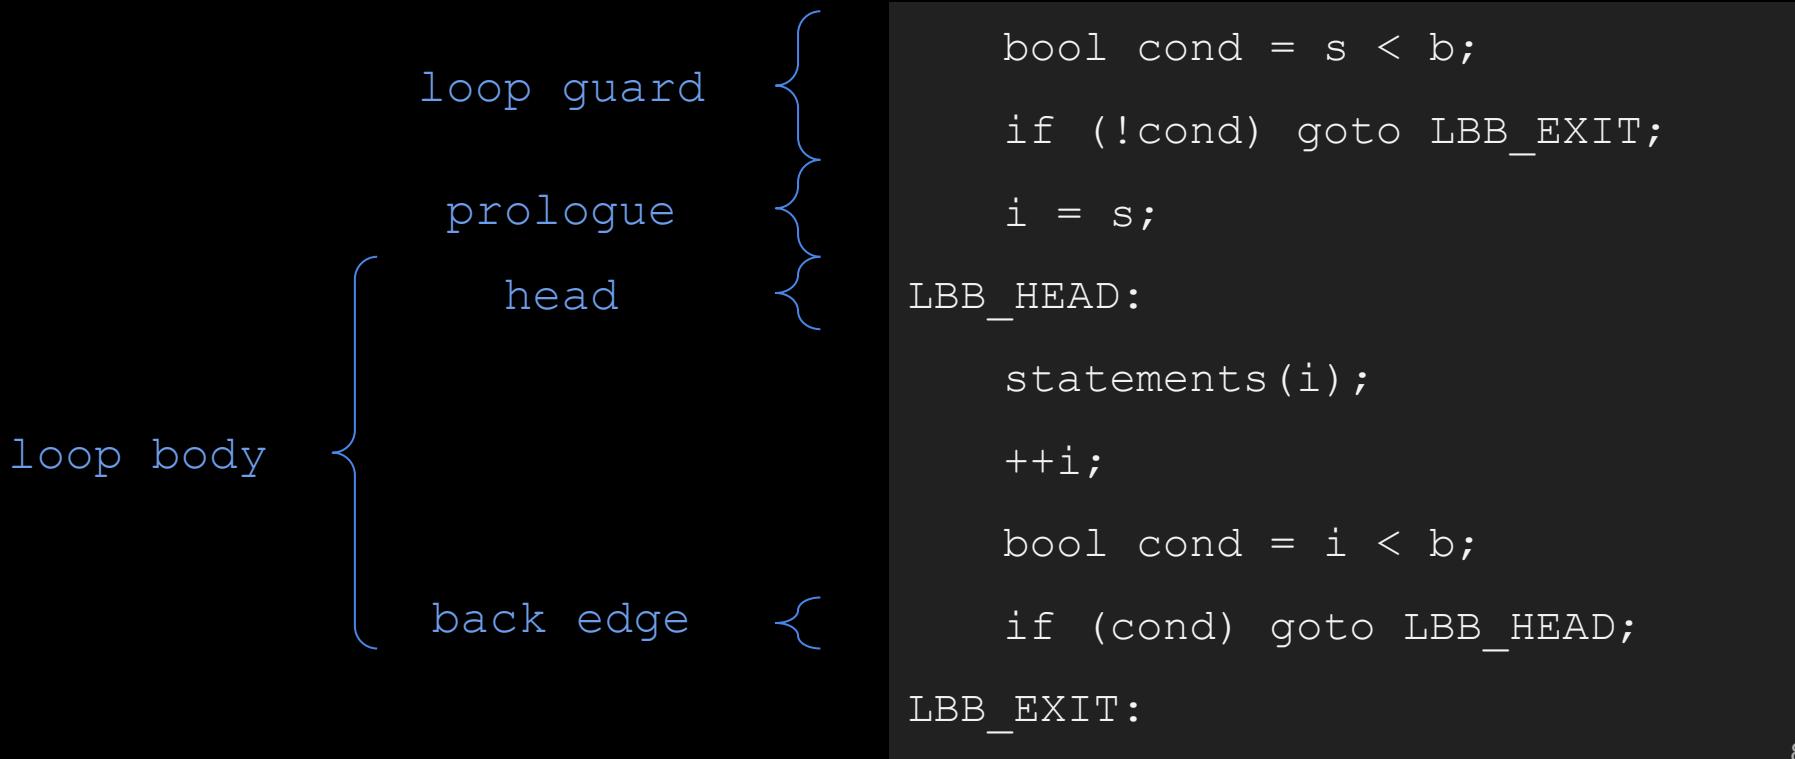
```
++i;
```

```
bool cond = i < b;
```

```
if (cond) goto LBB_HEAD;
```

LBB_EXIT:

Loop from compiler perspective



Loop payload and administration overhead

```
for (i = s; i < b; i++)  
    statements(i);
```

```
        bool cond = s < b;  
        if (!cond) goto LBB_EXIT;  
        i = s;  
  
LBB_HEAD:  
    statements(i);  
    ++i;  
    bool cond = i < b;  
    if (cond) goto LBB_HEAD;  
  
LBB_EXIT:
```

Loop payload and administration overhead

```
for (i = s; i < b; i++)  
    statements(i);
```

payload

```
    bool cond = s < b;  
    if (!cond) goto LBB_EXIT;  
    i = s;
```

LBB_HEAD:

```
    statements(i);  
    ++i;
```

```
    bool cond = i < b;
```

```
    if (cond) goto LBB_HEAD;
```

LBB_EXIT:

Loop payload and administration overhead

```
for (i = s; i < b; i++)
```

```
    statements(i);
```

payload

loop administration
overhead

```
    bool cond = s < b;  
    if (!cond) goto LBB_EXIT;  
    i = s;
```

LBB_HEAD:

```
    statements(i);  
    ++i;
```

```
    bool cond = i < b;  
    if (cond) goto LBB_HEAD;
```

LBB_EXIT:

Loop payload and administration overhead

```
for (i = s; i < b; i++)  
    statements(i);
```

payload

loop administration
overhead



```
bool cond = s < b;  
if (!cond) goto LBB_EXIT;  
i = s;
```

LBB_HEAD:

```
statements(i);  
++i;
```

```
bool cond = i < b;  
if (cond) goto LBB_HEAD;
```

LBB_EXIT:

Examples. Loop

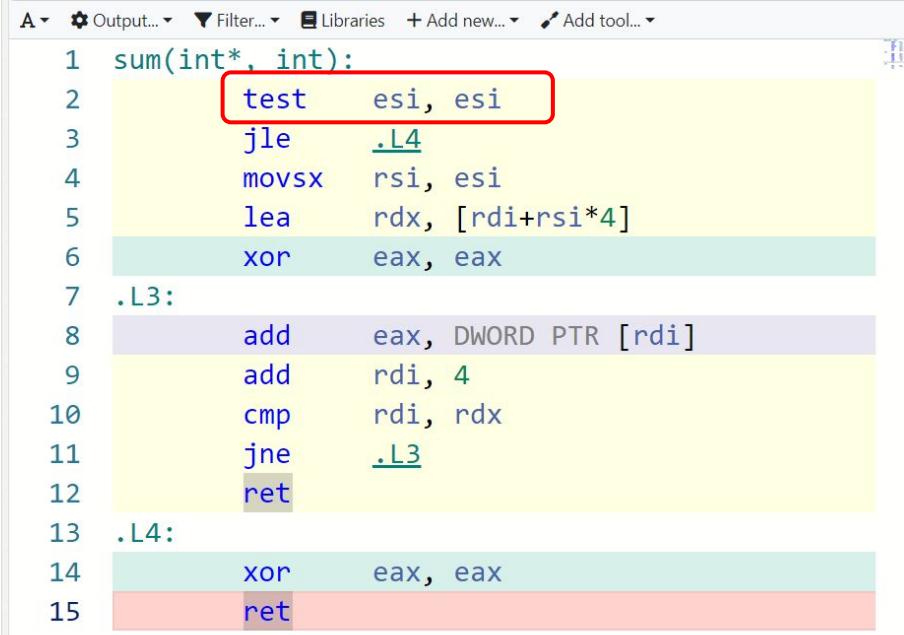
```
1 int sum(int *a, int n) {  
2     int x = 0;  
3     for (int i = 0; i < n; ++i)  
4         x += a[i];  
5     return x;  
6 }
```

```
A Output... Filter... Libraries + Add new... Add tool...  
1 sum(int*, int):  
2     test    esi, esi  
3     jle     .L4  
4     movsx   rsi, esi  
5     lea     rdx, [rdi+rsi*4]  
6     xor    eax, eax  
7 .L3:  
8     add    eax, DWORD PTR [rdi]  
9     add    rdi, 4  
10    cmp    rdi, rdx  
11    jne    .L3  
12    ret  
13 .L4:  
14    xor    eax, eax  
15    ret
```

Examples. Loop

TEST - Logical Compare

Computes the bitwise logical AND of the first operand and the second operand and sets the SF, ZF, and PF status flags.



The screenshot shows assembly code in a debugger interface. The code is color-coded by section:

- Line 1: `sum(int*, int):` (Yellow background)
- Line 2: `test esi, esi` (Yellow background, highlighted with a red box)
- Line 3: `jle .L4` (Yellow background)
- Line 4: `movsx rsi, esi` (Yellow background)
- Line 5: `lea rdx, [rdi+rsi*4]` (Yellow background)
- Line 6: `xor eax, eax` (Light blue background)
- Line 7: `.L3:` (Grey background)
- Line 8: `add eax, DWORD PTR [rdi]` (Grey background)
- Line 9: `add rdi, 4` (Yellow background)
- Line 10: `cmp rdi, rdx` (Yellow background)
- Line 11: `jne .L3` (Yellow background)
- Line 12: `ret` (Grey background)
- Line 13: `.L4:` (Grey background)
- Line 14: `xor eax, eax` (Light blue background)
- Line 15: `ret` (Red background)

Examples. Loop

JLE - Conditional Jump

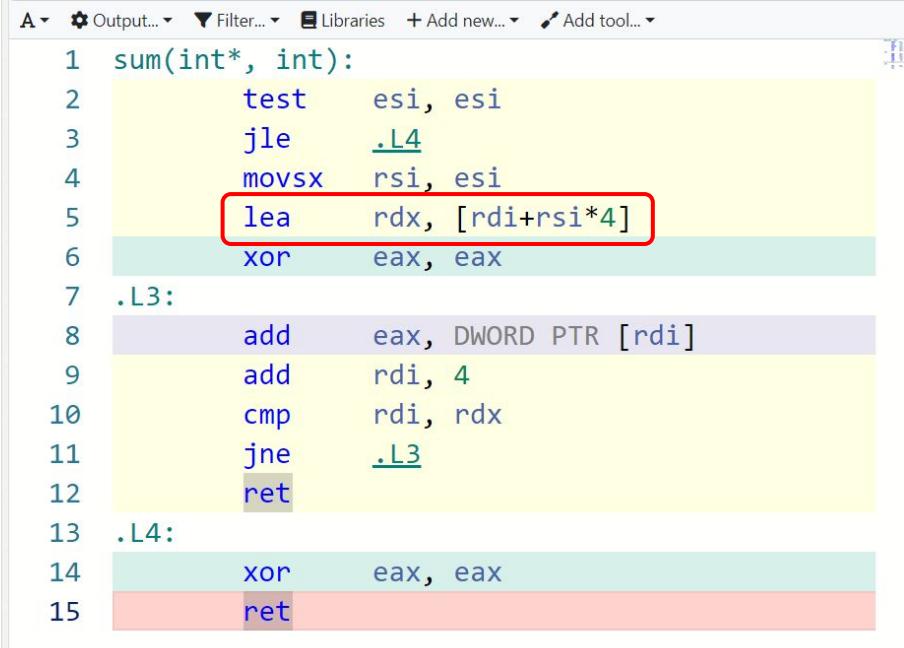
Jump to the destination if one or more of the status flags is set (... ZF ...).

```
1 sum(int*, int):
2     test    esi, esi
3     jle     .L4
4     movsx   rsi, esi
5     lea     rdx, [rdi+rsi*4]
6     xor    eax, eax
7 .L3:
8     add    eax, DWORD PTR [rdi]
9     add    rdi, 4
10    cmp    rdi, rdx
11    jne    .L3
12    ret
13 .L4:
14    xor    eax, eax
15    ret
```

Examples. Loop

LEA - Load Effective Address

Computes the effective address of the second operand and stores it in the first operand.

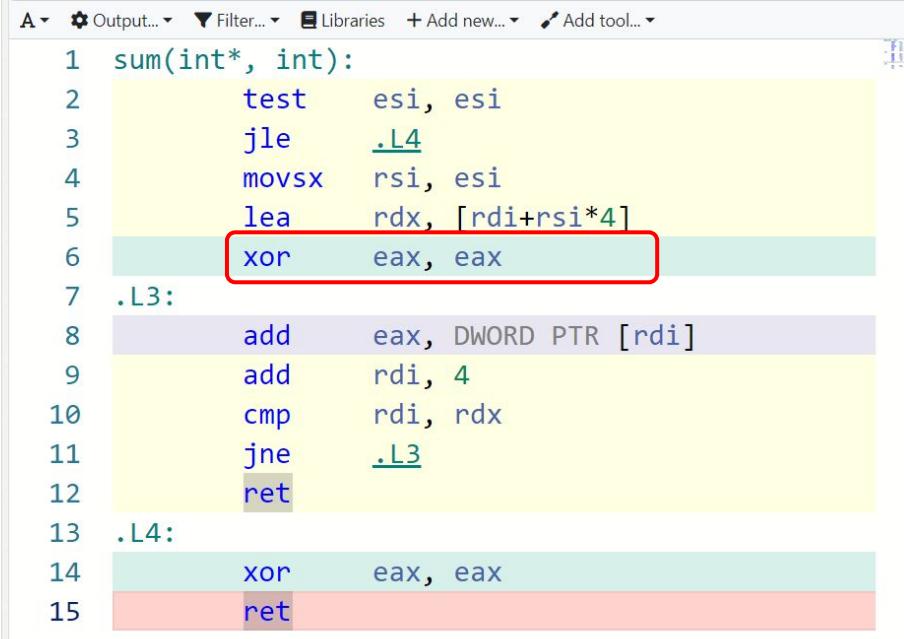


```
A Output... Filter... Libraries + Add new... Add tool...
1 sum(int*, int):
2     test    esi, esi
3     jle     .L4
4     movsx   rsi, esi
5     lea     rdx, [rdi+rsi*4] // Line 5 is highlighted
6     xor    eax, eax
7 .L3:
8     add    eax, DWORD PTR [rdi]
9     add    rdi, 4
10    cmp   rdi, rdx
11    jne   .L3
12    ret
13 .L4:
14    xor    eax, eax
15    ret
```

Examples. Loop

XOR

Performs a bitwise XOR operation on the first and second operands and stores the result in the first operand.



```
A Output... Filter... Libraries + Add new... Add tool...
1 sum(int*, int):
2     test    esi, esi
3     jle     .L4
4     movsx   rsi, esi
5     lea     rdx, [rdi+rsi*4]
6     xor     eax, eax
7 .L3:
8     add    eax, DWORD PTR [rdi]
9     add    rdi, 4
10    cmp    rdi, rdx
11    jne    .L3
12    ret
13 .L4:
14    xor    eax, eax
15    ret
```

Examples. Loop

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

```
A Output... Filter... Libraries + Add new... Add tool...  
1 sum(int*, int):  
2     test    esi, esi  
3     jle     .L4  
4     movsx   rsi, esi  
5     lea     rdx, [rdi+rsi*4]  
6     xor    eax, eax  
7 .L3:  
8     add    eax, DWORD PTR [rdi]  
9     add    rdi, 4  
10    cmp    rdi, rdx  
11    jne    .L3  
12    ret  
13 .L4:  
14    xor    eax, eax  
15    ret
```

Examples. Loop

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

Output... Filter... Libraries + Add new... Add tool...

```
1 sum(int*, int):  
2     test    esi, esi  
3     jle     .L4  
4     movsx   rsi, esi  
5     lea     rdx, [rdi+rsi*4]  
6     xor    eax, eax  
7 .L3:  
8     add    eax, DWORD PTR [rdi]  
9     add    rdi, 4  
10    cmp    rdi, rdx  
11    jne    .L3  
12    ret  
13 .L4:  
14    xor    eax, eax  
15    ret
```

loop guard

Examples. Loop

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

Output... Filter... Libraries + Add new... Add tool...

```
1 sum(int*, int):  
2     test    esi, esi  
3     jle     .L4  
4     movsx   rsi, esi  
5     lea     rdx, [rdi+rsi*4]  
6     xor    eax, eax  
7 .L3:  
8     add    eax, DWORD PTR [rdi]  
9     add    rdi, 4  
10    cmp    rdi, rdx  
11    jne    .L3  
12    ret  
13 .L4:  
14    xor    eax, eax  
15    ret
```

loop guard
prologue

Examples. Loop

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

Output... Filter... Libraries + Add new... Add tool...

```
1 sum(int*, int):  
2     test    esi, esi  
3     jle     .L4  
4     movsx   rsi, esi  
5     lea     rdx, [rdi+rsi*4]  
6     xor    eax, eax  
  
.L3:  
7     add    eax, DWORD PTR [rdi]  
8     add    rdi, 4  
9     cmp    rdi, rdx  
10    jne    .L3  
11    ret  
12  
.L4:  
13    xor    eax, eax  
14  
15    ret
```

loop guard
prologue
iteration check
back edge

Examples. Loop

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

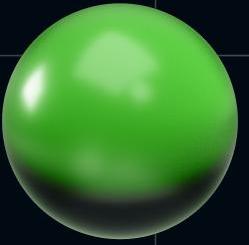
Output... Filter... Libraries + Add new... Add tool...

```
1 sum(int*, int):  
2     test    esi, esi  
3     jle     .L4  
4     movsx   rsi, esi  
5     lea     rdx, [rdi+rsi*4]  
6     xor    eax, eax  
  
.L3:  
7     add    eax, DWORD PTR [rdi]  
8     add    rdi, 4  
9     cmp    rdi, rdx  
10    jne    .L3  
11    ret  
12  
.L4:  
13    xor    eax, eax  
14  
15    ret
```

loop guard
prologue
payload
iteration check
back edge

Outline of the talk

- Loop from compiler perspective
- Loop unrolling basics
- Loop unrolling overhead
- GCC and CLANG unroll details
- New optimization opportunities after unrolling
- Example



Loop unrolling

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

Loop unrolling

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

N_jumps /= 4

N_checks /= 4

Loop unrolling

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

N_jumps /= 4

N_checks /= 4

(b - s) % 4 == 0 ?

Prolog + epilogue loop

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Overhead:

- Epilogue (remainder) loop
- Calculate b'

Prolog + epilogue loop

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Overhead:

- Epilogue (remainder) loop
- Calculate b'

Prolog + epilogue loop

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Overhead:

- Epilogue (remainder) loop
- Calculate b'

Prolog + epilogue loop

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$v1 = b - s$

$v2 = v1 \% 4$

$b' = b - v2$

Prolog + epilogue loop

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$\begin{aligned} v1 &= b - s \\ v2 &= v1 \% 4 \\ b' &= b - v2 \end{aligned}$$

stall!

(latency)

Prolog + epilogue loop

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$v1 = b - s$$

$$v2 = \boxed{v1 \% 4}$$

$$b' = b - v2$$

Prolog + epilogue loop

- `v1 % UNROLLED_STEP` is expensive in general case
- But ...

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$v1 = b - s$$

$$v2 = \boxed{v1 \% 4}$$

$$b' = b - v2$$

Prolog + epilogue loop

- $v1 \% \text{UNROLLED_STEP}$ is expensive in general case
- But ...
 - If $v1 \geq 0$

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$v1 = b - s$$

$$v2 = v1 \% 4$$

$$b' = b - v2$$

Prolog + epilogue loop

- $v1 \% \text{UNROLLED_STEP}$ is expensive in general case
- But ...
 - If $v1 \geq 0$
 - If $\text{UNROLLED_STEP} == 2^{**K}$

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$v1 = b - s$$

$$v2 = v1 \% 4$$

$$b' = b - v2$$

Prolog + epilogue loop

- `V1 % UNROLLED_STEP` is expensive in general case
- But ...
 - If `V1 >= 0`
 - If `UNROLLED_STEP == 2**K`

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$v1 = b - s$$

$$v2 = \boxed{v1 \% 4}$$

$$b' = b - v2$$

<code>V1</code>	0	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?
<code>2**K</code>	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0

Prolog + epilogue loop

- `V1 % UNROLLED_STEP` is expensive in general case
- But ...
 - If `V1 >= 0`
 - If `UNROLLED_STEP == 2**K`

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$v1 = b - s$$

$$v2 = \boxed{v1 \% 4}$$

$$b' = b - v2$$

<code>V1</code>	0	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?
<code>2**K</code>	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0

Prolog + epilogue loop

- $v1 \% \text{UNROLLED_STEP}$ is expensive in general case
- But ...
 - If $v1 \geq 0$
 - If $\text{UNROLLED_STEP} == 2^{**K}$

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$v1 = b - s$$

$$v2 = v1 \% 4$$

$$b' = b - v2$$

$v1$	0	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?
2^{**K}	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0
$2^{**K} - 1$	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1

Prolog + epilogue loop

- $v1 \% \text{UNROLLED_STEP}$ is expensive in general case
- But ...
 - If $v1 \geq 0$
 - If $\text{UNROLLED_STEP} == 2^{**K}$

$v1 \% 2^{**K} == v1 \& (2^{**K} - 1)$

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$v1 = b - s$$

$$v2 = \boxed{v1 \% 4}$$

$$b' = b - v2$$

$v1$	0	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?
2^{**K}	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0
$2^{**K} - 1$	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1

Prolog + epilogue loop

- If $s == 0$
for (i = 0; i < b; ++i)

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$b' = b - (b - s) \% 4$$

Prolog + epilogue loop

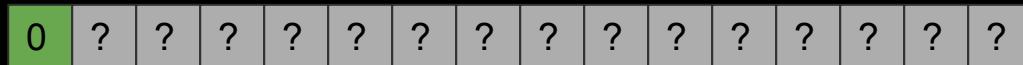
- If $s == 0$
for ($i = 0$; $i < b$; $++i$)
- If $b \geq 0$

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$b' = b - (b - s) \% 4$$

b



Prolog + epilogue loop

- If $s == 0$
`for (i = 0; i < b; ++i)`
- If $b \geq 0$
- If `UNROLLED_STEP == 2**K`

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$b' = b - (b - s) \% 4$$

b	0	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?
2^{**K}	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0

Prolog + epilogue loop

- If $s == 0$
- for ($i = 0; i < b; ++i$)
- If $b \geq 0$
- If $\text{UNROLLED_STEP} == 2^{**K}$
- $b' = b - b \% 2^{**K}$

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$b' = b - (b - s) \% 4$$

b	0	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?
2^{**K}	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0
b'	0	?	?	?	?	?	?	?	?	?	0	0	0	0	0	0

Prolog + epilogue loop

- If $s == 0$
- for ($i = 0; i < b; ++i$)
- If $b \geq 0$
- If $\text{UNROLLED_STEP} == 2^{**K}$
- $b' = b - b \% 2^{**K}$

$$b - b \% 2^{**K} == b \& \sim(2^{**K} - 1)$$

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$b' = b - (b - s) \% 4$$

b	0	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?
2^{**K}	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0
b'	0	?	?	?	?	?	?	?	?	?	0	0	0	0	0	0	0
$\sim(2^{**K} - 1)$	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0

Prolog + epilogue loop

- If $s == 0$
- for ($i = 0; i < b; ++i$)
- If $b \geq 0$
- If $\text{UNROLLED_STEP} == 2^{**K}$
- $b' = b - b \% 2^{**K}$

$$b - b \% 2^{**K} == b \& \sim(2^{**K} - 1)$$

Overhead:

- Epilogue (remainder) loop
- Calculate b'

$$b' = b - (b - s) \% 4$$

b	0	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?	?
2^{**K}	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0
b'	0	?	?	?	?	?	?	?	?	?	0	0	0	0	0	0	0
$\sim(2^{**K} - 1)$	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0

single
instruction
prologue!

Outline of the talk

- Loop from compiler perspective
- Loop unrolling basics
- Loop unrolling overhead
- GCC and CLANG unroll details
- New optimization opportunities after unrolling
- Example

origin loop

```
for (i = s; i < b; i++)
    statements(i);
```

origin loop

unroll +
runtime remainder

```
for (i = s; i < b; i++)  
    statements(i);
```

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

origin loop

```
for (i = s; i < b; i++)  
    statements(i);
```

unroll + runtime remainder

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

unroll + no remainder guarantee

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

origin loop

```
for (i = s; i < b; i++)  
    statements(i);
```

unroll + runtime remainder

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Case: $b - s == 100$
Overhead:

unroll + no remainder guarantee

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

origin loop

```
for (i = s; i < b; i++)  
    statements(i);
```

unroll + runtime remainder

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

unroll + no remainder guarantee

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

Case: $b - s == 100$
Overhead:

```
N_jumps = 99  
N_checks = 101
```

origin loop

```
for (i = s; i < b; i++)  
    statements(i);
```

unroll + runtime remainder

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Case: $b - s == 100$
Overhead:

```
N_jumps = 99  
N_checks = 101
```

unroll + no remainder guarantee

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

```
N_jumps = 25  
N_checks = 27  
extra prologue
```

origin loop

```
for (i = s; i < b; i++)  
    statements(i);
```

unroll + runtime remainder

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Case: $b - s == 100$
Overhead:

```
N_jumps = 99  
N_checks = 101
```

unroll + no remainder guarantee

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

```
N_jumps = 25  
N_checks = 27  
extra prologue
```

```
N_jumps = 24  
N_checks = 26
```

origin loop

```
for (i = s; i < b; i++)  
    statements(i);
```

unroll + runtime remainder

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Case: $b - s == 100$
Overhead:

N_jumps = 99
N_checks = 101

unroll + no remainder guarantee

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

N_jumps = 25
N_checks = 27
extra prologue



N_jumps = 24
N_checks = 26



origin loop

```
for (i = s; i < b; i++)  
    statements(i);
```

unroll + runtime remainder

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Case: $b - s == 4$

Overhead:

unroll + no remainder guarantee

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

origin loop

```
for (i = s; i < b; i++)  
    statements(i);
```

unroll + runtime remainder

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Case: $b - s == 4$

Overhead:

```
N_jumps = 3  
N_checks = 5
```

unroll + no remainder guarantee

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

origin loop

```
for (i = s; i < b; i++)  
    statements(i);
```

unroll + runtime remainder

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Case: $b - s == 4$

Overhead:

```
N_jumps = 3  
N_checks = 5
```

unroll + no remainder guarantee

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

```
N_jumps = 1  
N_checks = 4  
extra prologue
```

origin loop

```
for (i = s; i < b; i++)  
    statements(i);
```

unroll + runtime remainder

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Case: $b - s == 4$

Overhead:

```
N_jumps = 3  
N_checks = 5
```

unroll + no remainder guarantee

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

```
N_jumps = 1  
N_checks = 4  
extra prologue
```

```
N_jumps = 0  
N_checks = 2
```

origin loop

```
for (i = s; i < b; i++)  
    statements(i);
```

unroll + runtime remainder

```
b' = b - (b - s) % 4;  
  
for (i = s; i < b'; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}  
  
for (i = b'; i < b; i++)  
    statements(i);
```

Case: $b - s == 4$

Overhead:

N_jumps = 3
N_checks = 5



unroll + no remainder guarantee

```
for (i = s; i < b; i += 4) {  
    statements(i);  
    statements(i + 1);  
    statements(i + 2);  
    statements(i + 3);  
}
```

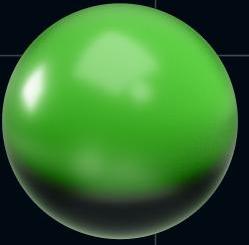
N_jumps = 1
N_checks = 4
extra prologue



N_jumps = 0
N_checks = 2

Outline of the talk

- Loop from compiler perspective
- Loop unrolling basics
- Loop unrolling overhead
- GCC and CLANG unroll details
- New optimization opportunities after unrolling
- Example



Outline of the talk

- Loop from compiler perspective
- Loop unrolling basics
- Loop unrolling overhead
- GCC and CLANG unroll details
- New optimization opportunities after unrolling
- Example

GCC 12.2

-O2
-march=skylake
-fno-tree-vectorize

CLANG 15.0

-O2
-march=skylake
-fno-vectorize
-fno-slp-vectorize

Examples. Loop (GCC vs CLANG)

```
1 int sum(int *a, int n) {  
2     int x = 0;  
3     for (int i = 0; i < n; ++i)  
4         x += a[i];  
5     return x;  
6 }  
7
```

Examples. Loop (GCC vs CLANG)

x86-64 clang 15.0.0 -O2 -march=skylake -fno-vectorize -fno-sse4

```
1 sum(int*, int): # @sum(int*, int)
2     test    esi, esi
3     jle     .LBB0_1
4     mov     esi, esi
5     lea     rax, [rsi - 1]
6     mov     ecx, esi
7     and    ecx, 7
8     cmp     rax, 7
9     jae     .LBB0_8
10    xor    edx, edx
11    xor    eax, eax
12    jmp     .LBB0_4
13 .LBB0_1:
14    xor    eax, eax
15    ret
16 .LBB0_8:
17    and    esi, -8
18    xor    edx, edx
19    xor    eax, eax
20 .LBB0_9:           # =>This Inner Loop Header: Depth=1
21    add    eax, dword ptr [rdi + 4*rdx]
22    add    eax, dword ptr [rdi + 4*rdx + 4]
23    add    eax, dword ptr [rdi + 4*rdx + 8]
24    add    eax, dword ptr [rdi + 4*rdx + 12]
25    add    eax, dword ptr [rdi + 4*rdx + 16]
26    add    eax, dword ptr [rdi + 4*rdx + 20]
27    add    eax, dword ptr [rdi + 4*rdx + 24]
28    add    eax, dword ptr [rdi + 4*rdx + 28]
29    add    rdx, 8
30    cmp    rsi, rdx
31    jne     .LBB0_9
32 .LBB0_4:
33    test   rcx, rcx
34    je      .LBB0_7
35    lea     rdx, [rdi + 4*rdx]
36    xor    esi, esi
37 .LBB0_6:           # =>This Inner Loop Header: Depth=1
38    add    eax, dword ptr [rdx + 4*rsi]
39    inc    rsi
40    cmp    rcx, rsi
41    jne     .LBB0_6
42 .LBB0_7:
43    ret
```

Examples. Loop (GCC vs CLANG)

x86-64 clang 15.0.0 -O2 -march=skylake -fno-vectorize -fno-sse4.2

```
1 sum(int*, int): # @sum(int*, int)
2     test    esi, esi
3     jle     .LBB0_1
4     mov     esi, esi
5     lea     rax, [rsi - 1]
6     mov     ecx, esi
7     and    ecx, 7
8     cmp     rax, 7
9     jae     .LBB0_8
10    xor    edx, edx
11    xor    eax, eax
12    jmp     .LBB0_4
13 .LBB0_1:
14    xor    eax, eax
15    ret
16 .LBB0_8:
17    and    esi, -8
18    xor    edx, edx
19    xor    eax, eax
20 .LBB0_9:           # =>This Inner Loop Header: Depth=1
21    add    eax, dword ptr [rdi + 4*rdx]
22    add    eax, dword ptr [rdi + 4*rdx + 4]
23    add    eax, dword ptr [rdi + 4*rdx + 8]
24    add    eax, dword ptr [rdi + 4*rdx + 12]
25    add    eax, dword ptr [rdi + 4*rdx + 16]
26    add    eax, dword ptr [rdi + 4*rdx + 20]
27    add    eax, dword ptr [rdi + 4*rdx + 24]
28    add    eax, dword ptr [rdi + 4*rdx + 28]
29    add    rdx, 8
30    cmp    rsi, rdx
31    jne     .LBB0_9
32 .LBB0_4:
33    test   rcx, rcx
34    je     .LBB0_7
35    lea     rdx, [rdi + 4*rdx]
36    xor    esi, esi
37 .LBB0_6:           # =>This Inner Loop Header: Depth=1
38    add    eax, dword ptr [rdx + 4*rsi]
39    inc    rsi
40    cmp    rcx, rsi
41    jne     .LBB0_6
42 .LBB0_7:
43    ret
```

loop guard

loop guard

Examples. Loop (GCC vs CLANG)

x86-64 clang 15.0.0 -O2 -march=skylake -fno-vectorize -fno-sse4.2

```
1 sum(int*, int): # @sum(int*, int)
2     test    esi, esi
3     jle     .LBB0_1
4     mov     esi, esi
5     lea     rax, [rsi - 1]
6     mov     ecx, esi
7     and    ecx, 7
8     cmp     rax, 7
9     jae     .LBB0_8
10    xor    edx, edx
11    xor    eax, eax
12    jmp     .LBB0_4
13 .LBB0_1:
14    xor    eax, eax
15    ret
16 .LBB0_8:
17    and    esi, -8
18    xor    edx, edx
19    xor    eax, eax
20 .LBB0_9: # =>This Inner Loop Header: Depth=1
21    add    eax, dword ptr [rdi + 4*rdx]
22    add    eax, dword ptr [rdi + 4*rdx + 4]
23    add    eax, dword ptr [rdi + 4*rdx + 8]
24    add    eax, dword ptr [rdi + 4*rdx + 12]
25    add    eax, dword ptr [rdi + 4*rdx + 16]
26    add    eax, dword ptr [rdi + 4*rdx + 20]
27    add    eax, dword ptr [rdi + 4*rdx + 24]
28    add    eax, dword ptr [rdi + 4*rdx + 28]
29    add    rdx, 8
30    cmp    rsi, rdx
31    jne     .LBB0_9
32 .LBB0_4:
33    test   rcx, rcx
34    je     .LBB0_7
35    lea     rdx, [rdi + 4*rdx]
36    xor    esi, esi
37 .LBB0_6: # =>This Inner Loop Header: Depth=1
38    add    eax, dword ptr [rdx + 4*rsi] p1
39    inc    rsi
40    cmp    rcx, rsi
41    jne     .LBB0_6
42 .LBB0_7:
43    ret
```

loop guard

p1 x8

loop guard

65

Examples. Loop (GCC vs CLANG)

x86-64 clang 15.0.0 -O2 -march=skylake -fno-vectorize -fno-sse4.2

```
1 sum(int*, int):          # @sum(int*, int)
2     test    esi, esi
3     jle     .LBB0_1
4     mov     esi, esi
5     lea     rax, [rsi - 1]
6     mov     ecx, esi
7     and    ecx, 7
8     cmp     rax, 7
9     jae     .LBB0_8
10    xor    edx, edx
11    xor    eax, eax
12    jmp     .LBB0_4
13 .LBB0_1:
14    xor    eax, eax
15    ret
16 .LBB0_8:
17    and    esi, -8
18    xor    edx, edx
19    xor    eax, eax
20 .LBB0_9:                 # =>This Inner Loop Header: Depth=1
21    add    eax, dword ptr [rdi + 4*rdx]
22    add    eax, dword ptr [rdi + 4*rdx + 4]
23    add    eax, dword ptr [rdi + 4*rdx + 8]
24    add    eax, dword ptr [rdi + 4*rdx + 12]
25    add    eax, dword ptr [rdi + 4*rdx + 16]
26    add    eax, dword ptr [rdi + 4*rdx + 20]
27    add    eax, dword ptr [rdi + 4*rdx + 24]
28    add    eax, dword ptr [rdi + 4*rdx + 28]
29    add    rdx, 8
30    cmp    rsi, rdx
31    jne     .LBB0_9
32 .LBB0_4:
33    test   rcx, rcx
34    je     .LBB0_7
35    lea     rdx, [rdi + 4*rdx]
36    xor    esi, esi
37 .LBB0_6:                 # =>This Inner Loop Header: Depth=1
38    add    eax, dword ptr [rdx + 4*rsi]
39    inc    rsi
40    cmp    rcx, rsi
41    jne     .LBB0_6
42 .LBB0_7:
43    ret
```

loop guard

pl x8

iter x8
check + be

loop guard

pl
iter
check + be

Examples. Loop (GCC vs CLANG)

x86-64 clang 15.0.0 -O2 -march=skylake -fno-vectorize -fno-sse4.2

```
1 sum(int*, int):          # @sum(int*, int)
2     test    esi, esi
3     jle     .LBB0_1
4     mov     esi, esi
5     lea     rax, [rsi - 1]
6     mov     ecx, esi
7     and    ecx, 7
8     cmp     rax, 7
9     jae     .LBB0_8
10    xor    edx, edx
11    xor    eax, eax
12    jmp     .LBB0_4
13 .LBB0_1:
14    xor    eax, eax
15    ret
16 .LBB0_8:
17    and    esi, -8
18    xor    edx, edx
19    xor    eax, eax
20 .LBB0_9:                 # =>This Inner Loop Header: Depth=1
21    add    eax, dword ptr [rdi + 4*rdx]
22    add    eax, dword ptr [rdi + 4*rdx + 4]
23    add    eax, dword ptr [rdi + 4*rdx + 8]
24    add    eax, dword ptr [rdi + 4*rdx + 12]
25    add    eax, dword ptr [rdi + 4*rdx + 16]
26    add    eax, dword ptr [rdi + 4*rdx + 20]
27    add    eax, dword ptr [rdi + 4*rdx + 24]
28    add    eax, dword ptr [rdi + 4*rdx + 28]
29    add    rdx, 8
30    cmp    rsi, rdx
31    jne     .LBB0_9
32 .LBB0_4:
33    test   rcx, rcx
34    je     .LBB0_7
35    lea     rdx, [rdi + 4*rdx]
36    xor    esi, esi
37 .LBB0_6:                 # =>This Inner Loop Header: Depth=1
38    add    eax, dword ptr [rdx + 4*rsi]
39    inc    rsi
40    cmp    rcx, rsi
41    jne     .LBB0_6
42 .LBB0_7:
43    ret
```

loop guard

prologue

pl x8

iter x8
check + be

loop guard
prologue

pl
iter
check + be

Examples. Loop (GCC vs CLANG)

x86-64 clang 15.0.0 -O2 -march=skylake -fno-vectorize -fno-sse4.2

```
1 sum(int*, int): # @sum(int*, int)
2     test    esi, esi
3     jle     .LBB0_1
4     mov     esi, esi
5     lea     rax, [rsi - 1]
6     mov     ecx, esi
7     and    ecx, 7
8     cmp     rax, 7
9     jae     .LBB0_8
10    xor    edx, edx
11    xor    eax, eax
12    jmp     .LBB0_4
13 .LBB0_1:
14    xor    eax, eax
15    ret
16 .LBB0_8:
17    and    esi, -8
18    xor    edx, edx
19    xor    eax, eax
20 .LBB0_9: # => This is Inner Loop Header: Depth=1
21    add    eax, dword ptr [rdi + 4*rdx]
22    add    eax, dword ptr [rdi + 4*rdx + 4]
23    add    eax, dword ptr [rdi + 4*rdx + 8]
24    add    eax, dword ptr [rdi + 4*rdx + 12]
25    add    eax, dword ptr [rdi + 4*rdx + 16]
26    add    eax, dword ptr [rdi + 4*rdx + 20]
27    add    eax, dword ptr [rdi + 4*rdx + 24]
28    add    eax, dword ptr [rdi + 4*rdx + 28]
29    add    rdx, 8
30    cmp    rsi, rdx
31    jne     .LBB0_9
32 .LBB0_4:
33    test   rcx, rcx
34    je     .LBB0_7
35    lea     rdx, [rdi + 4*rdx]
36    xor    esi, esi
37 .LBB0_6: # => This is Inner Loop Header: Depth=1
38    add    eax, dword ptr [rdx + 4*rsi]
39    inc    rsi
40    cmp    rcx, rsi
41    jne     .LBB0_6
42 .LBB0_7:
43    ret
```

loop guard
extra prologue
(%, etc)
prologue
pl x8
iter x8
check + be
loop guard
prologue
pl
iter
check + be

Examples. Loop (GCC vs CLANG)

What about GCC ?

Examples. Loop (GCC vs CLANG)

The image shows two assembly code editors side-by-side, comparing the output of GCC 12.2 and clang 15.0 for a simple sum function. The clang output is annotated with red boxes and text labels to highlight various loop constructs.

GCC 12.2 Assembly:

```
1 sum(int*, int):
2     test    esi, esi
3     jle     .L4
4     movsx   rsi, esi
5     lea     rdx, [rdi+rsi*4]
6     xor    eax, eax
7 .L3:
8     add    eax, DWORD PTR [rdi]
9     add    rdi, 4
10    cmp    rdi, rdx
11    jne    .L3
12    ret
13 .L4:
14    xor    eax, eax
15    ret
```

clang 15.0 Assembly:

```
1 sum(int*, int):
2     test    esi, esi
3     jle     .LBB0_1
4     mov    esi, esi
5     lea    rax, [rsi - 1]
6     mov    ecx, esi
7     and    ecx, 7
8     cmp    rax, 7
9     jae    .LBB0_8
10    xor    edx, edx
11    xor    eax, eax
12    jmp    .LBB0_4
13 .LBB0_1:
14    xor    eax, eax
15    ret
16 .LBB0_8:
17    and    esi, -8
18    xor    edx, edx
19    xor    eax, eax
20 .LBB0_9:           # => This Inner Loop Header: Depth=1
21    add    eax, dword ptr [rdi + 4*rdx]
22    add    eax, dword ptr [rdi + 4*rdx + 4]
23    add    eax, dword ptr [rdi + 4*rdx + 8]
24    add    eax, dword ptr [rdi + 4*rdx + 12]
25    add    eax, dword ptr [rdi + 4*rdx + 16]
26    add    eax, dword ptr [rdi + 4*rdx + 20]
27    add    eax, dword ptr [rdi + 4*rdx + 24]
28    add    eax, dword ptr [rdi + 4*rdx + 28]
29    add    rdx, 8
30    cmp    rsi, rdx
31    jne    .LBB0_9
32 .LBB0_4:
33    test   rcx, rcx
34    je     .LBB0_7
35    lea    rdx, [rdi + 4*rdx]
36    xor    esi, esi
37 .LBB0_6:           # => This Inner Loop Header: Depth=1
38    add    eax, dword ptr [rdx + 4*rsi]
39    inc    rsi
40    cmp    rcx, rsi
41    jne    .LBB0_6
42 .LBB0_7:
43    ret
```

Annotated clang assembly:

- loop guard**: `jle .LBB0_1`
- extra prologue**: `(%, etc)`
- prologue**: `and esi, -8`
- pl x8**: `add eax, dword ptr [rdi + 4*rdx]` to `add eax, dword ptr [rdi + 4*rdx + 28]`
- iter x8 check + be**: `add rdx, 8`, `cmp rsi, rdx`, `jne .LBB0_9`
- loop guard prologue**: `test rcx, rcx`, `je .LBB0_7`
- pl iter check + be**: `add eax, dword ptr [rdx + 4*rsi]` to `add eax, dword ptr [rdx + 4*rsi]`

Examples. Loop (GCC vs CLANG)

The image shows two assembly code editors side-by-side, comparing the output of GCC 12.2 and clang 15.0 for a simple sum function. The clang output is annotated with red boxes and text labels to highlight various loop constructs.

GCC 12.2 Assembly:

```
1 sum(int*, int):
2     test    esi, esi
3     jle     .L4
4     movsx   rsi, esi
5     lea     rdx, [rdi+rsi*4]
6     xor    eax, eax
7 .L3:
8     add    eax, DWORD PTR [rdi]
9     add    rdi, 4
10    cmp    rdi, rdx
11    jne     .L3
12    ret
13 .L4:
14    xor    eax, eax
15    ret
```

clang 15.0 Assembly:

```
1 sum(int*, int):
2     test    esi, esi
3     jle     .LBB0_1
4     mov    esi, esi
5     lea    rax, [rsi - 1]
6     mov    ecx, esi
7     and    ecx, 7
8     cmp    rax, 7
9     jae     .LBB0_8
10    xor    edx, edx
11    xor    eax, eax
12    jmp     .LBB0_4
13 .LBB0_1:
14    xor    eax, eax
15    ret
16 .LBB0_8:
17    and    esi, -8
18    xor    edx, edx
19    xor    eax, eax
20 .LBB0_9:           # => This Inner Loop Header: Depth=1
21    add    eax, dword ptr [rdi + 4*rdx]
22    add    eax, dword ptr [rdi + 4*rdx + 4]
23    add    eax, dword ptr [rdi + 4*rdx + 8]
24    add    eax, dword ptr [rdi + 4*rdx + 12]
25    add    eax, dword ptr [rdi + 4*rdx + 16]
26    add    eax, dword ptr [rdi + 4*rdx + 20]
27    add    eax, dword ptr [rdi + 4*rdx + 24]
28    add    eax, dword ptr [rdi + 4*rdx + 28]
29    add    rdx, 8
30    cmp    rsi, rdx
31    jne     .LBB0_9
32 .LBB0_4:
33    test   rcx, rcx
34    je     .LBB0_7
35    lea    rdx, [rdi + 4*rdx]
36    xor    esi, esi
37 .LBB0_6:           # => This Inner Loop Header: Depth=1
38    add    eax, dword ptr [rdx + 4*rsi]
39    inc    rsi
40    cmp    rcx, rsi
41    jne     .LBB0_6
42 .LBB0_7:
43    ret
```

Annotations for clang assembly:

- loop guard**: Points to the `jle .LBB0_1` instruction.
- extra prologue**: Points to the `and esi, -8` instruction.
- (%, etc)**: Points to the `and esi, -8` instruction.
- prologue**: Points to the `test rcx, rcx` instruction.
- pl x8**: Points to the first `add eax, ...` instruction in the inner loop header.
- iter x8 check + be**: Points to the `add rdx, 8` instruction.
- loop guard prologue**: Points to the `je .LBB0_7` instruction.
- pl iter**: Points to the `add eax, ...` instruction in the inner loop header.
- iter check + be**: Points to the `inc rsi` instruction.

Text in the center:

Where is
unrolling ?!

GCC vs CLANG (-O2)

GCC:

- using PGO data or pragma

Clang:

- auto detect unroll and unroll count

GCC vs CLANG (-O2)

```
1 int sum(int *a, int n) {  
2     int x = 0;  
3     #pragma GCC unroll 4  
4     for (int i = 0; i < n; ++i)  
5         x += a[i];  
6     return x;  
7 }  
8
```

force GCC to
unroll loop

GCC vs CLANG (-O2)

x86-64 gcc 12.2 -O2 -march=skylake -fno-tree-vec

```
1 sum(int*, int):
2     test    esi, esi
3     jle     .L4
4     movsx   rsi, esi
5     lea     rdx, [-4*rsi*4]
6     shr     rdx, 2
7     inc     rdx
8     lea     rcx, [rdi+rsi*4]
9     xor     eax, eax
10    and    edx, 3
11    je     .L3
12    cmp    rdx, 1
13    je     .L15
14    cmp    rdx, 2
15    je     .L16
16    mov    eax, DWORD PTR [rdi]
17    add    rdi, 4
18 .L16:
19    add    eax, DWORD PTR [rdi]
20    add    rdi, 4
21 .L15:
22    add    eax, DWORD PTR [rdi]
23    add    rdi, 4
24    cmp    rdi, rcx
25    je     .L22
26 .L3:
27    add    eax, DWORD PTR [rdi]
28    add    eax, DWORD PTR [rdi+4]
29    add    eax, DWORD PTR [rdi+8]
30    add    eax, DWORD PTR [rdi+12]
31    add    rdi, 16
32    cmp    rdi, rcx
33    jne   .L3
34    ret
35 .L22:
36    ret
37 .L4:
38    xor    eax, eax
39    ret
```

GCC vs CLANG (-O2)

x86-64 gcc 12.2 -O2 -march=skylake -fno-tree-vec

```
1 sum(int*, int):
2     test    esi, esi
3     jle     .L4
4     movsx   rsi, esi
5     lea     rdx, [-4+rsi*4]
6     shr     rdx, 2
7     inc     rdx
8     lea     rcx, [rdi+rsi*4]
9     xor     eax, eax
10    and    edx, 3
11    je     .L3
12    cmp    rdx, 1
13    je     .L15
14    cmp    rdx, 2
15    je     .L16
16    mov    eax, DWORD PTR [rdi]
17    add    rdi, 4
.L16:
18    add    eax, DWORD PTR [rdi]
19    add    rdi, 4
20
.L15:
21    add    eax, DWORD PTR [rdi]
22    add    rdi, 4
23    cmp    rdi, rcx
24    je     .L22
25
.L3:
26    add    eax, DWORD PTR [rdi]
27    add    eax, DWORD PTR [rdi+4]
28    add    eax, DWORD PTR [rdi+8]
29    add    eax, DWORD PTR [rdi+12]
30
31    add    rdi, 16
32    cmp    rdi, rcx
33    jne   .L3
34    ret
35
.L22:
36    ret
37
.L4:
38    xor    eax, eax
39    ret
```

x4

GCC vs CLANG (-O2)

GCC:

- using PGO data or pragma
- trick with jumps and body cloning

Clang:

- auto detect unroll and unroll count
- “naive”

GCC vs CLANG (-O2)

```
C++  
int sum(int *a, int n) {  
    int x = 0;  
    #pragma GCC unroll 5  
    for (int i = 0; i < n; ++i)  
        x += a[i];  
    return x;  
}
```

force to
unroll 5

x86-64 gcc 12.2 -O2 -march=skylake -fno-vector

```
14     cmp    rdx, 2  
15     je     .L16  
16     mov    eax, DWORD PTR [rdi]  
17     add    rdi, 4  
.L16:  
19     add    eax, DWORD PTR [rdi]  
20     add    rdi, 4  
.L15:  
22     add    eax, DWORD PTR [rdi]  
23     add    rdi, 4  
24     cmp    rdi, rcx  
25     je     .L22  
.L3:  
27     add    eax, DWORD PTR [rdi]  
28     add    eax, DWORD PTR [rdi+4]  
29     add    eax, DWORD PTR [rdi+8]  
30     add    eax, DWORD PTR [rdi+12]  
31     add    rdi, 16  
32     cmp    rdi, rcx  
33     jne    .L3      x4 !  
34     ret  
.L22:  
36     ret  
.L4:  
38     xor    eax, eax  
39     ret
```

x86-64 clang 15.0.0 -O2 -march=skylake -fno-vector

```
18     cmp    eax, 4  
19     jae    .LBB0_8  
20     xor    edx, edx  
21     xor    eax, eax  
22     jmp    .LBB0_4  
.LBB0_1:  
23     xor    eax, eax  
24     xor    eax, eax  
25     ret  
.LBB0_8:  
27     sub    esi, ecx  
28     xor    edx, edx  
29     xor    eax, eax  
30     .LBB0_9:          # =>This Inr  
31     add    eax, dword ptr [rdi + 4*rdx]  
32     add    eax, dword ptr [rdi + 4*rdx + 4]  
33     add    eax, dword ptr [rdi + 4*rdx + 8]  
34     add    eax, dword ptr [rdi + 4*rdx + 12]  
35     add    eax, dword ptr [rdi + 4*rdx + 16]  
36     add    rdx, 5  
37     cmp    esi, edx  
38     jne    .LBB0_9      x5  
.LBB0_4:  
40     test   ecx, ecx  
41     je     .LBB0_7  
42     lea    rdx, [rdi + 4*rdx]  
43     mov    ecx, ecx  
44     xor    esi, esi  
45     .LBB0_6:          # =>This Inr  
46     add    eax, dword ptr [rdx + 4*rsi]  
47     inc    rsi
```

GCC vs CLANG (-O2)

```
C++  
int sum(int *a, int n) {  
    int x = 0;  
    #pragma GCC unroll 5  
    for (int i = 0; i < n; ++i)  
        x += a[i];  
    return x;  
}
```

force to
unroll 5

No GCC
warning!

x86-64 gcc 12.2 -O2 -march=skylake -fno-vector

```
14     cmp    rdx, 2  
15     je     .L16  
16     mov    eax, DWORD PTR [rdi]  
17     add    rdi, 4  
.L16:  
19     add    eax, DWORD PTR [rdi]  
20     add    rdi, 4  
.L15:  
22     add    eax, DWORD PTR [rdi]  
23     add    rdi, 4  
24     cmp    rdi, rcx  
25     je     .L22  
.L3:  
27     add    eax, DWORD PTR [rdi]  
28     add    eax, DWORD PTR [rdi+4]  
29     add    eax, DWORD PTR [rdi+8]  
30     add    eax, DWORD PTR [rdi+12]  
31     add    rdi, 16  
32     cmp    rdi, rcx  
33     jne    .L3           x4 !  
34     ret  
.L22:  
36     ret  
.L4:  
38     xor    eax, eax  
39     ret
```

x86-64 clang 15.0.0 -O2 -march=skylake -fno-vector

```
18     cmp    eax, 4  
19     jae    .LBB0_8  
20     xor    edx, edx  
21     xor    eax, eax  
22     jmp    .LBB0_4  
.LBB0_1:  
23     xor    eax, eax  
24     xor    eax, eax  
25     ret  
.LBB0_8:  
27     sub    esi, ecx  
28     xor    edx, edx  
29     xor    eax, eax  
30     .LBB0_9:          # =>This Inr  
31     add    eax, dword ptr [rdi + 4*rdx]  
32     add    eax, dword ptr [rdi + 4*rdx + 4]  
33     add    eax, dword ptr [rdi + 4*rdx + 8]  
34     add    eax, dword ptr [rdi + 4*rdx + 12]  
35     add    eax, dword ptr [rdi + 4*rdx + 16]  
36     add    rdx, 5  
37     cmp    esi, edx  
38     jne    .LBB0_9           x5  
.LBB0_4:  
40     test   ecx, ecx  
41     je     .LBB0_7  
42     lea    rdx, [rdi + 4*rdx]  
43     mov    ecx, ecx  
44     xor    esi, esi  
45     .LBB0_6:          # =>This Inr  
46     add    eax, dword ptr [rdx + 4*rsi]  
47     inc    rsi
```

GCC vs CLANG (-O2)

The screenshot shows three panes of assembly code side-by-side, comparing the output of GCC 12.2 and Clang 15.0.0 with optimization level -O2, targeting Skylake architecture.

Left Pane (GCC Output):

```
1 int sum(int *a, int n) {
2     int x = 0;
3     #pragma GCC unroll 4
4     for (int i = 0; i < n; i += 3)
5         x += a[i];
6     return x;
7 }
```

Middle Pane (GCC Assembly):

```
1 sum(int*, int):
2     test    esi, esi
3     jle     .L4
4     xor    eax, eax
5     xor    edx, edx
6 .L3:
7     lea     rcx, [rax+3]
8     add    edx, DWORD PTR [rdi+rcx*4]
9     cmp    esi, ecx
10    jle    .L1
11    lea     rax, [rcx+3]
12    add    edx, DWORD PTR [rdi+rcx*4]
13    cmp    esi, eax
14    jle    .L1
15    add    edx, DWORD PTR [rdi+rcx*4]
16    lea     rax, [rcx+6]
17    cmp    esi, eax
18    jle    .L1
19    add    edx, DWORD PTR [rdi+rcx*4]
20    lea     rax, [rcx+9]
21    cmp    esi, eax
22    jg     .L3
23 .L1:
24     mov    eax, edx
25     ret
26 .L4:
27     xor    edx, edx
28     mov    eax, edx
29     ret
```

Right Pane (Clang Assembly):

```
16     ret
17 .LBB0_8:
18     mov    edx, r8d
19     and    edx, 2147483644
20     xor    esi, esi
21     xor    eax, eax
22 .LBB0_9:                                # =>This Inr
23     add    eax, dword ptr [rdi + 4*rsi]
24     add    eax, dword ptr [rdi + 4*rsi + 12]
25     add    eax, dword ptr [rdi + 4*rsi + 24]
26     add    eax, dword ptr [rdi + 4*rsi + 36]
27     add    rsi, 12
28     add    edx, -4
29     jne    .LBB0_9
30 .LBB0_4:
31     test   r8b, 3
32     je     .LBB0_7
33     lea    rdx, [rdi + 4*rsi]
34     inc    cl
35     movzx  ecx, cl
36     and    ecx, 3
37     shl    rcx, 2
38     lea    rcx, [rcx + 2*rcx]
39     xor    esi, esi
40 .LBB0_6:                                # =>This Inr
41     add    eax, dword ptr [rdx + rsi]
42     add    rsi, 12
43     cmp    ecx, esi
44     jne    .LBB0_6
45 .LBB0_7:
```

Annotations:

- A red box highlights the loop step `i += 3` in the C code.
- The text "non-power-of 2 loop step" is overlaid in red on the left side of the image.

GCC vs CLANG (-O2)

non-power-of 2
loop step

Line Number	Code (GCC)	Code (CLANG)
1	int sum(int *a, int n) {	
2	int x = 0;	
3	#pragma GCC unroll 4	
4	for (int i = 0; i < n; i += 3)	test esi, esi jle .L4 xor eax, eax xor edx, edx
5	x += a[i];	.L3:
6	return x;	lea rcx, [rax+3] add edx, DWORD PTR [rdi+rax*4] cmp esi, ecx jle .L1 lea rax, [rcx+3] add edx, DWORD PTR [rdi+rcx*4] cmp esi, eax jle .L1 add edx, DWORD PTR [rdi+rax*4] lea rax, [rcx+6] cmp esi, eax jle .L1 add edx, DWORD PTR [rdi+rax*4] lea rax, [rcx+9] cmp esi, eax jg .L3
7	}	.L1: mov eax, edx ret
8		.L4: xor edx, edx mov eax, edx ret
		16 ret 17 .LBB0_8: 18 mov edx, r8d 19 and edx, 2147483644 20 xor esi, esi 21 xor eax, eax
		22 .LBB0_9: # =>This Inr 23 add eax, dword ptr [rdi + 4*rsi] 24 add eax, dword ptr [rdi + 4*rsi + 12] 25 add eax, dword ptr [rdi + 4*rsi + 24] 26 add eax, dword ptr [rdi + 4*rsi + 36] 27 add rsi, 12 28 add edx, -4 29 jne .LBB0_9
		30 .LBB0_4: 31 test r8b, 3 32 je .LBB0_7 33 lea rdx, [rdi + 4*rsi] 34 inc cl 35 movzx ecx, cl 36 and ecx, 3 37 shl rcx, 2 38 lea rcx, [rcx + 2*rcx] 39 xor esi, esi
		40 .LBB0_6: # =>This Inr 41 add eax, dword ptr [rdx + rsi] 42 add rsi, 12 43 cmp ecx, esi 44 jne .LBB0_6
		45 .LBB0_7:

GCC vs CLANG (-O2)

GCC:

- using PGO data or pragma
- trick with jumps and body cloning
- Power-of-2 step && unroll count
(rejection / body cloning)

Clang:

- auto detect unroll and unroll count
- “naive”
- any step / unroll count
(prologue cost)

Outline of the talk

- Loop from compiler perspective
- Loop unrolling basics
- Loop unrolling overhead
- GCC and CLANG unroll details
- New optimization opportunities after unrolling
- Example

New opportunities

- Loop vectorization
- SLP vectorization
- Full unroll / loop deletion
- (Mem2Reg) Promote Memory To Register /
(SROA) Scalar Replacement Of Aggregates
- ...

Full unroll / loop deletion

The image shows three windows of a debugger comparing assembly output from two different compilers for the same C++ code.

C++ Code:

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

Compiler 1: x86-64 gcc 12.2 -O2 -march=skylake -fno-tree-vectorize

```
1 sum(int*):
2     mov    eax, DWORD PTR [rdi+4]
3     add    eax, DWORD PTR [rdi]
4     add    eax, DWORD PTR [rdi+8]
5     add    eax, DWORD PTR [rdi+12]
6     ret
```

Compiler 2: x86-64 clang 15.0.0 -O2 -march=skylake -fno-slp-vectorize -fno-vectorize

```
1 sum(int*):
2     mov    eax, dword ptr [rdi + 4]
3     add    eax, dword ptr [rdi]
4     add    eax, dword ptr [rdi + 8]
5     add    eax, dword ptr [rdi + 12]
6     ret
```

Full unroll / loop deletion

The image shows three windows of a debugger comparing the assembly output of two compilers for the same C code. The leftmost window shows the C code:1 int sum(int *a) {
2 int x = 0;
3 for (int i = 0; i < 4; i++)
4 x += a[i];
5 return x;
6 }
7 compile time boundsThe loop iteration count (4) is highlighted with a red box. The middle window shows the assembly for x86-64 gcc 12.2, which includes a loop:

```
1 sum(int*):  
2     mov    eax, DWORD PTR [rdi+4]  
3     add    eax, DWORD PTR [rdi]  
4     add    eax, DWORD PTR [rdi+8]  
5     add    eax, DWORD PTR [rdi+12]  
6     ret
```

The rightmost window shows the assembly for x86-64 clang 15.0.0, which has unrolled the loop:

```
1 sum(int*):  
2     mov    eax, dword ptr [rdi + 4] # @sum(ir  
3     add    eax, dword ptr [rdi]  
4     add    eax, dword ptr [rdi + 8]  
5     add    eax, dword ptr [rdi + 12]  
6     ret
```

Full unroll / loop deletion

The image shows three windows of a debugger side-by-side, comparing the assembly output of two compilers for the same C++ code.

Left Window (GCC 12.2):

```
1 int sum(int *a) {
2     int x = 0;
3     for (int i = 0; i < 4; i++)
4         x += a[i];
5     return x;
6 }
7 compile time bounds
```

Middle Window (GCC 12.2 Assembly):

```
1 sum(int*):
2     mov    eax, DWORD PTR [rdi+4]
3     add    eax, DWORD PTR [rdi]
4     add    eax, DWORD PTR [rdi+8]
5     add    eax, DWORD PTR [rdi+12]
6     ret
loop is deleted
```

Right Window (Clang 15.0.0 Assembly):

```
1 sum(int*):
2     mov    eax, dword ptr [rdi + 4]
3     add    eax, dword ptr [rdi]
4     add    eax, dword ptr [rdi + 8]
5     add    eax, dword ptr [rdi + 12]
6     ret
loop is deleted
```

In the assembly code, the loop body (lines 2-5) is highlighted with a green box, and the text "loop is deleted" is displayed at the bottom of the assembly window.

Full unroll / loop deletion

The image shows three code editors side-by-side, illustrating the compilation of a simple C++ function into assembly code. The leftmost editor shows the C++ source code:

```
1 int sum(int *a) {
2     int x = 0;
3     for (int i = 0; i < 8; i++)
4         x += a[i];
5     return x;
6 }
```

The loop iteration count (8) is highlighted with a red rectangle. The middle editor shows the assembly output for x86-64 using gcc 12.2, generated with optimization flags -O2, -march=skylake, and -fno-tree-vectorize. The assembly code is:1 sum(int*):
2 lea rdx, [rdi+32]
3 xor eax, eax
4 .L2:
5 add eax, DWORD PTR [rdi]
6 add rdi, 4
7 cmp rdi, rdx
8 jne .L2
9 ret

The rightmost editor shows the assembly output for x86-64 using clang 15.0.0, also with -O2, -march=skylake, and -fno-tree-vectorize. The assembly code is:1 sum(int*):
2 mov eax, dword ptr [rdi + 4]
3 add eax, dword ptr [rdi]
4 add eax, dword ptr [rdi + 8]
5 add eax, dword ptr [rdi + 12]
6 add eax, dword ptr [rdi + 16]
7 add eax, dword ptr [rdi + 20]
8 add eax, dword ptr [rdi + 24]
9 add eax, dword ptr [rdi + 28]
10 ret

A red box highlights the assembly code for the loop iteration count (8), which corresponds to the highlighted line in the C++ source code.

compile time bounds

Full unroll / loop deletion

compile time bounds

here is the loop
(no loop guard)

loop is deleted

```
x86-64 gcc 12.2 -O2 -march=skylake -fno-tree-vectorize
```

```
1 int sum(int *a) {  
2     int x = 0;  
3     for (int i = 0; i < 8; i++)  
4         x += a[i];  
5     return x;  
6 }
```

```
x86-64 clang 15.0.0 -O2 -march=skylake -fno-slp-vectorize -fno-vectorize
```

```
1 sum(int*):  
2     lea    rdx, [rdi+32]  
3     xor    eax, eax  
4     .L2:  
5     add    eax, DWORD PTR [rdi]  
6     add    rdi, 4  
7     cmp    rdi, rdx  
8     jne    .L2  
9     ret
```

Full unroll / loop deletion

The image shows three code editors side-by-side, illustrating the assembly output for two compilers (GCC 12.2 and Clang 15.0) when given the same C++ code with a specific pragma.

Left Editor (C++ Code):

```
1 int sum(int *a) {  
2     int x = 0;  
3     #pragma GCC unroll 8  
4     for (int i = 0; i < 8; i++)  
5         x += a[i];  
6     return x;  
7 }
```

Middle Editor (GCC 12.2 Assembly):

```
1 sum(int*):  
2     mov    eax, DWORD PTR [rdi+4]  
3     add    eax, DWORD PTR [rdi]  
4     add    eax, DWORD PTR [rdi+8]  
5     add    eax, DWORD PTR [rdi+12]  
6     add    eax, DWORD PTR [rdi+16]  
7     add    eax, DWORD PTR [rdi+20]  
8     add    eax, DWORD PTR [rdi+24]  
9     add    eax, DWORD PTR [rdi+28]  
10    ret
```

Right Editor (Clang 15.0 Assembly):

```
1 sum(int*):  
2     mov    eax, dword ptr [rdi + 4]  
3     add    eax, dword ptr [rdi]  
4     add    eax, dword ptr [rdi + 8]  
5     add    eax, dword ptr [rdi + 12]  
6     add    eax, dword ptr [rdi + 16]  
7     add    eax, dword ptr [rdi + 20]  
8     add    eax, dword ptr [rdi + 24]  
9     add    eax, dword ptr [rdi + 28]  
10    ret
```

Annotations:

- The **for** loop body in the C++ code is highlighted with a red box.
- The assembly code for both compilers shows a sequence of 8 **add** instructions, each adding a memory location to the **eax** register. The first **add** instruction is at **[rdi+4]**, and subsequent ones are at **[rdi]**, **[rdi+8]**, **[rdi+12]**, **[rdi+16]**, **[rdi+20]**, **[rdi+24]**, and **[rdi+28]**.
- The assembly code for both compilers ends with a **ret** instruction.

Text at the bottom left:

compile time bounds
+ pragma hint

Full unroll / loop deletion

```
A C++  
1 int sum(int *a) {  
2     int x = 0;  
3     for (int i = 0; i < 1024; i++)  
4         x += a[i];  
5     return x;  
6 }
```

compile time bounds

```
x86-64 gcc 12.2 -O2 -march=skylake -fno-tree-vectorize  
1 sum(int*):  
2     lea    rdx, [rdi+4096]  
3     xor    eax, eax  
4     .L2:  
5     add    eax, DWORD PTR [rdi]  
6     add    rdi, 4  
7     cmp    rdi, rdx  
8     jne    .L2  
9     ret
```

no unroll
no loop guard

```
x86-64 clang 15.0.0 -O2 -march=skylake -fno-slp-vectorize -fno-vectorize  
1 sum(int*): # @sum(int*)  
2     xor    ecx, ecx  
3     xor    eax, eax  
4     .LBB0_1: # =>This Inner  
5     add    eax, dword ptr [rdi + 4*rcx]  
6     add    eax, dword ptr [rdi + 4*rcx + 4]  
7     add    eax, dword ptr [rdi + 4*rcx + 8]  
8     add    eax, dword ptr [rdi + 4*rcx + 12]  
9     add    eax, dword ptr [rdi + 4*rcx + 16]  
10    add   eax, dword ptr [rdi + 4*rcx + 20]  
11    add   eax, dword ptr [rdi + 4*rcx + 24]  
12    add   eax, dword ptr [rdi + 4*rcx + 28]  
13    add   eax, dword ptr [rdi + 4*rcx + 32]  
14    add   eax, dword ptr [rdi + 4*rcx + 36]  
15    add   eax, dword ptr [rdi + 4*rcx + 40]  
16    add   eax, dword ptr [rdi + 4*rcx + 44]  
17    add   eax, dword ptr [rdi + 4*rcx + 48]  
18    add   eax, dword ptr [rdi + 4*rcx + 52]  
19    add   eax, dword ptr [rdi + 4*rcx + 56]  
20    add   eax, dword ptr [rdi + 4*rcx + 60]  
21    add   rcx, 16  
22    cmp   rcx, 1024  
23    jne   .LBB0_1  
24    ret
```

unroll x16
no loop guard
no extra prologue

GCC vs CLANG (-O2)

GCC:

- using PGO data or pragma
- trick with jumps and body cloning
- Power-of-2 step && unroll count
(rejection / body cloning)
- Less aggressive full unroll

Clang:

- auto detect unroll and unroll count
- “naive”
- any step / unroll count
(prologue cost)
- More aggressive full unroll

Mem2Reg / SROA (scalar replacement of aggregates)

The image shows a debugger interface with two panes. The left pane displays the C++ source code for a function named `sum`. The right pane shows the generated assembly code for the same function, compiled with clang 15.0.0.

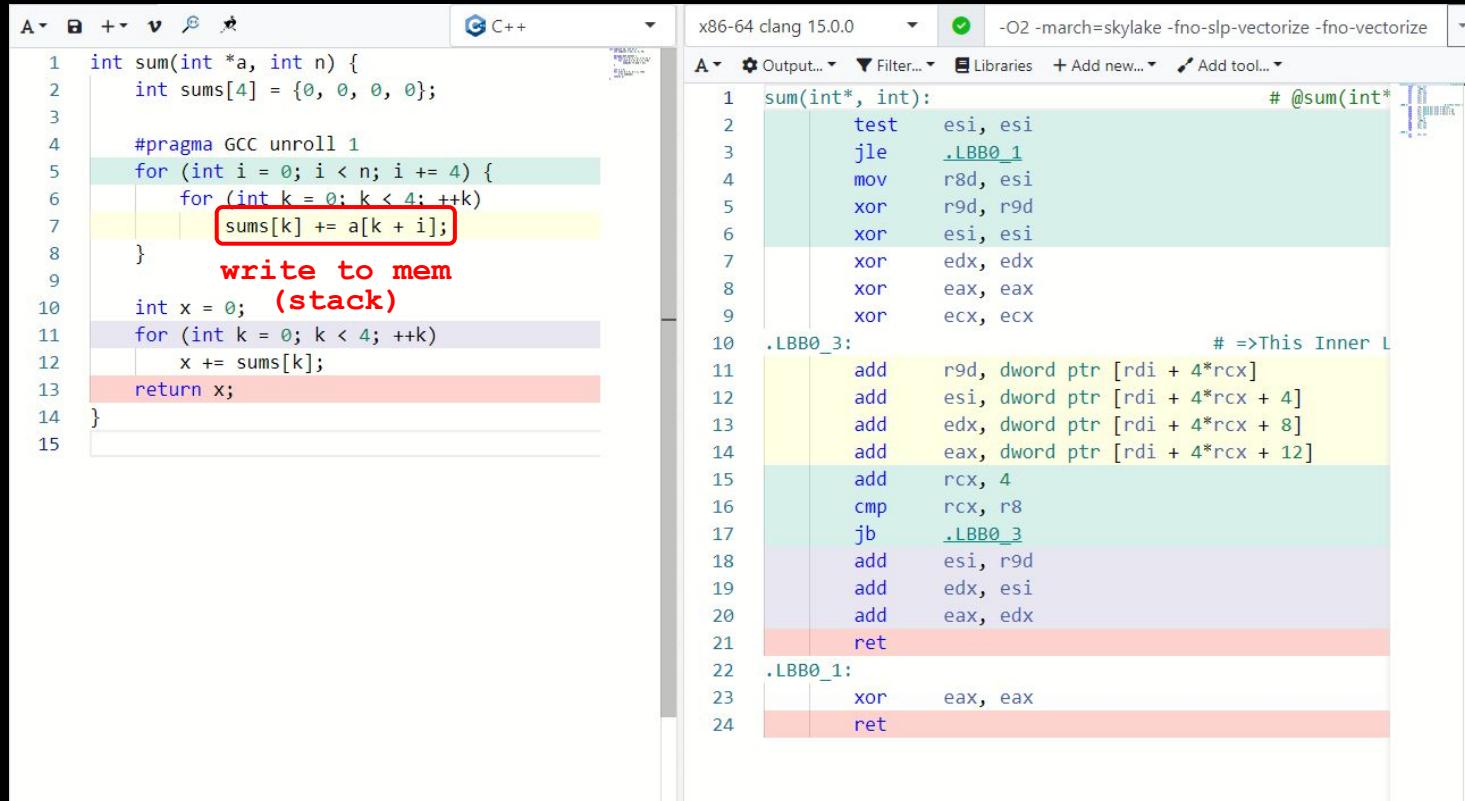
C++ Source Code:

```
1 int sum(int *a, int n) {
2     int sums[4] = {0, 0, 0, 0};
3
4     #pragma GCC unroll 1
5     for (int i = 0; i < n; i += 4) {
6         for (int k = 0; k < 4; ++k)
7             sums[k] += a[k + i];
8     }
9
10    int x = 0;
11    for (int k = 0; k < 4; ++k)
12        x += sums[k];
13    return x;
14}
15
```

Assembly Output:

```
x86-64 clang 15.0.0 -O2 -march=skylake -fno-slp-vectorize -fno-vectorize
1 sum(int*, int):                                # @sum(int*)
2     test    esi, esi
3     jle     .LBB0_1
4     mov     r8d, esi
5     xor     r9d, r9d
6     xor     esi, esi
7     xor     edx, edx
8     xor     eax, eax
9     xor     ecx, ecx
10    .LBB0_3:                                     # =>This Inner L
11        add    r9d, dword ptr [rdi + 4*rcx]
12        add    esi, dword ptr [rdi + 4*rcx + 4]
13        add    edx, dword ptr [rdi + 4*rcx + 8]
14        add    eax, dword ptr [rdi + 4*rcx + 12]
15        add    rcx, 4
16        cmp    rcx, r8
17        jb     .LBB0_3
18        add    esi, r9d
19        add    edx, esi
20        add    eax, edx
21    ret
22    .LBB0_1:
23        xor    eax, eax
24        ret
```

Mem2Reg / SROA (scalar replacement of aggregates)



The image shows the LLVM IR viewer interface with two panes. The left pane displays the original C++ code, and the right pane shows the generated assembly code.

Left Pane (C++ Code):

```
1 int sum(int *a, int n) {
2     int sums[4] = {0, 0, 0, 0};
3
4     #pragma GCC unroll 1
5     for (int i = 0; i < n; i += 4) {
6         for (int k = 0; k < 4; ++k)
7             sums[k] += a[k + i];
8     }
9     write to mem
10    int x = 0;      (stack)
11    for (int k = 0; k < 4; ++k)
12        x += sums[k];
13    return x;
14 }
15
```

Right Pane (Assembly Output):

```
x86-64 clang 15.0.0 -O2 -march=skylake -fno-slp-vectorize -fno-vectorize
1 sum(int*, int):                                # @sum(int*)
2     test    esi, esi
3     jle     .LBB0_1
4     mov     r8d, esi
5     xor     r9d, r9d
6     xor     esi, esi
7     xor     edx, edx
8     xor     eax, eax
9     xor     ecx, ecx
10    .LBB0_3:                                     # =>This Inner L
11        add    r9d, dword ptr [rdi + 4*rcx]
12        add    esi, dword ptr [rdi + 4*rcx + 4]
13        add    edx, dword ptr [rdi + 4*rcx + 8]
14        add    eax, dword ptr [rdi + 4*rcx + 12]
15        add    rcx, 4
16        cmp    rcx, r8
17        jb     .LBB0_3
18        add    esi, r9d
19        add    edx, esi
20        add    eax, edx
21    ret
22    .LBB0_1:
23        xor    eax, eax
24        ret
```

Mem2Reg / SROA (scalar replacement of aggregates)

The image shows two windows side-by-side. The left window is a C++ code editor with the following code:

```
1 int sum(int *a, int n) {
2     int sums[4] = {0, 0, 0, 0};
3
4     #pragma GCC unroll 1
5     for (int i = 0; i < n; i += 4) {
6         for (int k = 0; k < 4; ++k)
7             sums[k] += a[k + i];
8     }
9     write to mem
10    int x = 0;      (stack)
11    for (int k = 0; k < 4; ++k)
12        x += sums[k];
13    return x;
14 }
15
```

The line `sums[k] += a[k + i];` is highlighted with a red box and labeled "write to mem". The line `int x = 0;` is highlighted with a purple box and labeled "(stack)". The line `return x;` is highlighted with a pink box.

The right window is a LLVM IR viewer showing the assembly output for the same code. The assembly is:

```
x86-64 clang 15.0.0 -O2 -march=skylake -fno-slp-vectorize -fno-vectorize
1 sum(int*, int):                                # @sum(int*)
2     test    esi, esi
3     jle     .LBB0_1
4     mov     r8d, esi
5     xor     r9d, r9d
6     xor     esi, esi
7     xor     edx, edx
8     xor     eax, eax
9     xor     ecx, ecx
10    .LBB0_3:                                     # =>This Inner L
11        add    r9d, dword ptr [rdi + 4*rcx]
12        add    esi, dword ptr [rdi + 4*rcx + 4]
13        add    edx, dword ptr [rdi + 4*rcx + 8]
14        add    eax, dword ptr [rdi + 4*rcx + 12]
15        add    rcx, 4
16        cmp    rcx, r8
17        jb     .LBB0_3
18        add    esi, r9d
19        add    edx, esi
20        add    eax, edx
21        ret
22    .LBB0_1:
23        xor    eax, eax
24        ret
```

The assembly code corresponds to the C++ code, showing the scalar replacement of the aggregate `sums` array. The assembly output includes comments indicating the original C code: `# @sum(int*)`, `read`, and `# =>This Inner L`.

Mem2Reg / SROA (scalar replacement of aggregates)

The image shows two windows side-by-side. The left window is a C++ code editor with the following code:

```
1 int sum(int *a, int n) {
2     int sums[4] = {0, 0, 0, 0};
3
4     #pragma GCC unroll 1
5     for (int i = 0; i < n; i += 4) {
6         for (int k = 0; k < 4; ++k)
7             sums[k] += a[k + i];
8     }
9     write to mem
10    int x = 0;      (stack)
11    for (int k = 0; k < 4; ++k)
12        x += sums[k];
13    return x;
14 }
15
```

The code is annotated with several color-coded regions: a green box highlights the inner loop body, a yellow box highlights the `sums[k] += a[k + i];` assignment, and a red box highlights the `sums` array. Red text overlays "write to mem" and "(stack)" are placed over the annotated regions.

The right window is a LLVM IR viewer showing the assembly output for the `sum` function:

```
x86-64 clang 15.0.0 -O2 -march=skylake -fno-slp-vectorize -fno-vectorize
```

```
1 sum(int*, int):                                # @sum(int*)
2     test    esi, esi
3     jle     .LBB0_1
4     mov     r8d, esi
5     xor     r9d, r9d
6     xor     esi, esi
7     xor     edx, edx   read
8     xor     eax, eax   read
9     xor     ecx, ecx
10    .LBB0_3:                                     # =>This Inner L
11    add    r9d, dword ptr [rdi + 4*rcx]
12    add    esi, dword ptr [rdi + 4*rcx + 4]
13    add    edx, dword ptr [rdi + 4*rcx + 8]
14    add    eax, dword ptr [rdi + 4*rcx + 12]
15    add    rcx, 4
16    cmp    rcx, r8
17    jb     .LBB0_3
18    add    esi, r9d
19    add    edx, esi
20    add    eax, edx
21    ret
22    .LBB0_1:
23    xor     eax, eax
24    ret
```

The assembly code is color-coded to match the C++ code. Yellow boxes highlight the `add` instructions in the inner loop body, and red boxes highlight the `ret` and `xor` instructions at the end of the function. The `test`, `jle`, and `mov` instructions are also highlighted in yellow.

Mem2Reg / SROA (scalar replacement of aggregates)

The image shows two windows side-by-side. The left window is a C++ code editor with the following code:

```
1 int sum(int *a, int n) {
2     int sums[4] = {0, 0, 0, 0};
3
4     #pragma GCC unroll 1
5     for (int i = 0; i < n; i += 4) {
6         for (int k = 0; k < 4; ++k)
7             sums[k] += a[k + i];
8     }
9     write to mem
10    int x = 0;      (stack)
11    for (int k = 0; k < 4; ++k)
12        x += sums[k];
13    return x;
14 }
15
```

The line `sums[k] += a[k + i];` is highlighted with a red box and labeled "write to mem". The line `int x = 0;` is highlighted with a purple box and labeled "(stack)". The line `return x;` is highlighted with a pink box.

The right window is a LLVM IR viewer showing the assembly output for the same code. The assembly is:

```
x86-64 clang 15.0.0 -O2 -march=skylake -fno-slp-vectorize -fno-vectorize
1 sum(int*, int):                                # @sum(int*)
2     test    esi, esi
3     jle     .LBB0_1
4     mov     r8d, esi
5     xor     r9d, r9d
6     xor     esi, esi
7     xor     edx, edx    read
8     xor     eax, eax    read
9     xor     ecx, ecx
10    .LBB0_3:                                     # =>This Inner L
11        add    r9d, dword ptr [rdi + 4*rcx]
12        add    esi, dword ptr [rdi + 4*rcx + 4]
13        add    edx, dword ptr [rdi + 4*rcx + 8]
14        add    eax, dword ptr [rdi + 4*rcx + 12]
15        add    rcx, 4
16        cmp    rcx, r8
17        jb     .LBB0_3
18        add    esi, r9d
19        add    edx, esi
20        add    eax, edx
21    ret
22    .LBB0_1:
23        xor    eax, eax
24        ret
```

The assembly code is color-coded to match the regions in the C++ code: the inner loop body (lines 11-14) is highlighted in yellow, the first iteration of the outer loop (lines 11-14) is highlighted in light green, and the remaining code (lines 15-24) is highlighted in light blue.

Mem2Reg / SROA (scalar replacement of aggregates)

The image shows two windows side-by-side. The left window is a C++ code editor with the following code:

```
1 int sum(int *a, int n) {
2     int sums[4] = {0, 0, 0, 0};
3
4     #pragma GCC unroll 1
5     for (int i = 0; i < n; i += 4) {
6         for (int k = 0; k < 4; ++k)
7             sums[k] += a[k + i];
8     }
9     write to mem
10    int x = 0;      (stack)
11    for (int k = 0; k < 4; ++k)
12        x += sums[k];
13    return x;
14 }
15
```

The line `sums[k] += a[k + i];` is highlighted with a red box and labeled "write to mem". The line `int x = 0;` is highlighted with a purple box and labeled "(stack)". The line `return x;` is highlighted with a pink box.

The right window is a LLVM IR viewer showing the assembly output for the same code. The assembly is:

```
x86-64 clang 15.0.0 -O2 -march=skylake -fno-slp-vectorize -fno-vectorize
1 sum(int*, int):                                # @sum(int*)
2     test    esi, esi
3     jle     .LBB0_1
4     mov     r8d, esi
5     xor     r9d, r9d
6     xor     esi, esi
7     xor     edx, edx      read
8     xor     eax, eax      read
9     xor     ecx, ecx
10    .LBB0_3:                                     # =>This Inner L
11        add    r9d, dword ptr [rdi + 4*rcx]
12        add    esi, dword ptr [rdi + 4*rcx + 4]
13        add    edx, dword ptr [rdi + 4*rcx + 8]
14        add    eax, dword ptr [rdi + 4*rcx + 12]
15        add    rcx, 4
16        cmp    rcx, r8
17        jb     .LBB0_3
18        add    esi, r9d
19        add    edx, esi
20        add    eax, edx
21    ret
22    .LBB0_1:
23        xor    eax, eax
24        ret
```

The assembly code is color-coded to match the regions in the C++ code: the inner loop body (lines 11-14) has a yellow background, the read operations (lines 7-9) have a green background, and the outer loop body (lines 15-21) has a light blue background. The `[rdi + 4*rcx]` and related memory addresses are highlighted with orange boxes.

Mem2Reg / SROA (scalar replacement of aggregates)

The image shows two side-by-side windows from the Clang LLVM IR viewer. The left window displays the original C++ code, and the right window shows the generated assembly code.

C++ Code:

```
1 int sum(int *a, int n) {
2     int sums[4] = {0, 0, 0, 0};
3
4     #pragma GCC unroll 1
5     for (int i = 0; i < n; i += 4) {
6         for (int k = 0; k < 4; ++k)
7             sums[k] += a[k + i];
8     }
9     write to mem
10    int x = 0;      (stack)
11    for (int k = 0; k < 4; ++k)
12        x += sums[k];
13    return x;
14 }
15
```

Assembly Output:

```
x86-64 clang 15.0.0 -O2 -march=skylake -fno-slp-vectorize -fno-vectorize
1 sum(int*, int):                                # @sum(int*)
2     test    esi, esi
3     jle     .LBB0_1
4     mov     r8d, esi
5     xor     r9d, r9d
6     xor     esi, esi
7     xor     edx, edx    read
8     xor     eax, eax    read
9     xor     ecx, ecx
10    .LBB0_3:                                         # =>This Inner L
11    add    r9d, dword ptr [rdi + 4*rcx]
12    add    esi, dword ptr [rdi + 4*rcx + 4]
13    add    edx, dword ptr [rdi + 4*rcx + 8]
14    add    eax, dword ptr [rdi + 4*rcx + 12]
15    add    rcx, 4
16    cmp    rcx, r8
17    jb     .LBB0_3
18    add    esi, r9d
19    add    edx, esi
20    add    eax, edx
21    ret
22    .LBB0_1:
23    xor     eax, eax
24    ret
```

The assembly code illustrates the scalar replacement of aggregates. It shows how the inner loop (lines 11-14) iterates over the four elements of the `sums` array, performing additions and storing the results back into memory. The assembly uses memory addresses like `[rdi + 4*rcx]` to access the elements. The outer loop (lines 11-14) accumulates the sum into the variable `x` on the stack. The assembly output also includes annotations indicating reads and writes to memory.

Mem2Reg / SROA (scalar replacement of aggregates)

```
int sum(int *a, int n) {
    int sums[4] = {0, 0, 0, 0};

    for (int i = 0; i < n; i += 4) {
        for (int k = 0; k < 4; ++k)
            sums[k] += a[k + i];
    }

    int x = 0;
    for (int k = 0; k < 4; ++k)
        x += sums[k];
    return x;
}
```

Mem2Reg / SROA (scalar replacement of aggregates)

```
int sum(int *a, int n) {
    int sums[4] = {0, 0, 0, 0};

    for (int i = 0; i < n; i += 4) {
        for (int k = 0; k < 4; ++k)
            sums[k] += a[k + i];
    }

    int x = 0;
    for (int k = 0; k < 4; ++k)
        x += sums[k];
    return x;
}
```

Mem2Reg / SROA (scalar replacement of aggregates)

```
int sum(int *a, int n) {
    int sums[4] = {0, 0, 0, 0};

    for (int i = 0; i < n; i += 4) {
        for (int k = 0; k < 4; k += 4) {
            sums[0] += a[i];
            sums[1] += a[i + 1];
            sums[2] += a[i + 2];
            sums[3] += a[i + 3];
        }
    }

    int x = 0;
    for (int k = 0; k < 4; k += 4) {
        x += sums[0];
        x += sums[1];
        x += sums[2];
        x += sums[3];
    }
    return x;
}
```

Mem2Reg / SROA (scalar replacement of aggregates)

```
int sum(int *a, int n) {
    int sums[4] = {0, 0, 0, 0};

    for (int i = 0; i < n; i += 4) {
        for (int k = 0; k < 4; k += 4) {
            sums[0] += a[i];
            sums[1] += a[i + 1];
            sums[2] += a[i + 2];
            sums[3] += a[i + 3];
        }
    }

    int x = 0;
    for (int k = 0; k < 4; k += 4) {
        x += sums[0];
        x += sums[1];
        x += sums[2];
        x += sums[3];
    }

    return x;
}
```

Mem2Reg / SROA (scalar replacement of aggregates)

```
int sum(int *a, int n) {
    int sums[4] = {0, 0, 0, 0};

    for (int i = 0; i < n; i += 4) {
        sums[0] += a[i];
        sums[1] += a[i + 1];
        sums[2] += a[i + 2];
        sums[3] += a[i + 3];
    }

    int x = 0;
    x += sums[0];
    x += sums[1];
    x += sums[2];
    x += sums[3];
    return x;
}
```

Mem2Reg / SROA (scalar replacement of aggregates)

```
int sum(int *a, int n) {
    int sums[4] = {0, 0, 0, 0};

    for (int i = 0; i < n; i += 4) {
        sums[0] += a[i];
        sums[1] += a[i + 1];
        sums[2] += a[i + 2];
        sums[3] += a[i + 3];
    }

    int x = 0;
    x += sums[0];
    x += sums[1];
    x += sums[2];
    x += sums[3];
    return x;
}
```

Mem2Reg / SROA (scalar replacement of aggregates)

```
int sum(int *a, int n) {
    int sums_0 = 0;
    int sums_1 = 0;
    int sums_2 = 0;
    int sums_3 = 0;

    for (int i = 0; i < n; i += 4) {
        sums_0 += a[i];
        sums_1 += a[i + 1];
        sums_2 += a[i + 2];
        sums_3 += a[i + 3];
    }

    int x = 0;
    x += sums_0;
    x += sums_1;
    x += sums_2;
    x += sums_3;
    return x;
}
```

Mem2Reg / SROA (scalar replacement of aggregates)

The screenshot shows a debugger interface with two panes. The left pane displays the C++ source code for a function named `sum`. The right pane shows the generated assembly code for the same function, compiled with optimization flags `-O2 -march=skylake -fno-tree-vectorize`.

C++ Source Code:

```
1 int sum(int *a, int n) {
2     int sums[4] = {0, 0, 0, 0};
3
4     #pragma GCC unroll 1
5     for (int i = 0; i < n; i += 4) {
6         for (int k = 0; k < 4; ++k)
7             sums[k] += a[k + i];
8     }
9
10    int x = 0;
11    for (int k = 0; k < 4; ++k)
12        x += sums[k];
13    return x;
14 }
```

Assembly Output:

```
1 sum(int*, int):
2     mov    QWORD PTR [rsp-24], 0
3     mov    QWORD PTR [rsp-16], 0
4     xor    edx, edx
5     xor    r9d, r9d
6     lea    r8, [rsp-8]
7     test   esi, esi
8     jle   .L9
9 .L5:
10    lea    rax, [rsp-24]
11    mov    rcx, rdi
12 .L6:
13    add    edx, DWORD PTR [rcx]
14    mov    DWORD PTR [rax], edx
15    add    rax, 4
16    add    rcx, 4
17    cmp    r8, rax
18    je     .L10
19    mov    edx, DWORD PTR [rax]
20    jmp    .L6
21 .L10:
22    add    r9d, 4
23    mov    edx, DWORD PTR [rsp-24]
24    add    rdi, 16
25    cmp    esi, r9d
26    jg    .L5
27    mov    eax, DWORD PTR [rsp-20]
28    add    eax, edx
29    add    eax, DWORD PTR [rsp-16]
30    add    eax, DWORD PTR [rsp-12]
31    ret
32 .L9:
33    xor    eax, eax
34    ret
```

Mem2Reg / SROA (scalar replacement of aggregates)

The screenshot shows two windows side-by-side. The left window is a C++ code editor with the following code:

```
1 int sum(int *a, int n) {
2     int sums[4] = {0, 0, 0, 0};
3
4     #pragma GCC unroll 1
5     for (int i = 0; i < n; i += 4) {
6         for (int k = 0; k < 4; ++k)
7             sums[k] += a[k + i];
8     }
9
10    int x = 0;
11    for (int k = 0; k < 4; ++k)
12        x += sums[k];
13    return x;
14 }
```

The right window is a debugger showing the assembly output for x86-64 gcc 12.2, generated with -O2 -march=skylake -fno-tree-vectorize. The assembly code is as follows:

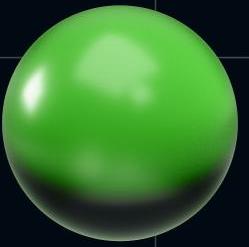
```
1 sum(int*, int):
2     mov    QWORD PTR [rsp-24], 0
3     mov    QWORD PTR [rsp-16], 0
4     xor    edx, edx
5     xor    r9d, r9d
6     lea    r8, [rsp-8]
7     test   esi, esi
8     jle   .L9
9 .L5:
10    lea    rax, [rsp-24]
11    mov    rcx, rdi
12 .L6:
13    add    edx, DWORD PTR [rcx]
14    mov    DWORD PTR [rax], edx
15    add    rax, 4
16    add    rcx, 4
17    cmp    r8, rax
18    je     .L10
19    mov    edx, DWORD PTR [rax]
20    jmp   .L6
21 .L10:
22    add    r9d, 4
23    mov    edx, DWORD PTR [rsp-24]
24    add    rdi, 16
25    cmp    esi, r9d
26    jg    .L5
27    mov    eax, DWORD PTR [rsp-20]
28    add    eax, edx
29    add    eax, DWORD PTR [rsp-16]
30    add    eax, DWORD PTR [rsp-12]
31    ret
32 .L9:
33    xor    eax, eax
34    ret
```

Annotations in red highlight specific instructions and memory locations:

- Annotations "write" are placed next to the `mov` instructions at lines 2 and 3.
- An annotation "write" is placed next to the `mov` instruction at line 14.
- A large yellow rectangular highlight covers the entire assembly code area.

Outline of the talk

- Loop from compiler perspective
- Loop unrolling basics
- Loop unrolling overhead
- GCC and CLANG unroll details
- New optimization opportunities after unrolling
- Example

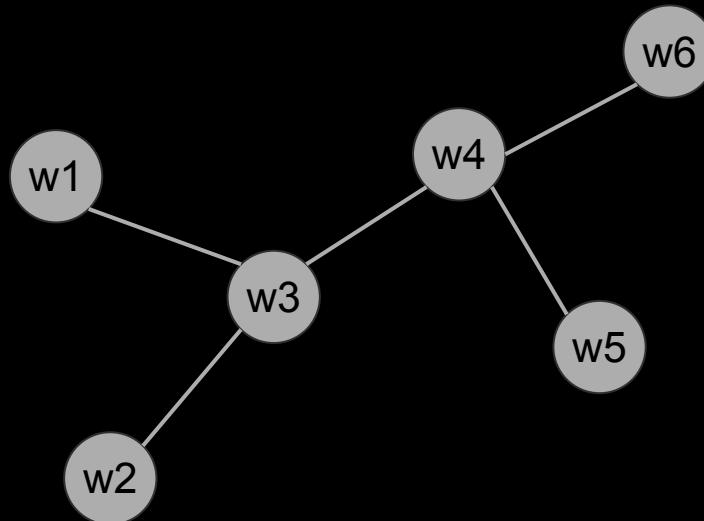


Example

Need to solve Minimum Weighted Vertex Cover Problem ASAP

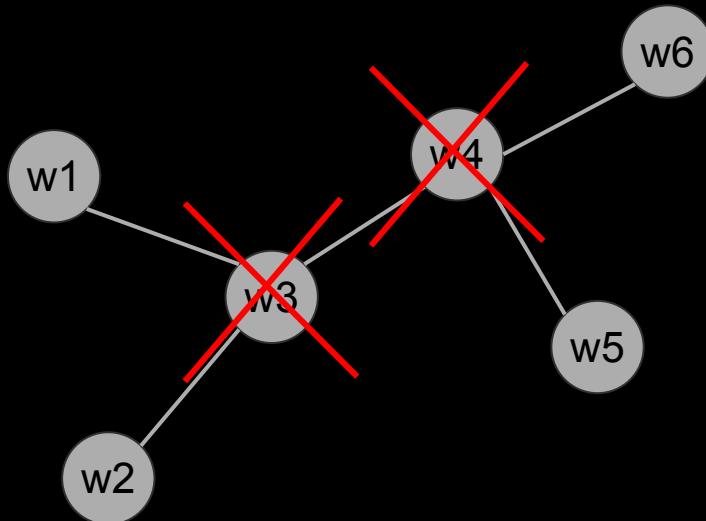
Example

Need to solve Minimum Weighted Vertex Cover Problem ASAP



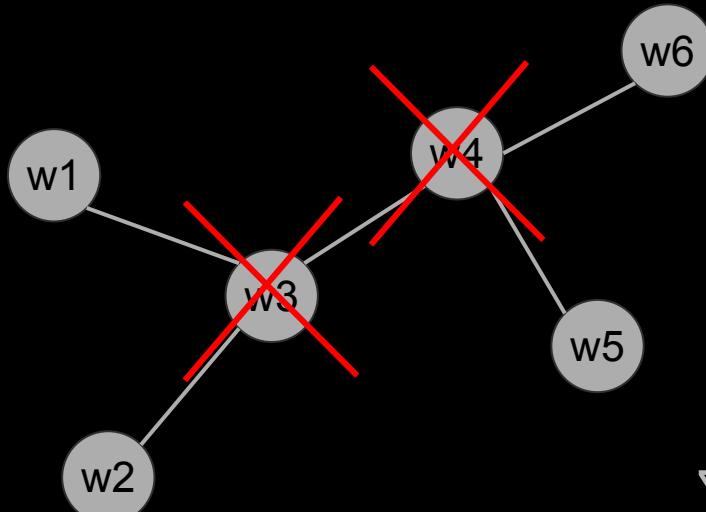
Example

Need to solve Minimum Weighted Vertex Cover Problem ASAP



Example

Need to solve Minimum Weighted Vertex Cover Problem ASAP

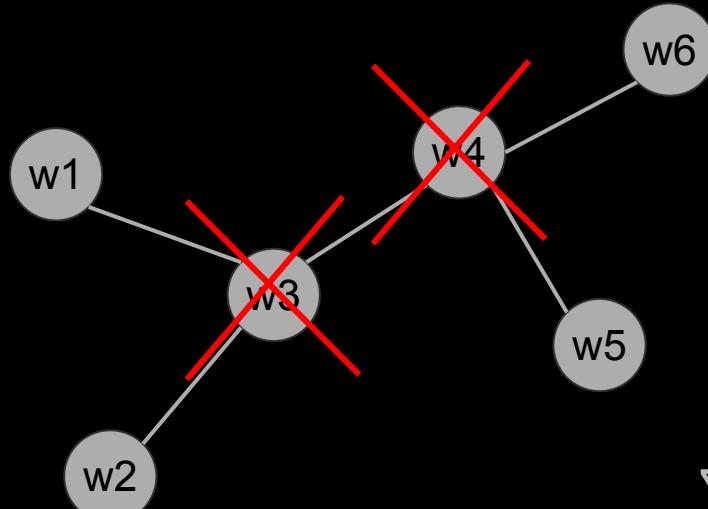


$$\sum_{\text{deleted}} w_i \rightarrow \min$$

Example

Need to solve Minimum Weighted Vertex Cover Problem ASAP

NP-complete!



$$\sum_{\text{deleted}} w_i \rightarrow \min$$

Example

Need to solve Minimum Weighted Vertex Cover Problem ASAP

But...

Example

Need to solve Minimum Weighted Vertex Cover Problem ASAP

But...

- Assume weights $\in \mathbb{N}$

Example

Need to solve Minimum Weighted Vertex Cover Problem ASAP

But...

- Assume weights $\in \mathbb{N}$
- Assume NumVertex ≤ 16

Example

Need to solve Minimum Weighted Vertex Cover Problem ASAP

But...

- Assume weights $\in \mathbb{N}$
- Assume NumVertex ≤ 16
- Encode task as... bits:
 - vertices set is given as bitmask (uint32_t)

Example

```
uint32_t getBestSubset(
    const vector<uint32_t> &InW,
    const vector<uint32_t> &InM) {
    array<uint32_t, 16> Weights = {0};
    array<uint32_t, 16> AdjMasks = {0};
    copy(begin(InW), end(InW), begin(Weights));
    copy(begin(InM), end(InM), begin(AdjMasks));
```

Example

```
const auto GetSubsetWeight = [&](uint32_t Subset) {
    uint32_t RV = 0;
    for (size_t I = 0; I < 16; ++I)
        RV += (Subset & (1 << I)) ? Weights[I] : 0;
    return RV;
};

const auto IsCompatible = [&](uint32_t Subset) {
    uint32_t IncompatibleVerMask = 0;
    for (size_t I = 0; I < 16; ++I)
        IncompatibleVerMask |= (Subset & (1 << I)) ? AdjMasks[I] : 0;
    return (Subset & IncompatibleVerMask) == 0;
};
```

Example

```
uint32_t BestWeight = 0;
uint32_t BestSubset = 0;
const uint32_t NumSubsets = 1 << InW.size();
for (uint32_t Subset = 1; Subset < NumSubsets; ++Subset) {

    const uint32_t Weight = GetSubsetWeight(Subset);
    if (Weight <= BestWeight)
        continue;

    if (!IsCompatible(Subset))
        continue;

    BestWeight = Weight;
    BestSubset = Subset;
}

return BestSubset;
}
```

Example

The image shows a debugger interface with two main panes. The left pane displays C++ source code, and the right pane shows the generated assembly code.

C++ Source Code:

```
17     const uint32_t NumVtx) {
18     array<uint32_t, N> Weights = {0};
19     array<uint32_t, N> AdjMasks = {0};
20     copy(begin(InW), end(InW), begin(Weights));
21     copy(begin(InM), end(InM), begin(AdjMasks));
22
23     const auto GetSubsetWeight = [&](uint32_t Subset) {
24         uint32_t RV = 0;
25         for (size_t I = 0; I < N; ++I)
26             RV += (Subset & (1 << I)) ? Weights[I] : 0;
27         return RV;
28     };
29
30     const auto IsCompatible = [&](uint32_t Subset) {
31         uint32_t IncompatMask = 0;
32         for (size_t I = 0; I < N; ++I)
33             IncompatMask |= (Subset & (1 << I)) ? AdjMasks[I] : 0;
34         return (Subset & IncompatMask) == 0;
35     };
36
37     uint32_t BestWeight = 0;
38     uint32_t BestSubset = 0;
39     const uint32_t NumSubsets = 1 << NumVtx;
40     for (uint32_t Subset = 1; Subset < NumSubsets; ++Subset) {
41         const uint32_t Weight = GetSubsetWeight(Subset);
42         if (Weight <= BestWeight)
43             continue;
44
45         if (!IsCompatible(Subset))
46             continue;
47
48         BestWeight = Weight;
49         BestSubset = Subset;
50     }
51     return BestSubset;
52 }
```

Assembly Code:

```
59    vpxor  xmm6, xmm6, xmm6
60    xor    eax, eax
61    jmp    .LBB0_8
62 .LBB0_10:                                # in Loop: Header=.LBB0_8 Depth=1
63    inc    ecx
64    shrx   esi, ecx, ebx
65    test   esi, esi
66    jne    .LBB0_6
67 .LBB0_8:                                 # =>This Inner Loop Header: Depth=1
68    vmovd  xmm7, ecx
69    vpbroadcastd ymm7, xmm7
70    vpand  ymm8, ymm7, ymm4
71    vpand  ymm7, ymm7, ymm5
72    vpcmpqd ymm5, ymm5, ymm7, ymm5
73    vpand  ymm9, ymm9, ymm1
74    vpcmpqd ymm10, ymm8, ymm4
75    vpand  ymm10, ymm10, ymm0
76    vpaddd ymm9, ymm10, ymm9
77    vextracti128 xmm2, ymm9, 1
78    vpaddd xmm2, xmm9, xmm2
79    vpshufd xmm3, xmm2, 238                # xmm3 = xmm2[2,3,2,3]
80    vpadddd xmm2, xmm2, xmm3
81    vpshufd xmm3, xmm2, 85                 # xmm3 = xmm2[1,1,1,1]
82    vpadddd xmm2, xmm2, xmm3
83    vmovd  esi, xmm2
84    cmp    esi, edx
85    jbe    .LBB0_10
86    vpcmpgtd ymm2, ymm8, ymm6
87    vpand  ymm2, ymm11, ymm2
88    vpcmpgtd ymm3, ymm7, ymm6
89    vpand  ymm3, ymm12, ymm3
90    vpor   ymm2, ymm2, ymm3
91    vextracti128 xmm3, ymm2, 1
92    vpor   xmm2, xmm2, xmm3
93    vpshufd xmm3, xmm2, 238                # xmm3 = xmm2[2,3,2,3]
94    vpor   xmm2, xmm2, xmm3
95    vpshufd xmm3, xmm2, 85                 # xmm3 = xmm2[1,1,1,1]
96    vpor   xmm2, xmm2, xmm3
97    vmovd  edi, xmm2
98    test   edi, ecx
99    cmovne eax, ecx
100   cmovne edx, esi
101   jmp    .LBB0_10
```

The assembly code is annotated with comments explaining the SIMD operations and loop structure. The debugger also shows memory dump panes for registers and memory locations.

Example

The image shows a debugger interface with two panes. The left pane displays the C++ source code, and the right pane displays the generated assembly code.

C++ Source Code:

```
17     const uint32_t NumVtx) {
18     array<uint32_t, N> Weights = {0};
19     array<uint32_t, N> AdjMasks = {0};
20     copy(begin(InW), end(InW), begin(Weights));
21     copy(begin(InM), end(InM), begin(AdjMasks));
22
23     const auto GetSubsetWeight = [&](uint32_t Subset) {
24         uint32_t RV = 0;
25         for (size_t I = 0; I < N; ++I)
26             RV += (Subset & (1 << I)) ? Weights[I] : 0;
27         return RV;
28     };
29
30     const auto IsCompatible = [&](uint32_t Subset) {
31         uint32_t IncompatMask = 0;
32         for (size_t I = 0; I < N; ++I)
33             IncompatMask |= (Subset & (1 << I)) ? AdjMasks[I] : 0;
34         return (Subset & IncompatMask) == 0;
35     };
36
37     uint32_t BestWeight = 0;
38     uint32_t BestSubset = 0;
39     const uint32_t NumSubsets = 1 << NumVtx;
40     for (uint32_t Subset = 1; Subset < NumSubsets; ++Subset) {
41         const uint32_t Weight = GetSubsetWeight(Subset);
42         if (Weight <= BestWeight)
43             continue;
44
45         if (!IsCompatible(Subset))
46             continue;
47
48         BestWeight = Weight;
49         BestSubset = Subset;
50     }
51     return BestSubset;
52 }
```

Assembly Code:

```
59    vpxor  xmm6, xmm6, xmm6
60    xor    eax, eax
61    jmp    .LBB0_8
62 .LBB0_10:                                # in Loop: Header=.LBB0_8 Depth=1
63    inc    ecx
64    shrx   esi, ecx, ebx
65    test   esi, esi
66    jne    .LBB0_6
67 .LBB0_8:                                 # =>This Inner Loop Header: Depth=1
68    vmovd  xmm7, ecx
69    vpbroadcastd ymm7, xmm7
70    vpand  ymm8, ymm7, ymm4
71    vpand  ymm7, ymm7, ymm5
72    vpcmpqd ymm5, ymm5, ymm7, ymm5
73    vpand  ymm9, ymm9, ymm1
74    vpcmpqd ymm10, ymm8, ymm4
75    vpand  ymm10, ymm10, ymm0
76    vpaddd ymm9, ymm10, ymm9
77    vextracti128 xmm2, ymm9, 1
78    vpaddd xmm2, xmm9, xmm2
79    vpshufd xmm3, xmm2, 238               # xmm3 = xmm2[2,3,2,3]
80    vpadddd xmm2, xmm2, xmm3
81    vpshufd xmm3, xmm2, 85                # xmm3 = xmm2[1,1,1,1]
82    vpadddd xmm2, xmm2, xmm3
83    vmovd  esi, xmm2
84    cmp    esi, edx
85    jbe    .LBB0_10
86    vpcmpgtd ymm2, ymm8, ymm6
87    vpand  ymm2, ymm11, ymm2
88    vpcmpgtd ymm3, ymm7, ymm6
89    vpand  ymm3, ymm12, ymm3
90    vpor   ymm2, ymm2, ymm3
91    vextracti128 xmm3, ymm2, 1
92    vpor   xmm2, xmm2, xmm3
93    vpshufd xmm3, xmm2, 238               # xmm3 = xmm2[2,3,2,3]
94    vpor   xmm2, xmm2, xmm3
95    vpshufd xmm3, xmm2, 85                # xmm3 = xmm2[1,1,1,1]
96    vpor   xmm2, xmm2, xmm3
97    vmovd  edi, xmm2
98    test   edi, ecx
99    cmovne eax, ecx
100   cmovne edx, esi
101   jmp    .LBB0_10
```

The assembly code is annotated with comments explaining the SIMD operations and loop structure based on the highlighted regions in the C++ code.

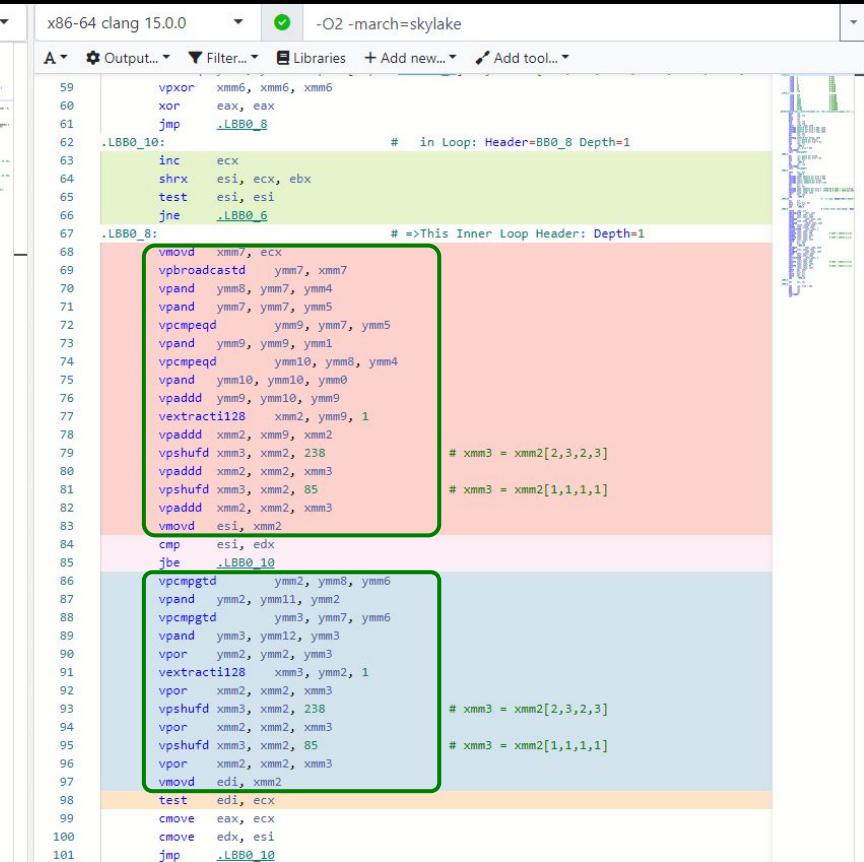
Example

Left pane: C++ code (Visual Studio)

```
17     const uint32_t NumVtx) {
18     array<uint32_t, N> Weights = {0};
19     array<uint32_t, N> AdjMasks = {0};
20     copy(begin(InW), end(InW), begin(Weights));
21     copy(begin(InM), end(InM), begin(AdjMasks));
22
23     const auto GetSubsetWeight = [&](uint32_t Subset) {
24         uint32_t RV = 0;
25         for (size_t I = 0; I < N; ++I)
26             RV += (Subset & (1 << I)) ? Weights[I] : 0;
27         return RV;
28     };
29
30     const auto IsCompatible = [&](uint32_t Subset) {
31         uint32_t IncompatMask = 0;
32         for (size_t I = 0; I < N; ++I)
33             IncompatMask |= (Subset & (1 << I)) ? AdjMasks[I] : 0;
34         return (Subset & IncompatMask) == 0;
35     };
36
37
38     uint32_t BestWeight = 0;
39     uint32_t BestSubset = 0;
40     const uint32_t NumSubsets = 1 << NumVtx;
41     for (uint32_t Subset = 1; Subset < NumSubsets; ++Subset) {
42         const uint32_t Weight = GetSubsetWeight(Subset);
43         if (Weight <= BestWeight)
44             continue;
45
46         if (!IsCompatible(Subset))
47             continue;
48
49         BestWeight = Weight;
50         BestSubset = Subset;
51     }
52     return BestSubset;
53 }
```

Right pane: Assembly output (x86-64 clang 15.0.0)

```
59     vpxor    xmm6, xmm6, xmm6
60     xor     eax, eax
61     jmp     .LBB0_8
62 .LBB0_10:                                # in Loop: Header=.LBB0_8 Depth=1
63     inc     ecx
64     shrx   esi, ecx, ebx
65     test   esi, esi
66     jne     .LBB0_6
67 .LBB0_8:                                # =>This Inner Loop Header: Depth=1
68     vmovd  xmm7, ecx
69     vpbroadcastd  ymm7, xmm7
70     vpand  ymm8, ymm7, ymm4
71     vpand  ymm7, ymm7, ymm5
72     vpcmpqd  ymm5, ymm7, ymm5
73     vpand  ymm9, ymm9, ymm1
74     vpcmpqd  ymm10, ymm8, ymm4
75     vpand  ymm10, ymm10, ymm0
76     vpaddd ymm9, ymm10, ymm9
77     vextracti128  xmm2, ymm9, 1
78     vpaddd xmm2, xmm9, xmm2
79     vpshufd xmm3, xmm2, 238
80     vpadddd ymm2, xmm2, xmm3
81     vpshufd xmm3, xmm2, 85
82     vpadddd ymm2, xmm2, xmm3
83     vmovd  esi, xmm2
84     cmp     esi, edx
85     jbe     .LBB0_10
86     vpcmpgtd  ymm2, ymm8, ymm6
87     vpand  ymm2, ymm11, ymm2
88     vpcmpgtd  ymm3, ymm7, ymm6
89     vpand  ymm3, ymm12, ymm3
90     vpor   ymm2, ymm2, ymm3
91     vextracti128  xmm3, ymm2, 1
92     vpor   xmm2, xmm2, xmm3
93     vpshufd xmm3, xmm2, 238
94     vpor   xmm2, xmm2, xmm3
95     vpshufd xmm3, xmm2, 85
96     vpor   xmm2, xmm2, xmm3
97     vmove  edi, xmm2
98     test   edi, ecx
99     cmovne eax, ecx
100    cmovne edx, esi
101    jmp     .LBB0_10
```



Example

The image shows a debugger interface with two panes. The left pane displays C++ code, and the right pane shows the generated assembly code.

C++ Code:

```
17     const uint32_t NumVtx) {  
18         array<uint32_t, N> Weights = {0};  
19         array<uint32_t, N> AdjMasks = {0};  
20         copy(begin(InW), end(InW), begin(Weights));  
21         copy(begin(InM), end(InM), begin(AdjMasks));  
22  
23     const auto GetSubsetWeight = [&](uint32_t Subset) {  
24         uint32_t RV = 0;  
25         for (size_t I = 0; I < N; ++I)  
26             RV += (Subset & (1 << I)) ? Weights[I] : 0;  
27         return RV;  
28     };  
29  
30     const auto IsCompatible = [&](uint32_t Subset) {  
31         uint32_t IncompatMask = 0;  
32         for (size_t I = 0; I < N; ++I)  
33             IncompatMask |= (Subset & (1 << I)) ? AdjMasks[I] : 0;  
34         return (Subset & IncompatMask) == 0;  
35     };  
36  
37     uint32_t BestWeight = 0;  
38     uint32_t BestSubset = 0;  
39     const uint32_t NumSubsets = 1 << NumVtx;  
40     for (uint32_t Subset = 1; Subset < NumSubsets; ++Subset) {  
41         const uint32_t Weight = GetSubsetWeight(Subset);  
42         if (Weight <= BestWeight)  
43             continue;  
44  
45         if (!IsCompatible(Subset))  
46             continue;  
47  
48         BestWeight = Weight;  
49         BestSubset = Subset;  
50     }  
51     return BestSubset;  
52 }
```

Assembly Output:

```
x86-64 clang 15.0.0 -O2 -march=skylake  
A A Output... Filter... Libraries + Add new...  
59    vpxor    xmm6, xmm6, xmm6  
60    xor     eax, eax  
61    jmp     .LBB0_8  
62 .LBB0_8:  
63    inc     ecx  
64    shrx   esi, ecx, ebx  
65    test    esi, esi  
66    jne     .LBB0_6  
67 .LBB0_6:  
68    vmovd  xmm7, ecx  
69    vpbroadcastd  ymm7, xmm7  
70    vpand   ymm8, ymm7, ymm4  
71    vpand   ymm7, ymm7, ymm5  
72    vpcmpqd  ymm9, ymm7, ymm5  
73    vpand   ymm9, ymm9, ymm1  
74    vpcmpqd  ymm10, ymm8, ymm4  
75    vpand  ymm10, ymm10, ymm0  
76    vpaddd  ymm9, ymm10, ymm9  
77    vextracti128  xmm2, ymm9, 1  
78    vpaddd  xmm2, xmm9, xmm9  
79    vpushufd  xmm3, xmm2, 238  
80    vpaddd  xmm2, xmm2, xmm3  
81    vpushufd  xmm3, xmm2, 85  
82    vpaddd  xmm2, xmm2, xmm3  
83    vmovd  esi, xmm2  
84    cmp     esi, edx  
85    jbe     .LBB0_10  
86    vpcmpgtd  ymm2, ymm8, ymm6  
87    vpand  ymm2, ymm11, ymm2  
88    vpcmpgtd  ymm3, ymm7, ymm6  
89    vpand  ymm3, ymm12, ymm3  
90    vpor    ymm2, ymm2, ymm3  
91    vextracti128  xmm3, ymm2, 1  
92    vpor    xmm2, xmm2, xmm3  
93    vpushufd  xmm3, xmm2, 238  
94    vpor    xmm2, xmm2, xmm3  
95    vpushufd  xmm3, xmm2, 85  
96    vpor    xmm2, xmm2, xmm3  
97    vmovd  edi, xmm2  
98    test    edi, ecx  
99    cmovne eax, ecx  
100   cmovne edx, esi  
101   jmp     .LBB0_10
```

Annotations: The code in the C++ pane has several sections highlighted with colored boxes: a green box around the loop in lines 25-34, a pink box around the inner loop in lines 32-33, and a light blue box around the outer loop in lines 40-51. The corresponding assembly code in the right pane also has these sections highlighted with matching colors: a green box around the first few instructions (lines 59-66), a pink box around the middle section of the loop (lines 68-82), and a light blue box around the final section of the loop (lines 86-101). The text "no loops" is displayed in green at the top right of the assembly pane.

Example

The image shows a debugger interface with two panes. The left pane displays C++ code, and the right pane displays the generated assembly code.

C++ Code:

```
17     const uint32_t NumVtx) {  
18         array<uint32_t, N> Weights = {0};  
19         array<uint32_t, N> AdjMasks = {0};  
20         copy(begin(InW), end(InW), begin(Weights));  
21         copy(begin(InM), end(InM), begin(AdjMasks));  
22  
23     const auto GetSubsetWeight = [&](uint32_t Subset) {  
24         uint32_t RV = 0;  
25         for (size_t I = 0; I < N; ++I)  
26             RV += (Subset & (1 << I)) ? Weights[I] : 0;  
27         return RV;  
28     };  
29  
30     const auto IsCompatible = [&](uint32_t Subset) {  
31         uint32_t IncompatMask = 0;  
32         for (size_t I = 0; I < N; ++I)  
33             IncompatMask |= (Subset & (1 << I)) ? AdjMasks[I] : 0;  
34         return (Subset & IncompatMask) == 0;  
35     };  
36  
37     uint32_t BestWeight = 0;  
38     uint32_t BestSubset = 0;  
39     const uint32_t NumSubsets = 1 << NumVtx;  
40     for (uint32_t Subset = 1; Subset < NumSubsets; ++Subset) {  
41         const uint32_t Weight = GetSubsetWeight(Subset);  
42         if (Weight <= BestWeight)  
43             continue;  
44  
45         if (!IsCompatible(Subset))  
46             continue;  
47  
48         BestWeight = Weight;  
49         BestSubset = Subset;  
50     }  
51     return BestSubset;  
52 }
```

Assembly Output:

```
59     vpxor    xmm6, xmm6, xmm6  
60     xor     eax, eax  
61     jmp     .LBB0_8  
62 .LBB0_8:  
63     inc     ecx  
64     shrx   esi, ecx, ebx  
65     test    esi, esi  
66     jne     .LBB0_6  
67 .LBB0_6:  
68     vmovd  xmm7, ecx  
69     vpbroadcastd  ymm7, xmm7  
70     vpand   ymm8, ymm7, ymm4  
71     vpand   ymm7, ymm7, ymm5  
72     vpcmpqd  ymm5, ymm7, ymm5  
73     vpand   ymm9, ymm9, ymm1  
74     vpcmpqd  ymm10, ymm8, ymm4  
75     vpand   ymm10, ymm10, ymm0  
76     vpaddd  ymm9, ymm10, ymm9  
77     vextracti128  xmm2, ymm9, 1  
78     vpaddd  xmm2, ymm9, ymm2  
79     vpushufd  xmm3, xmm2, 238  
80     vpaddd  xmm2, ymm2, ymm3  
81     vpushufd  xmm3, xmm2, 85  
82     vpaddd  xmm2, ymm2, ymm3  
83     vmovd  esi, xmm2  
84     cmp     esi, edx  
85     jbe     .LBB0_10  
86     vpcmpgtd  ymm2, ymm8, ymm6  
87     vpand   ymm2, ymm11, ymm2  
88     vpcmpgtd  ymm3, ymm7, ymm6  
89     vpand   ymm3, ymm12, ymm3  
90     vpor    ymm2, ymm2, ymm3  
91     vextracti128  xmm3, ymm2, 1  
92     vpor    xmm2, ymm2, xmm3  
93     vpushufd  xmm3, xmm2, 238  
94     vpor    xmm2, ymm2, ymm3  
95     vpushufd  xmm3, xmm2, 85  
96     vpor    xmm2, ymm2, ymm3  
97     vmovd  edi, xmm2  
98     test    edi, ecx  
99     cmovne eax, ecx  
100    cmovne edx, esi  
101    jmp     .LBB0_10
```

Annotations:

- A green box highlights the loop in the C++ code where the subset weight is calculated.
- A green box highlights the loop in the assembly code corresponding to the C++ loop.
- A pink box highlights the inner loop in the C++ code where the subset is checked for compatibility.
- A pink box highlights the inner loop in the assembly code corresponding to the C++ compatibility check.
- A yellow box highlights the final loop in the C++ code where the best subset is found.
- A yellow box highlights the final loop in the assembly code corresponding to the C++ best subset finding logic.

Text on the right:

no loops
no memory reads

Example

```
C++  
x86-64 clang 15.0.0 -O2 -march=skylar  
  
17     const uint32_t NumVtx) {  
18     array<uint32_t, N> Weights = {0};  
19     array<uint32_t, N> AdjMasks = {0};  
20     copy(begin(InW), end(InW), begin(Weights));  
21     copy(begin(InM), end(InM), begin(AdjMasks));  
22  
23     const auto GetSubsetWeight = [&](uint32_t Subset) {  
24         uint32_t RV = 0;  
25         for (size_t I = 0; I < N; ++I)  
26             RV += (Subset & (1 << I)) ? Weights[I] : 0;  
27         return RV;  
28     };  
29  
30     const auto IsCompatible = [&](uint32_t Subset) {  
31         uint32_t IncompatMask = 0;  
32         for (size_t I = 0; I < N; ++I)  
33             IncompatMask |= (Subset & (1 << I)) ? AdjMasks[I] : 0;  
34         return (Subset & IncompatMask) == 0;  
35     };  
36  
37  
38     uint32_t BestWeight = 0;  
39     uint32_t BestSubset = 0;  
40     const uint32_t NumSubsets = 1 << NumVtx;  
41     for (uint32_t Subset = 1; Subset < NumSubsets; ++Subset) {  
42         const uint32_t Weight = GetSubsetWeight(Subset);  
43         if (Weight <= BestWeight)  
44             continue;  
45  
46         if (!IsCompatible(Subset))  
47             continue;  
48  
49         BestWeight = Weight;  
50         BestSubset = Subset;  
51     }  
52     return BestSubset;  
53 }
```

no loops
no memory reads
vectorization
(xmm/ymm)

Example

N = 14	time
clang O2 <ul style="list-style-type: none">● unroll: off● vectorization: off	
clang O2 <ul style="list-style-type: none">● unroll: on● vectorization: off	
clang O2 <ul style="list-style-type: none">● unroll: on● vectorization: on	

Example

N = 14	time
clang O2 <ul style="list-style-type: none">● unroll: off● vectorization: off	668 mcs
clang O2 <ul style="list-style-type: none">● unroll: on● vectorization: off	
clang O2 <ul style="list-style-type: none">● unroll: on● vectorization: on	

Example

N = 14	time
clang O2 • unroll: off • vectorization: off	668 mcs
clang O2 • unroll: on • vectorization: off	93 mcs
clang O2 • unroll: on • vectorization: on	

x 7

Example

N = 14	time	
clang O2 • unroll: off • vectorization: off	668 mcs	
clang O2 • unroll: on • vectorization: off	93 mcs	x 7
clang O2 • unroll: on • vectorization: on	26 mcs	x 26

Outcome

- Unroll is generally:
 - useful with large trip count
 - harmful with small trip count
- Trust defaults... but...
 - Clang likely unrolls by default (small trip count case suffers)
 - GCC (O2) likely does not unroll by default. Use PGO / pragma.
- Step and unroll count should be power of 2
- Provide a hint for full unroll for GCC (O2)

Thank you


```

int sum(int *a, int s, int b) {
    int sum(int *a, int n) {
        int x = 0;
        // #pragma GCC unroll 8
        for (int i = 0; i < n; ++i)
            x += a[i];
        return x;
    }
    int x = 0;
    for (int i = s; i < b; ++i)
        x += a[i];
    return x;
}

int sum(int *a, int n) {
    int sums[4] = {0, 0, 0, 0};
    #pragma GCC unroll(1)
    for (int i = 0; i < n; i += 4) {
        #pragma GCC unroll(4)
        for (int k = 0; k < 4; ++k)
            sums[k] += a[k + i];
    }
    int x = 0;
    #pragma GCC unroll(4)
    for (int k = 0; k < 4; ++k)
        x += sums[k];
    return x;
}

```

-O2 -march=skylake -fno-tree-vectorize
-02 -march=skylake -fno-vectorize -fno-slp-vectorize