



Institute for Computing
Systems Architecture

PROFILE-DIRECTED SEMI-AUTOMATIC PARALLELISATION

Björn Franke

Institute for Computing Systems Architecture
University of Edinburgh

Erlangen, 11 July 2011



Overview

- Motivation
- Part 1: Profile Directed Dependence Analysis
- Part 2: Machine Learning Based Mapping
- Part 3: Extraction of Pipeline Parallelism
- Conclusions



Motivation

- Multi-cores are ubiquitous
 - Try buying a single-core mobile phone, netbook, PC, or whatsoever!
- Legacy code base of sequential applications
 - Writing parallel applications is hard!
- No single parallel machine model
 - Different parallel programs for different parallel machines
 - Parallelisation is **not** a one-off activity: Need to parallelise each application for each new platform again



Motivation

- Multi-cores are ubiquitous
 - Try buying a single-core mobile phone netbook

Tool Support for Parallelisation Increases Programmer Efficiency, Reduces Time-to-Market, Reduces Number of Bugs, Secures Software Investments etc.

BUT...

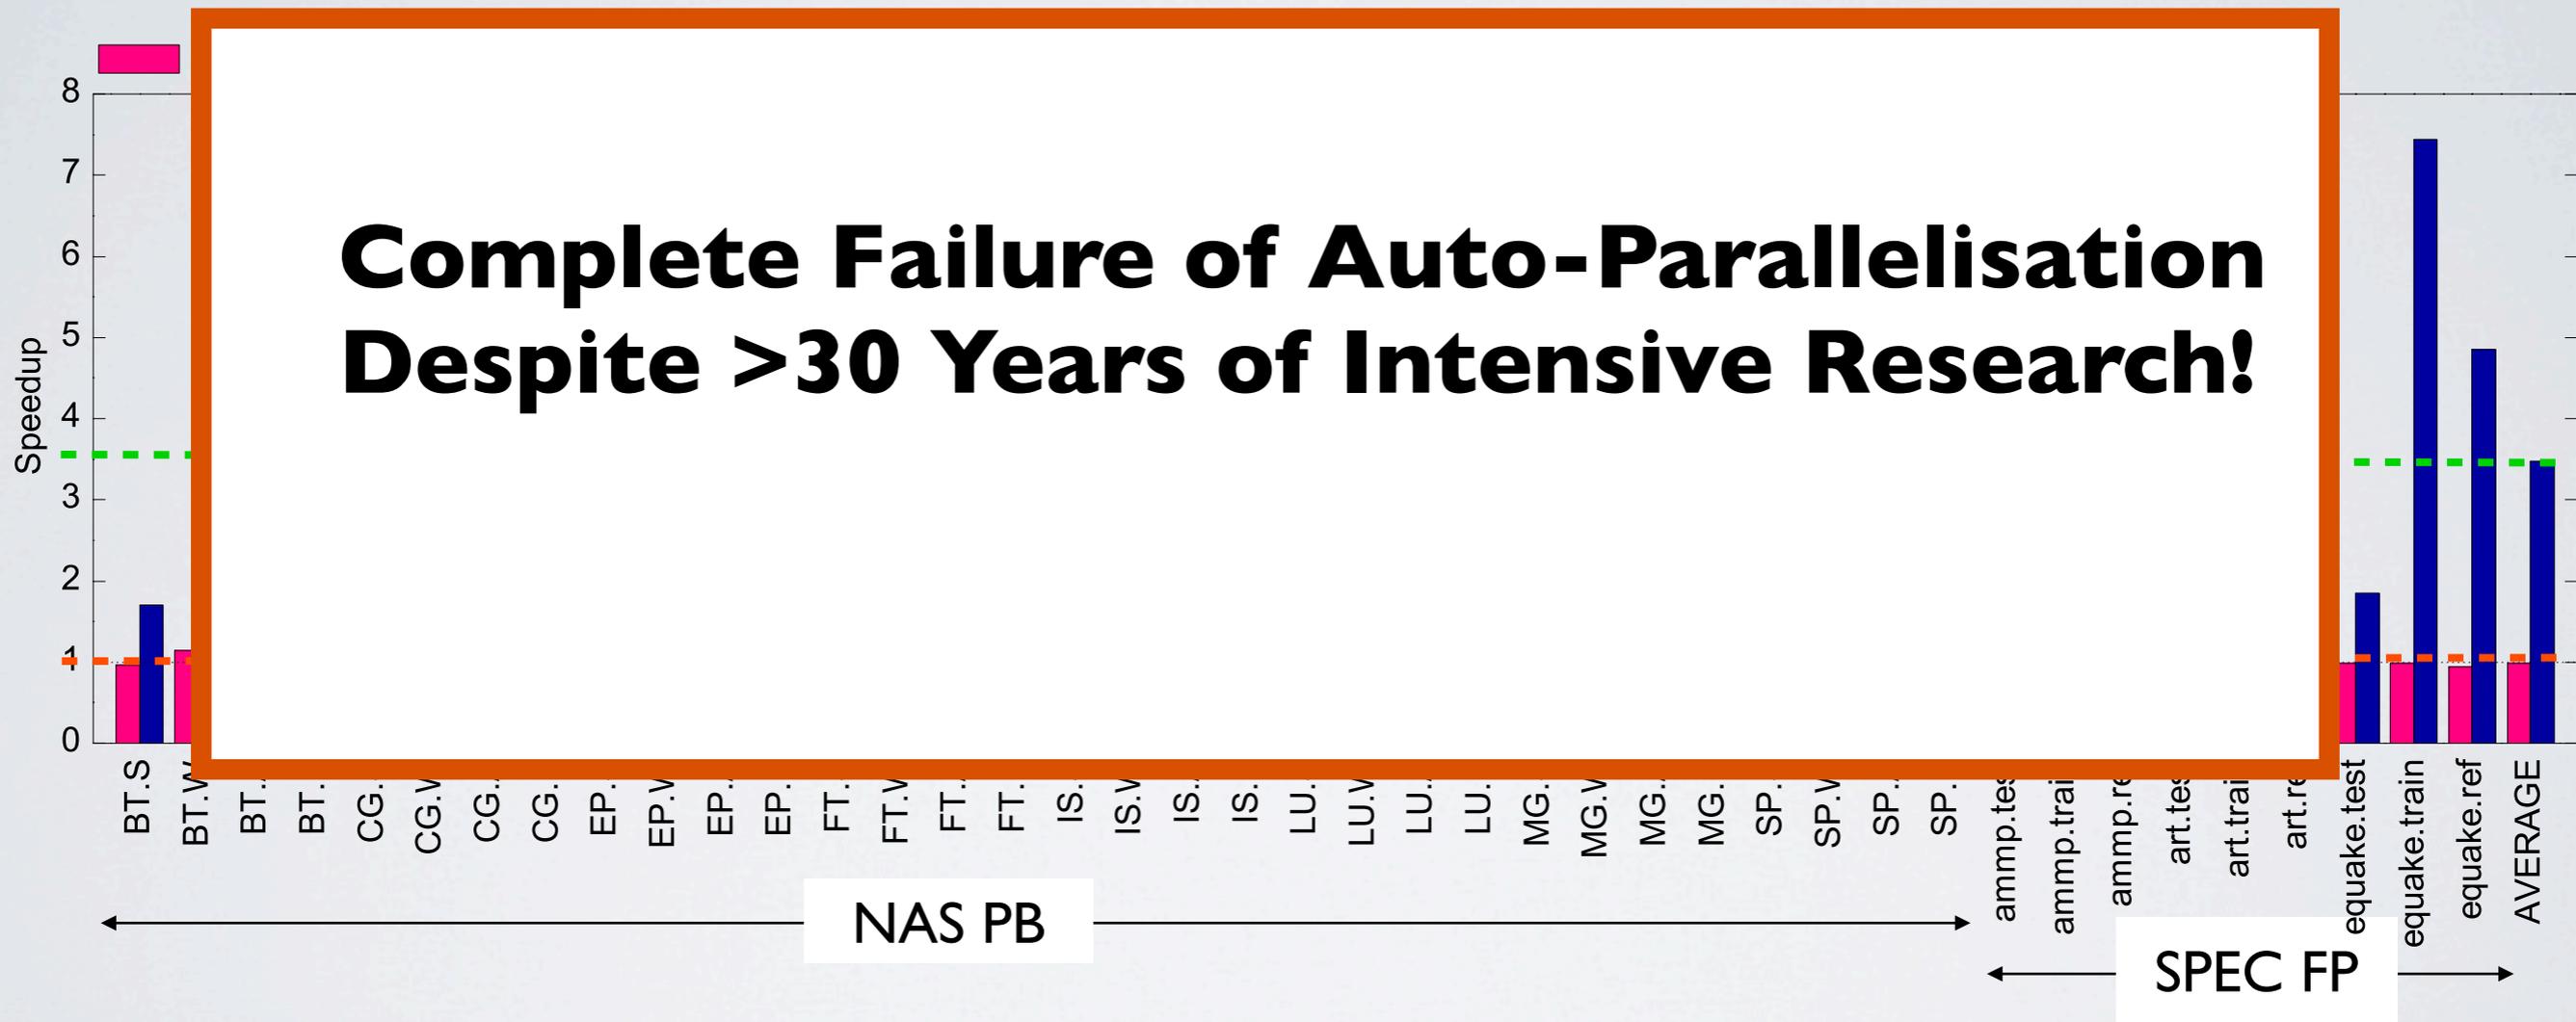
machines

- Parallelisation is **not** a one-off activity: Need to parallelise each application for each new platform again



State-of-the-Art

**Complete Failure of Auto-Parallelisation
Despite >30 Years of Intensive Research!**



NAS NPB 2.3 OMP-C and SPEC CFP2000
 2 Quad-cores (8 cores in total) Intel Xeon X5450 @ 3.00GHz
 Intel icc 10.1 -O2 -xT -axT -ipo



Observations

- Static Dependence Analysis Doesn't Work
 - ➔ Part 1: Profile Directed Dependence Analysis
- Mapping of Parallelism is Really Hard
 - ➔ Part 2: Machine Learning Based Mapping
- Parallelising FOR Loops is Not Enough
 - ➔ Part 3: Extraction of Pipeline Parallelism



PART I: PROFILE DIRECTED DEPENDENCE ANALYSIS



Motivating Example

SPEC quake (~75% of total exec. time)

```
for (i = 0; i < nodes; i++) {
  Anext = Aindex[i];
  Alast = Aindex[i + 1];

  sum0 = A[Anext][0][0]*v[i][0] +
         A[Anext][0][1]*v[i][1] +
         A[Anext][0][2]*v[i][2];
  sum1 = ...

  Anext++;
  while (Anext < Alast) {
    col = Acol[Anext];

    sum0 += A[Anext][0][0]*v[col][0] +
            A[Anext][0][1]*v[col][1] +
            A[Anext][0][2]*v[col][2];
    sum1 += ...

    w[col][0] += A[Anext][0][0]*v[i][0] +
                A[Anext][1][0]*v[i][1] +
                A[Anext][2][0]*v[i][2];
    w[col][1] += ...
    Anext++;
  }
  w[i][0] += sum0;
  w[i][1] += ...
}
```

- Static analysis fails to detect any parallelism
- Problems :
 - indirect array accesses 
 - compl. array reductions 
 - variable iteration count 
 - pointer aliasing
 - dynamic memory allocation
- But: Loop is parallel for all legal data inputs!

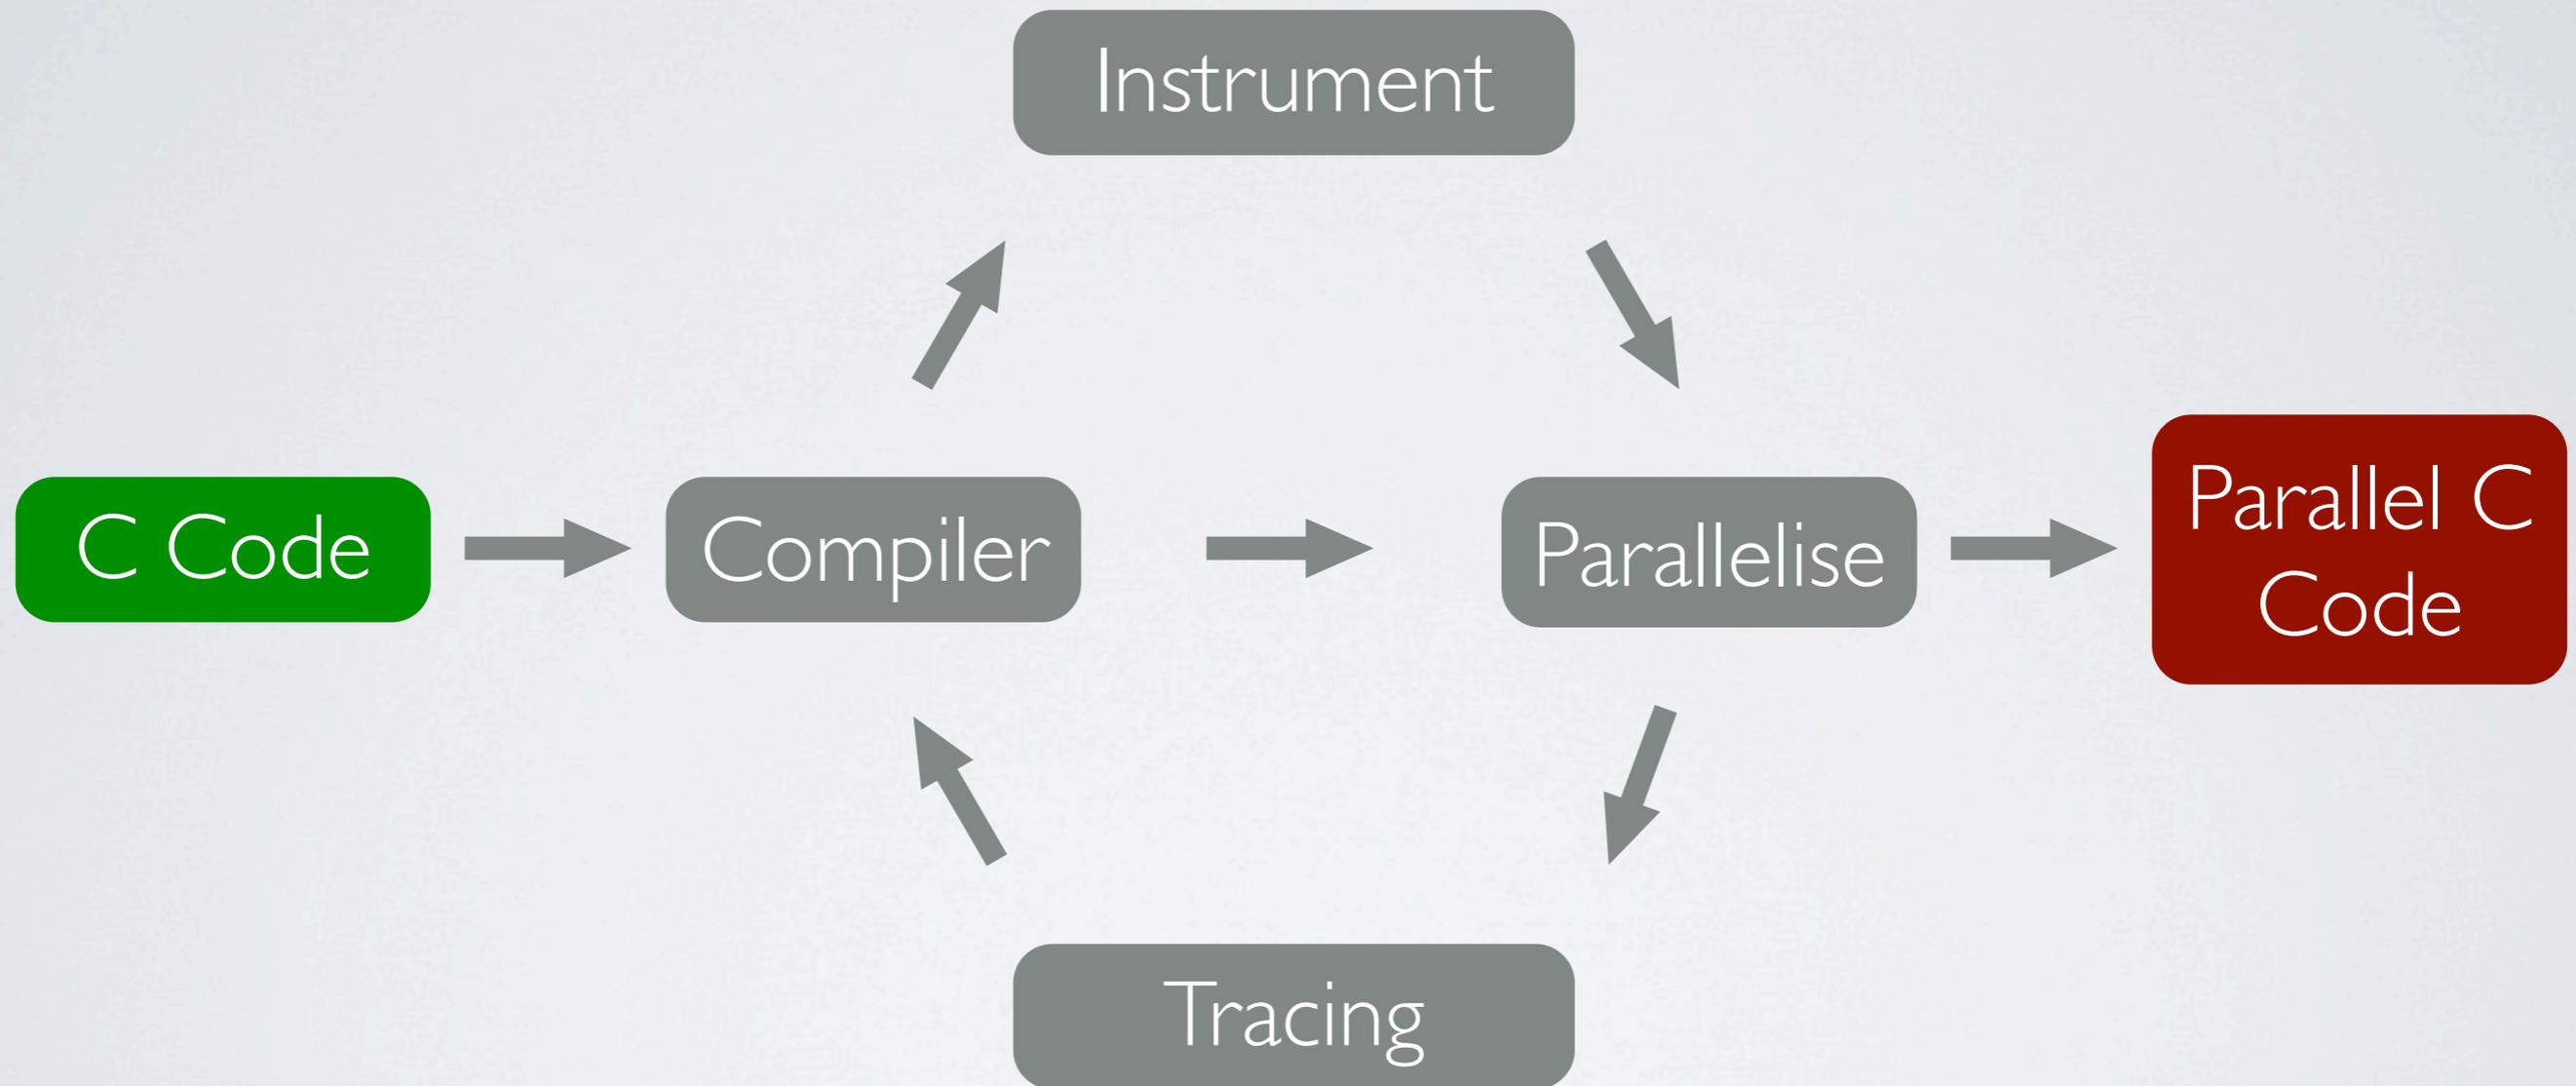


Profile Driven Parallelism Detection

- Use of profiling to capture data and control flow
 - Directly observe dependences → accuracy
- But: Need to solve two important problems
 - No general correctness proof
 - May have missed dependences
 - Assisted user validation → **semi**-automatic
 - Use low-level profiling information in compiler?
 - Instrumentation of intermediate representation
 - No actual ISA idiosyncracies



Approach





Approach

- Instrument using high-level intermed. representation
 - Access to source-level information (memory accesses, loops, induction/reduction vars)
- Avoids ISA obfuscation
- Execute natively
 - Generates data and control flow traces
- Straightforward back-annotation
 - References to symbol table, IR nodes
- **Combination with conventional static analyses**



PART 2: MACHINE LEARNING BASED MAPPING



Motivating Example

NAS CG

```
#pragma omp for reduction(+:sum) private(d)
for (j=1; j <= lastcol-firstcol-1; j++) {
    d = x[j] - r[j];
    sum = sum + d * d;
}
```

OpenMP

	Platform	
Scheduling	Cell BE	Intel Xeon
STATIC	slowdown	2.3x
DYNAMIC	slowdown	slowdown

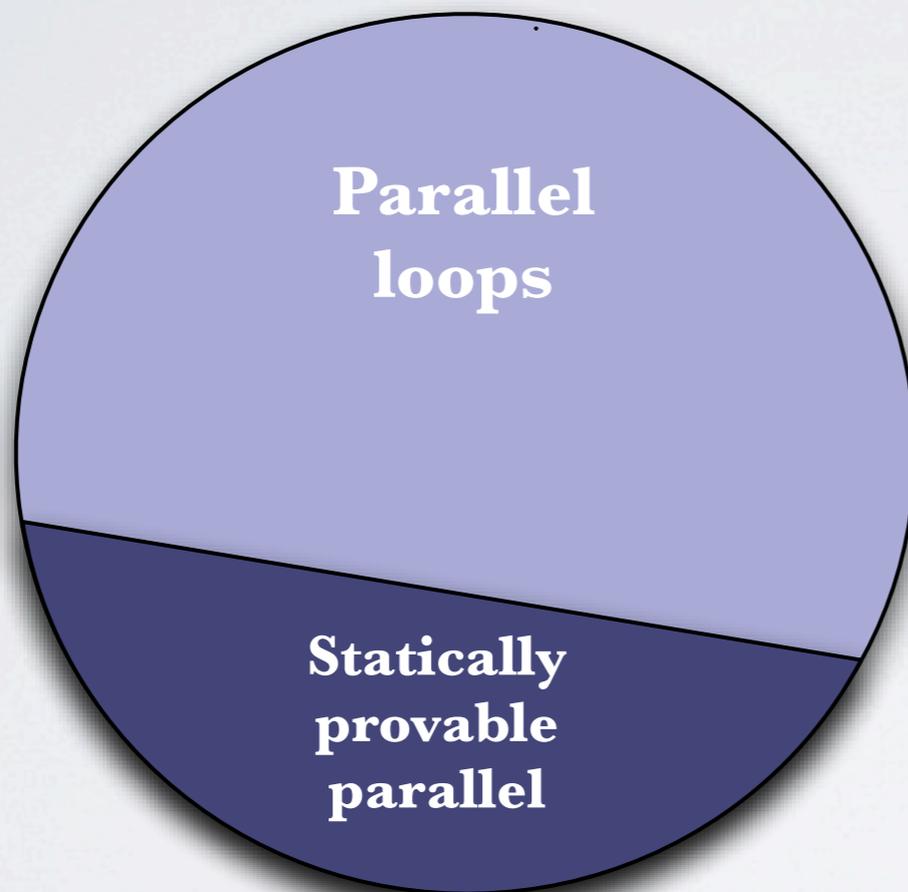


Problem Statement

**Parallel
loops**

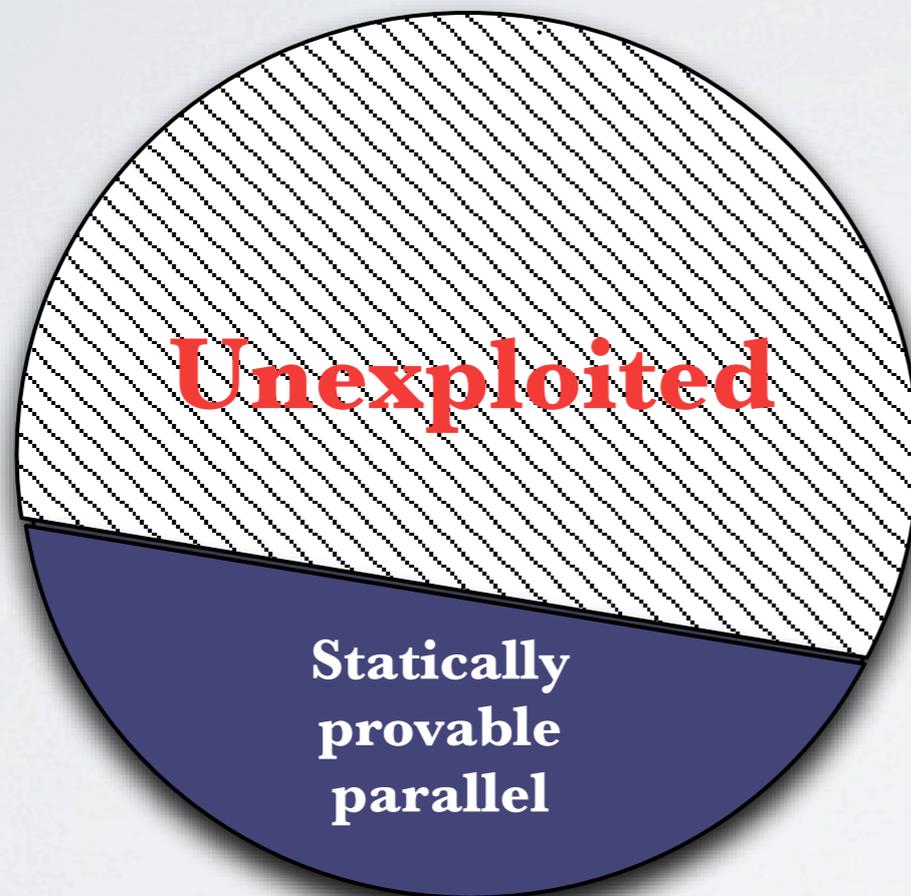


Problem Statement



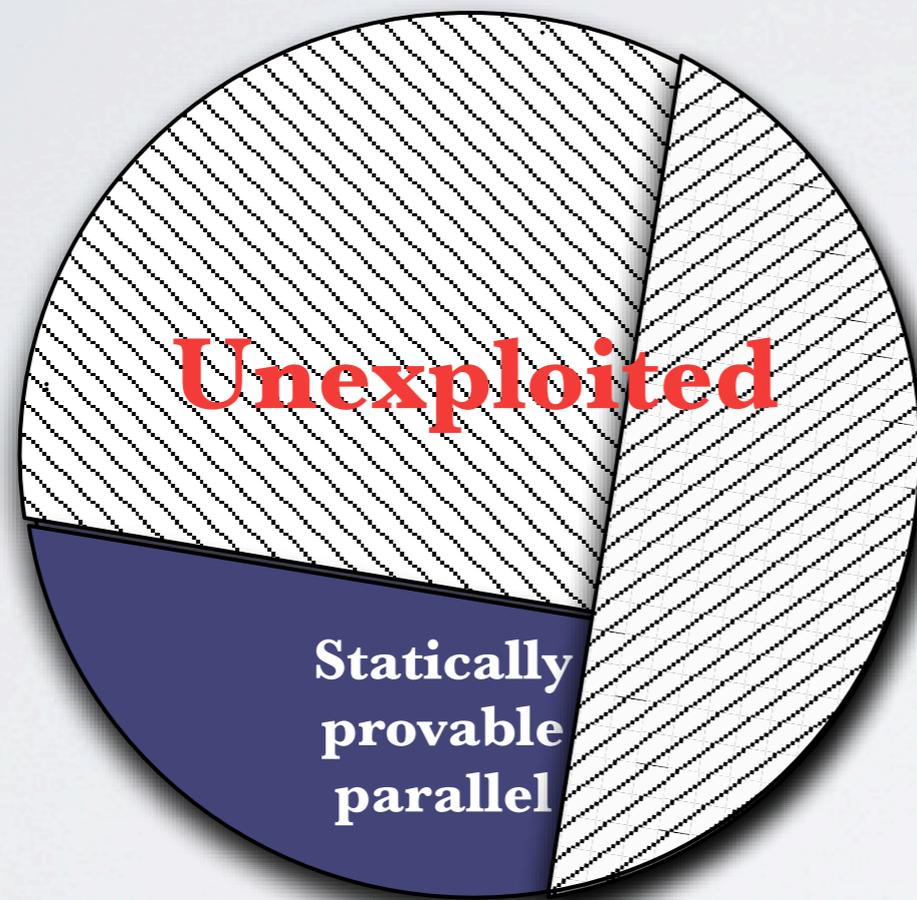


Problem Statement





Problem Statement

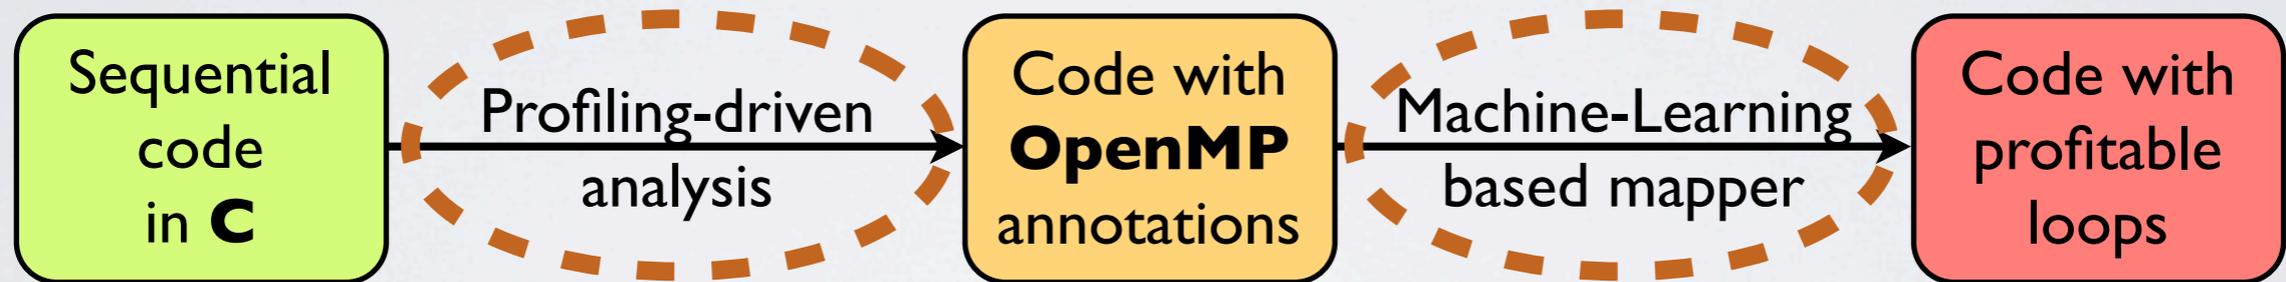


Profitable

A green arrow pointing from the word 'Profitable' towards the 'Unexploited' segment of the pie chart.

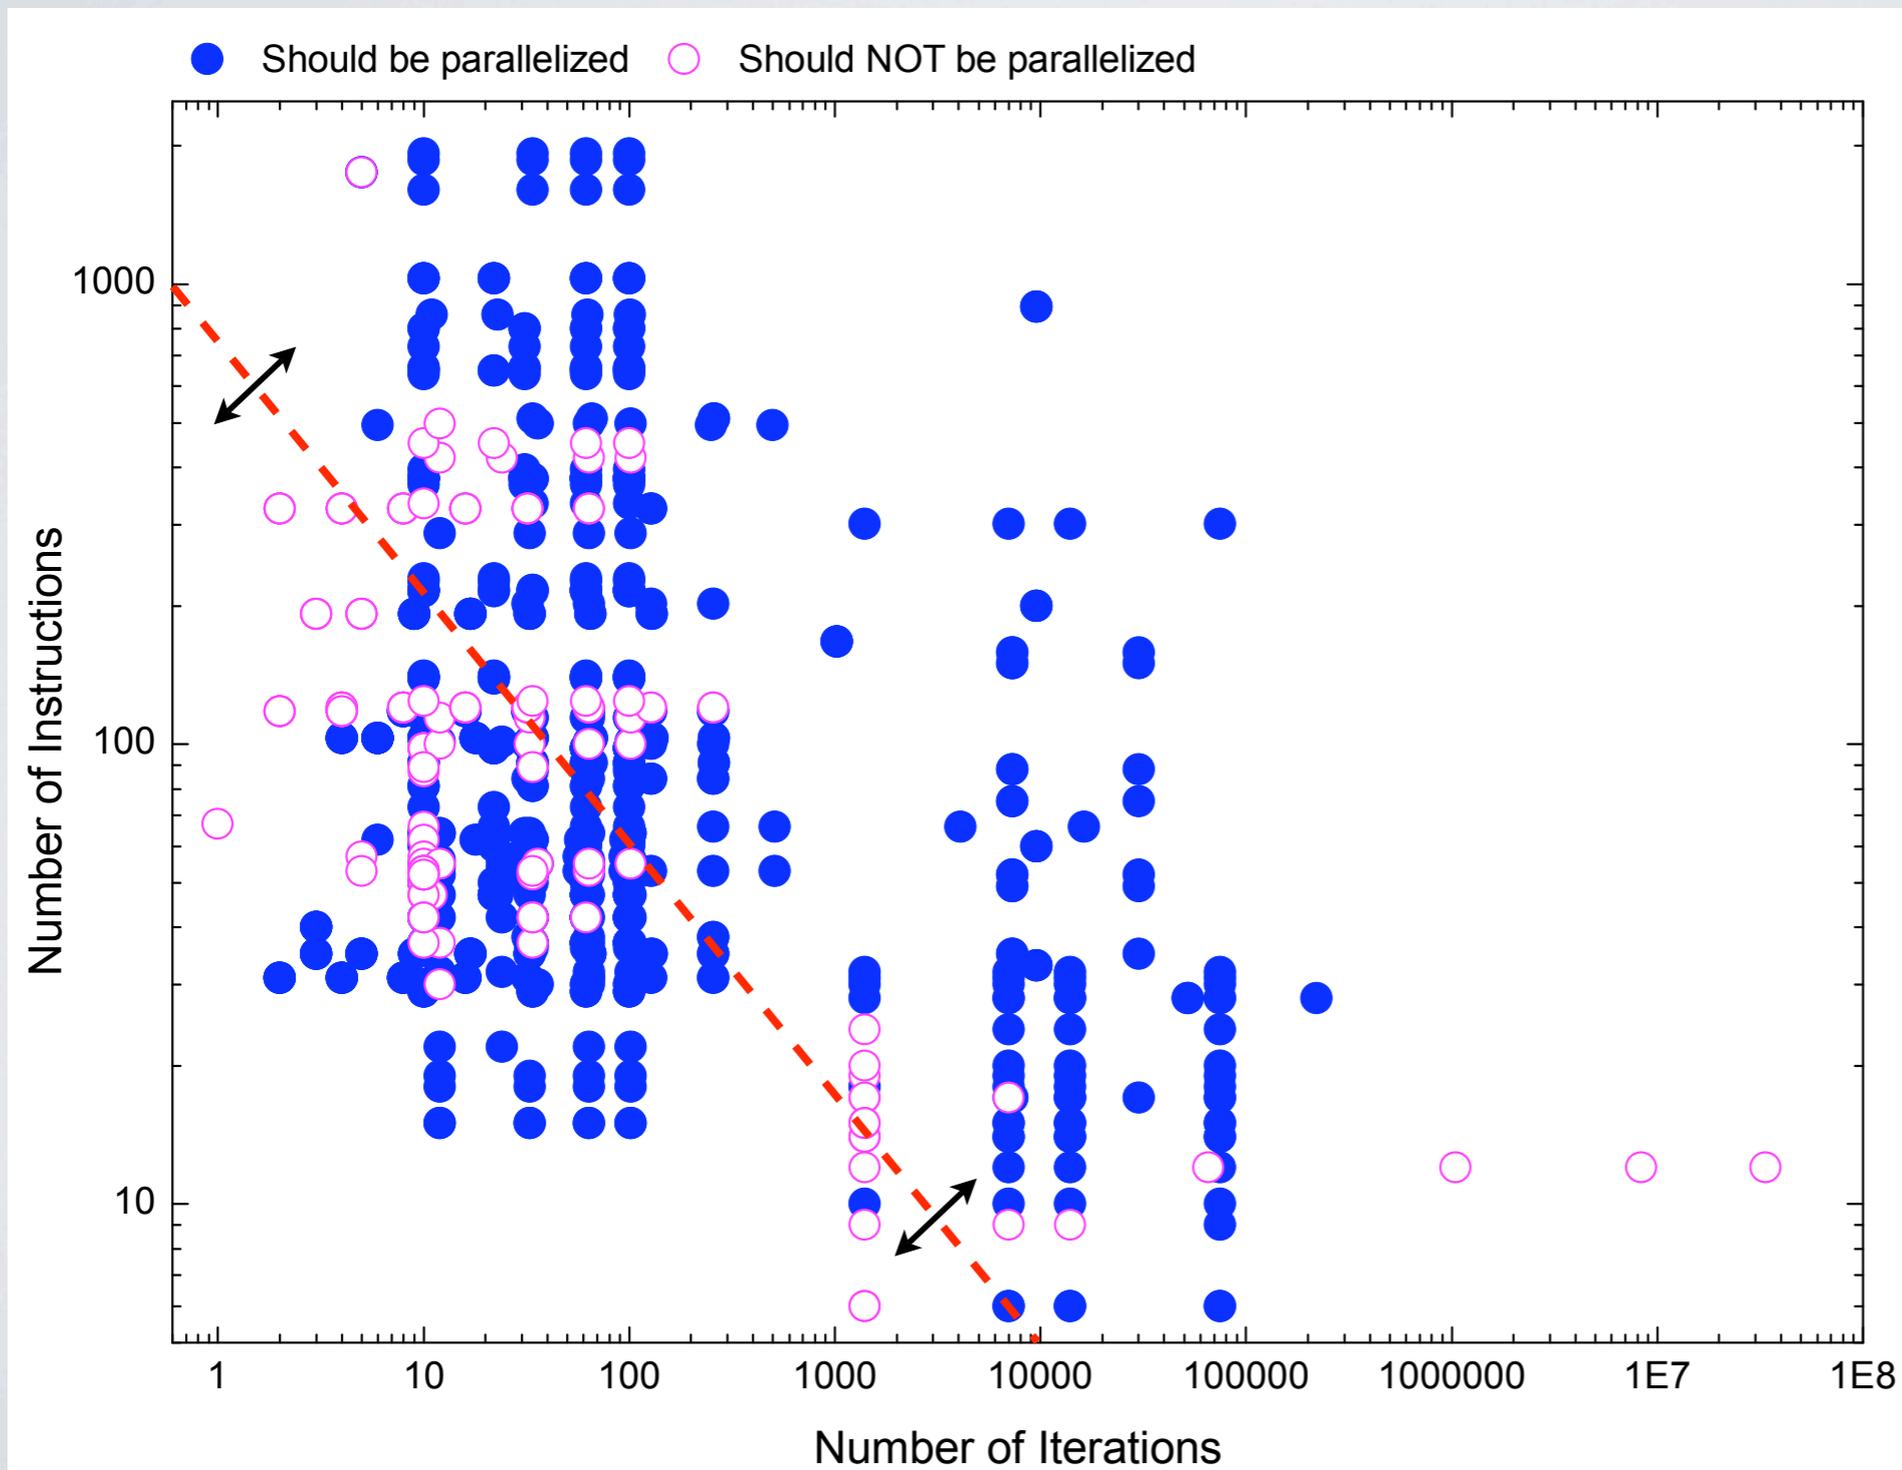


Holistic Approach: Detection & Mapping





Conventional Mapping Heuristics



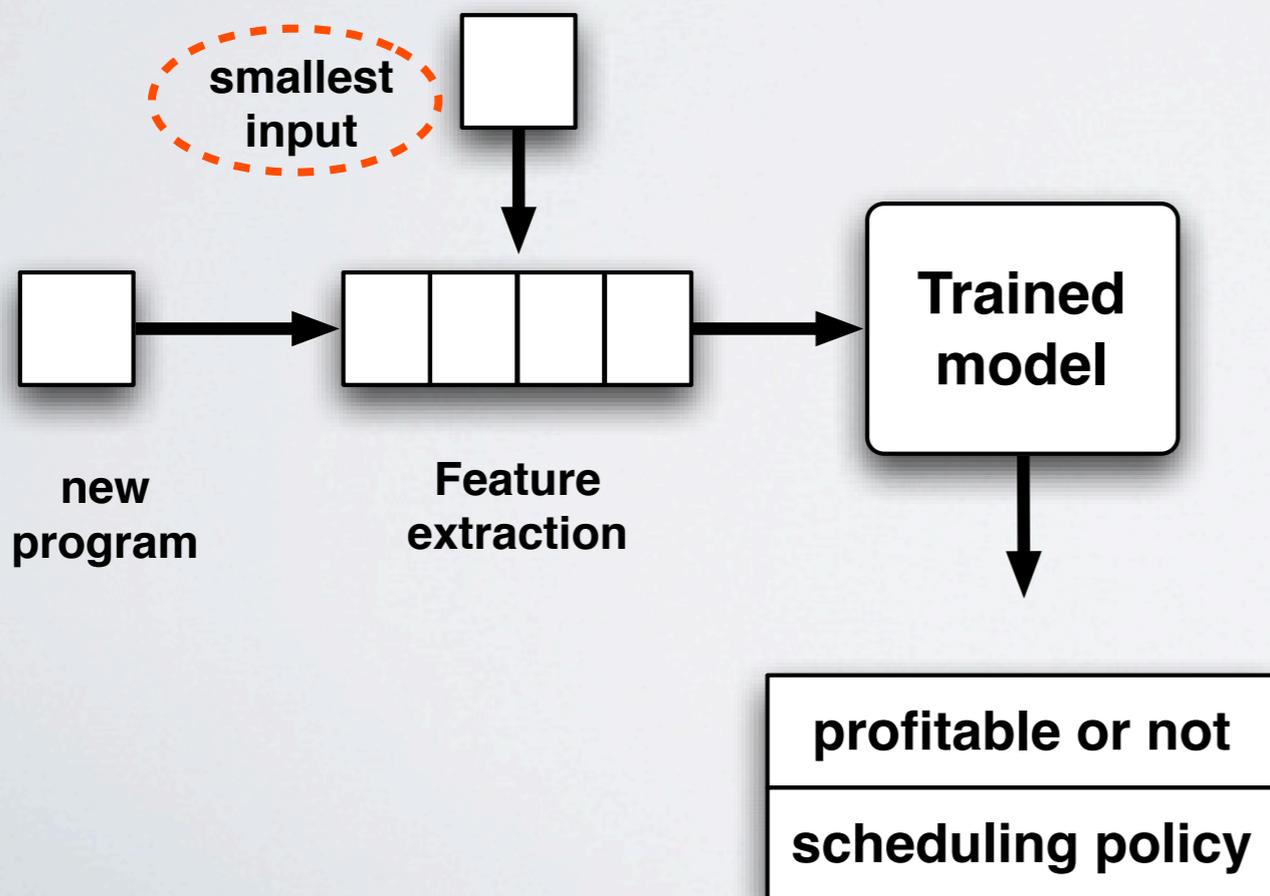
Static not suitable for separation.

Linear models not adequate either!



Machine Learning Based Mapping

- **Off-line** learning
- Predict using **smallest** input



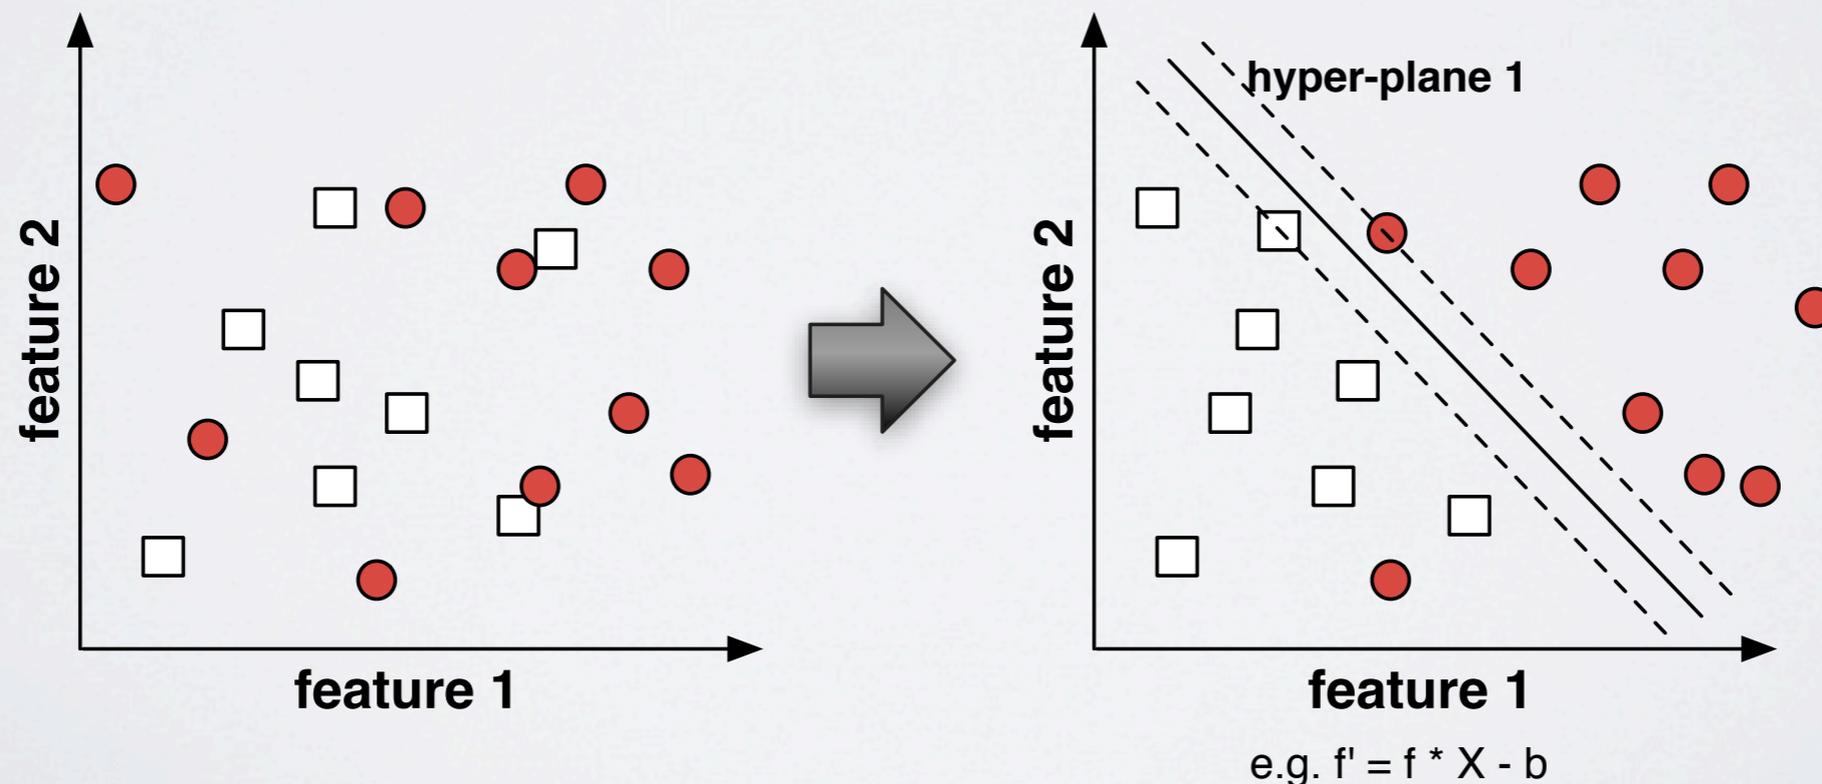
Features	
Static	IR instruction count IR Load/Store count IR Branch count Loop iteration count
Dynamic	Data access count Instruction count Branch count



Predictive Modelling

Support Vector Machine (SVM)

- Decide **(i)** profitability, **(ii)** loop scheduling
- Hyperplanes in **transformed higher-dimensional** space
- Non-linear & multi-class extensions





Experimental Evaluation

- 2 sets of applications
 - NAS Parallel Benchmarks 2.3
 - SPEC FP2000
- Sequential code in C
- Manually parallelized using OpenMP by expert programmers
- Use of multiple input datasets



Parallelism Coverage

- Profile-driven: almost **no lost opportunities**

Application	FP	FN
bt	0	0
cg	0	0
ep	0	0
ft	0	0
is	0	0
lu	0	0
mg	0	3
sp	0	0
equake	0	0
art	0	0
ammp	0	1



Parallelism Coverage

- *MG*: 3 loops never execute for all inputs
- *ammp*: critical loops require reshaping & locking

Application	FP	FN
bt	0	0
cg	0	0
ep	0	0
ft	0	0
is	0	0
lu	0	0
mg	0	3
sp	0	0
equake	0	0
art	0	0
ammp	0	1



Parallelism Coverage

- ICC finds many parallel loops

	icc		Profile-driven		Manual	
Applicatio	#loops(%cov)	#loops(%cov)	#loops(%cov)	#loops(%cov)	#loops(%cov)	#loops(%cov)
bt	72	(18.6%)	205	(99.9%)	54	(99.9%)
cg	16	(1.10%)	28	(93.1%)	22	(93.1%)
ep	6	(<1%)	8	(99.9%)	1	(99.9%)
ft	3	(<1%)	37	(88.2%)	6	(88.2%)
is	8	(29.4%)	9	(28.5%)	1	(27.3%)
lu	88	(65.9%)	54	(99.7%)	29	(81.5%)
mg	9	(4.70%)	48	(77.7%)	12	(77.7%)
sp	178	(88.0%)	287	(99.6%)	70	(61.8%)
equake	29	(23.8%)	69	(98.1%)	11	(98.0%)
art	16	(30.0%)	31	(85.6%)	5	(65.0%)
ammp	43	(<1%)	21	(1.40%)	7	(84.4%)



Parallelism Coverage

- BUT: low sequential time coverage
- ICC: majority of loops **too short** to be profitable

	icc		Profile-driven		Manual	
Applicatio	#loops(%cov)	#loops(%cov)	#loops(%cov)	#loops(%cov)	#loops(%cov)	#loops(%cov)
bt	72	(18.6%)	205	(99.9%)	54	(99.9%)
cg	16	(1.10%)	28	(93.1%)	22	(93.1%)
ep	6	(<1%)	8	(99.9%)	1	(99.9%)
ft	3	(<1%)	37	(88.2%)	6	(88.2%)
is	8	(29.4%)	9	(28.5%)	1	(27.3%)
lu	88	(65.9%)	54	(99.7%)	29	(81.5%)
mg	9	(4.70%)	48	(77.7%)	12	(77.7%)
sp	178	(88.0%)	287	(99.6%)	70	(61.8%)
equake	29	(23.8%)	69	(98.1%)	11	(98.0%)
art	16	(30.0%)	31	(85.6%)	5	(65.0%)
ammp	43	(<1%)	21	(1.40%)	7	(84.4%)



Parallelism Coverage

- Profile-driven: coverage close to manually parallelized

	icc		Profile-driven		Manual	
Applicatio	#loops(%cov)		#loops(%cov)		#loops(%cov)	
bt	72	(18.6%)	205	(99.9%)	54	(99.9%)
cg	16	(1.10%)	28	(93.1%)	22	(93.1%)
ep	6	(<1%)	8	(99.9%)	1	(99.9%)
ft	3	(<1%)	37	(88.2%)	6	(88.2%)
is	8	(29.4%)	9	(28.5%)	1	(27.3%)
lu	88	(65.9%)	54	(99.7%)	29	(81.5%)
mg	9	(4.70%)	48	(77.7%)	12	(77.7%)
sp	178	(88.0%)	287	(99.6%)	70	(61.8%)
equake	29	(23.8%)	69	(98.1%)	11	(98.0%)
art	16	(30.0%)	31	(85.6%)	5	(65.0%)
ammp	43	(<1%)	21	(1.40%)	7	(84.4%)



Parallelism Coverage

- ammp: we fail to parallelize the critical loop

	icc		Profile-driven		Manual	
Applicatio	#loops(%cov)		#loops(%cov)		#loops(%cov)	
bt	72	(18.6%)	205	(99.9%)	54	(99.9%)
cg	16	(1.10%)	28	(93.1%)	22	(93.1%)
ep	6	(<1%)	8	(99.9%)	1	(99.9%)
ft	3	(<1%)	37	(88.2%)	6	(88.2%)
is	8	(29.4%)	9	(28.5%)	1	(27.3%)
lu	88	(65.9%)	54	(99.7%)	29	(81.5%)
mg	9	(4.70%)	48	(77.7%)	12	(77.7%)
sp	178	(88.0%)	287	(99.6%)	70	(61.8%)
equake	29	(23.8%)	69	(98.1%)	11	(98.0%)
art	16	(30.0%)	31	(85.6%)	5	(65.0%)
ammp	43	(<1%)	21	(1.40%)	7	(84.4%)



Safety

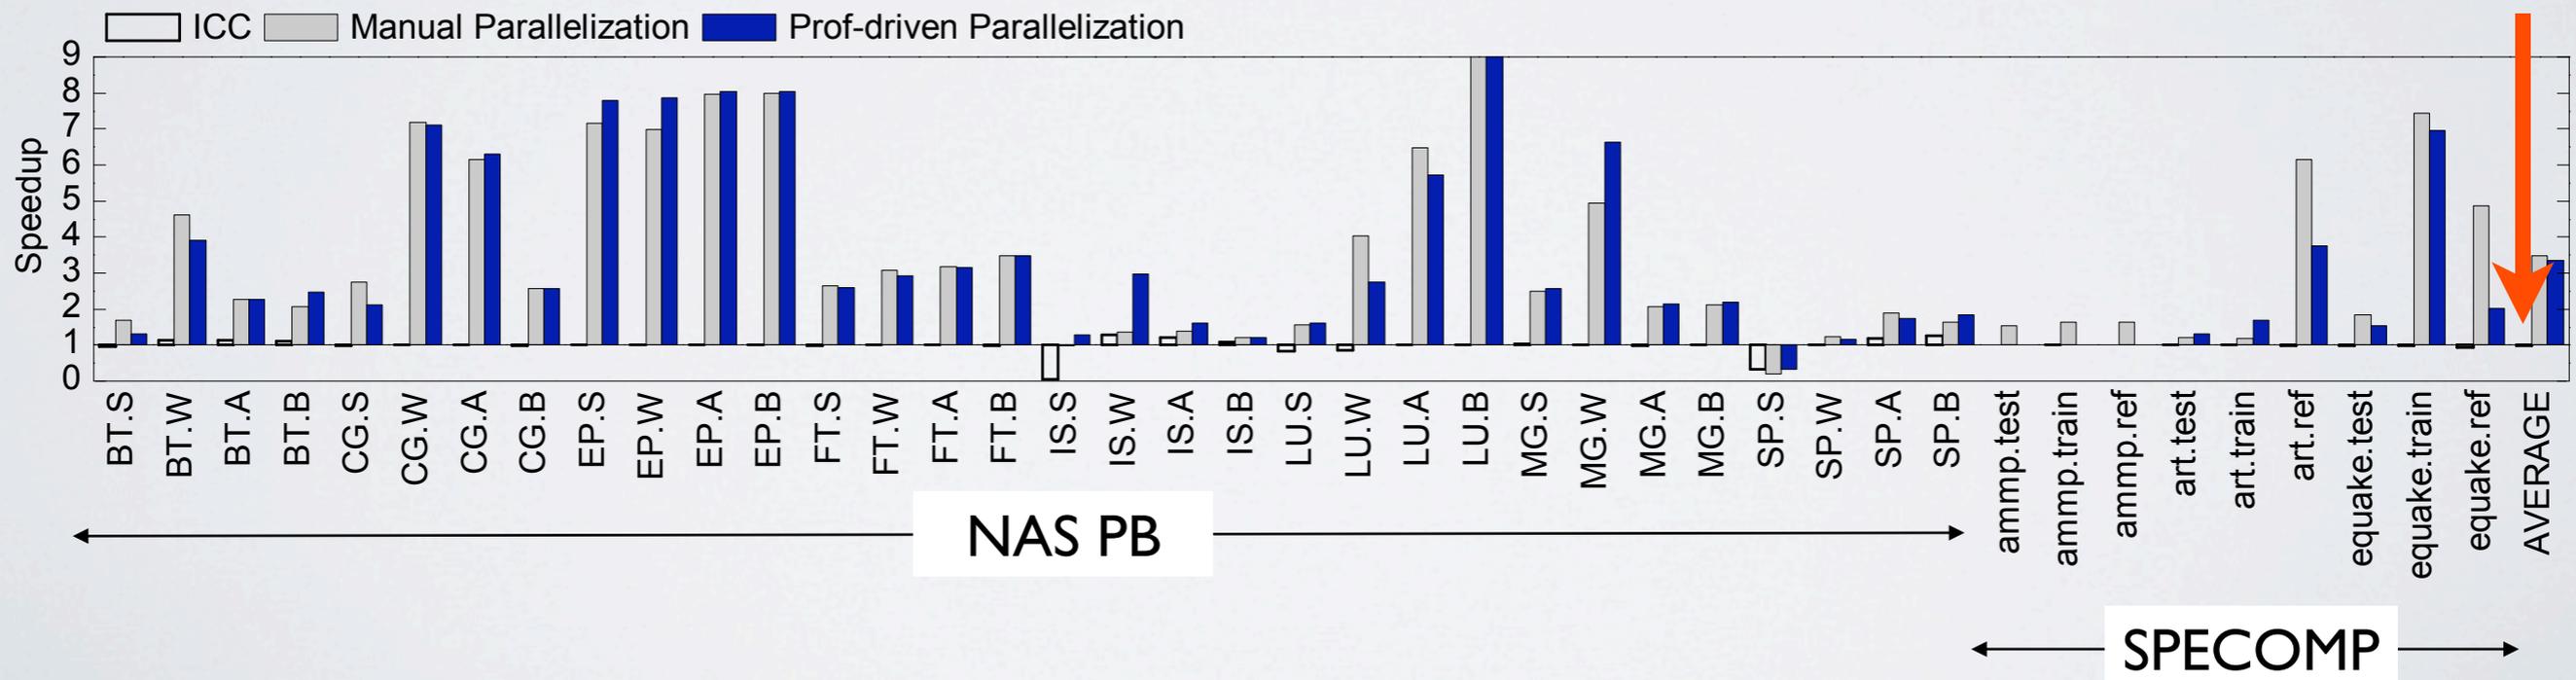
- Inherently unsafe, but surprisingly **no FP**
- Even when trained on the **smallest** dataset

Application	FP	FN
bt	0	0
cg	0	0
ep	0	0
ft	0	0
is	0	0
lu	0	0
mg	0	3
sp	0	0
equake	0	0
art	0	0
ammp	0	1



Speedup (Intel Xeon)

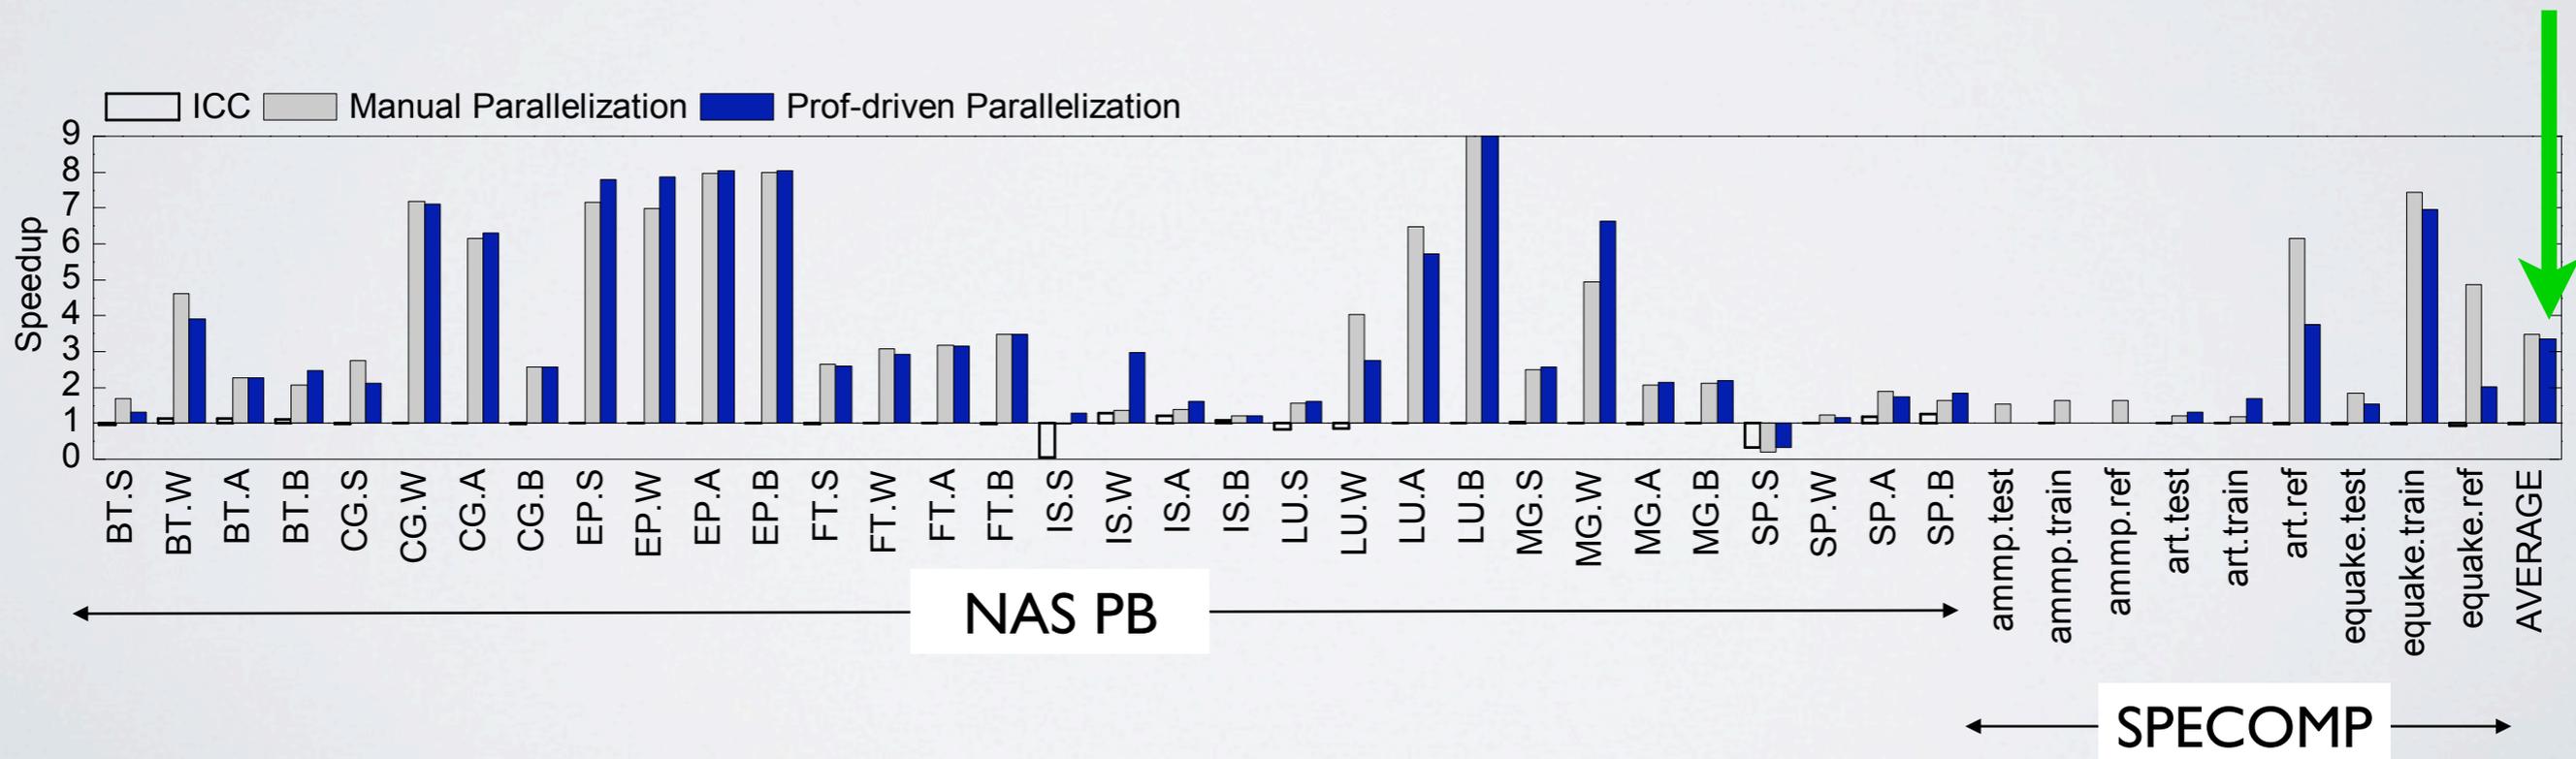
- Intel ICC **fails** to deliver any performance gain
- Even **slowdown** for some benchmarks





Speedup (Intel Xeon)

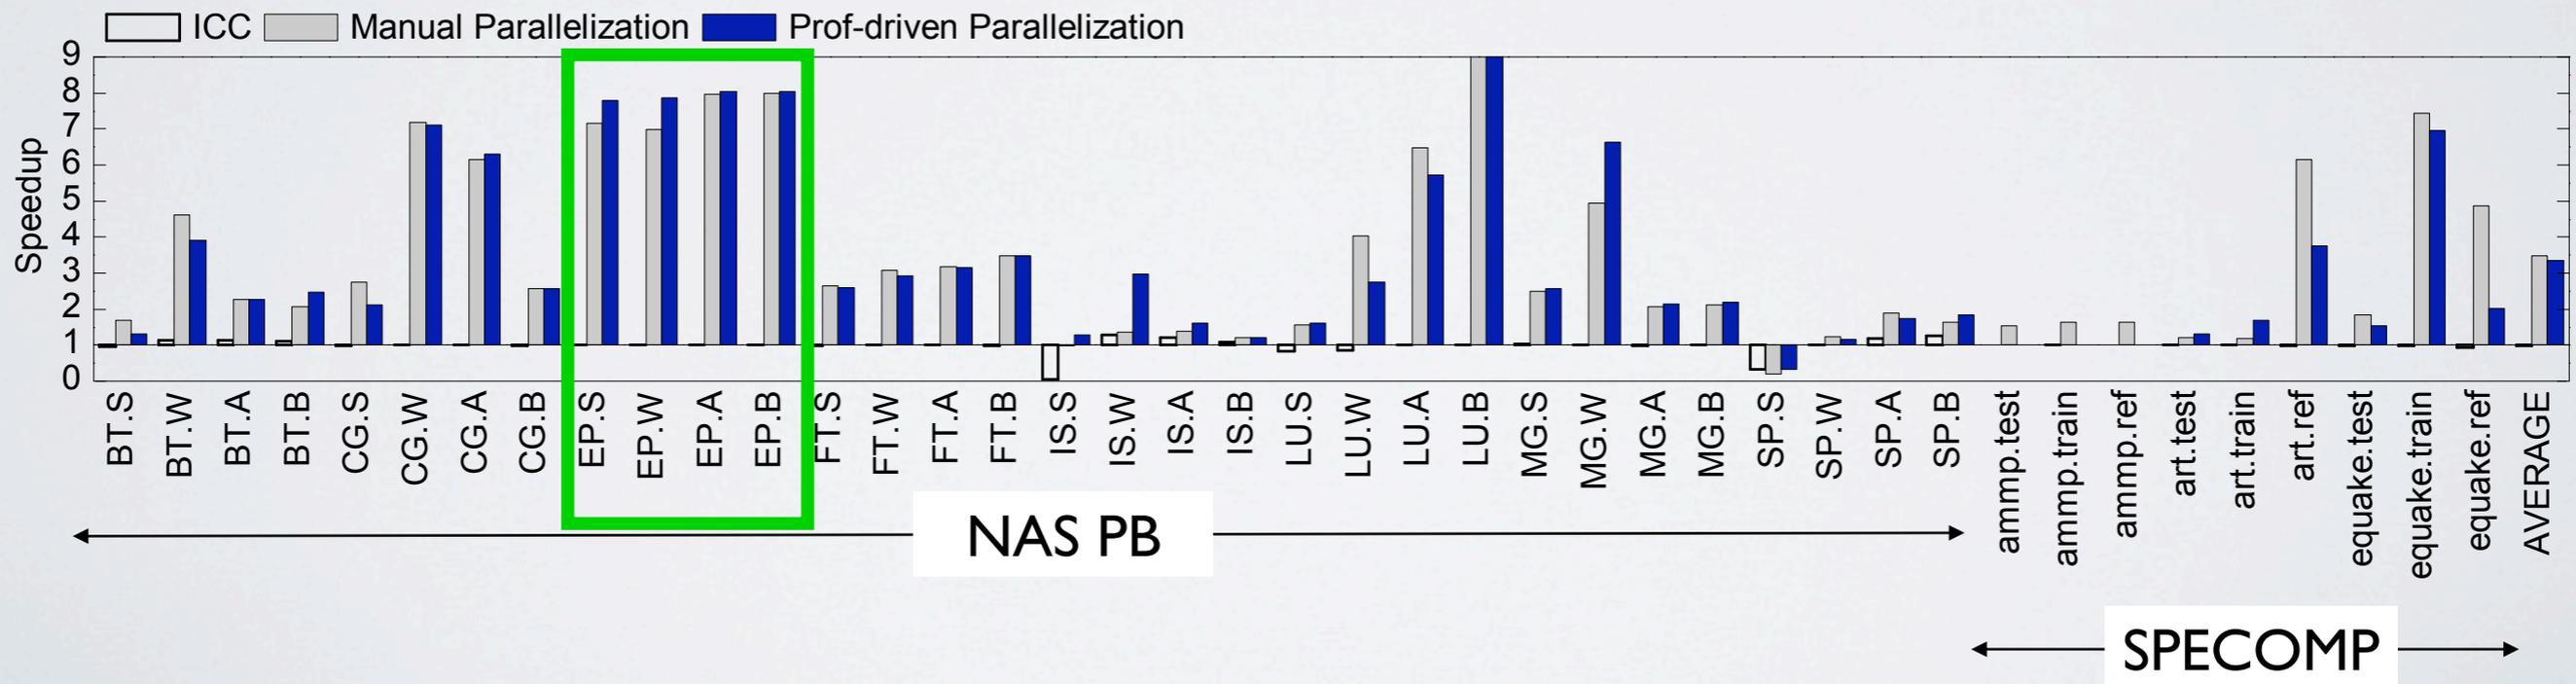
- Profile-driven parallelization achieves **96%** of the performance of manually parallelized benchmarks!





Speedup (Intel Xeon)

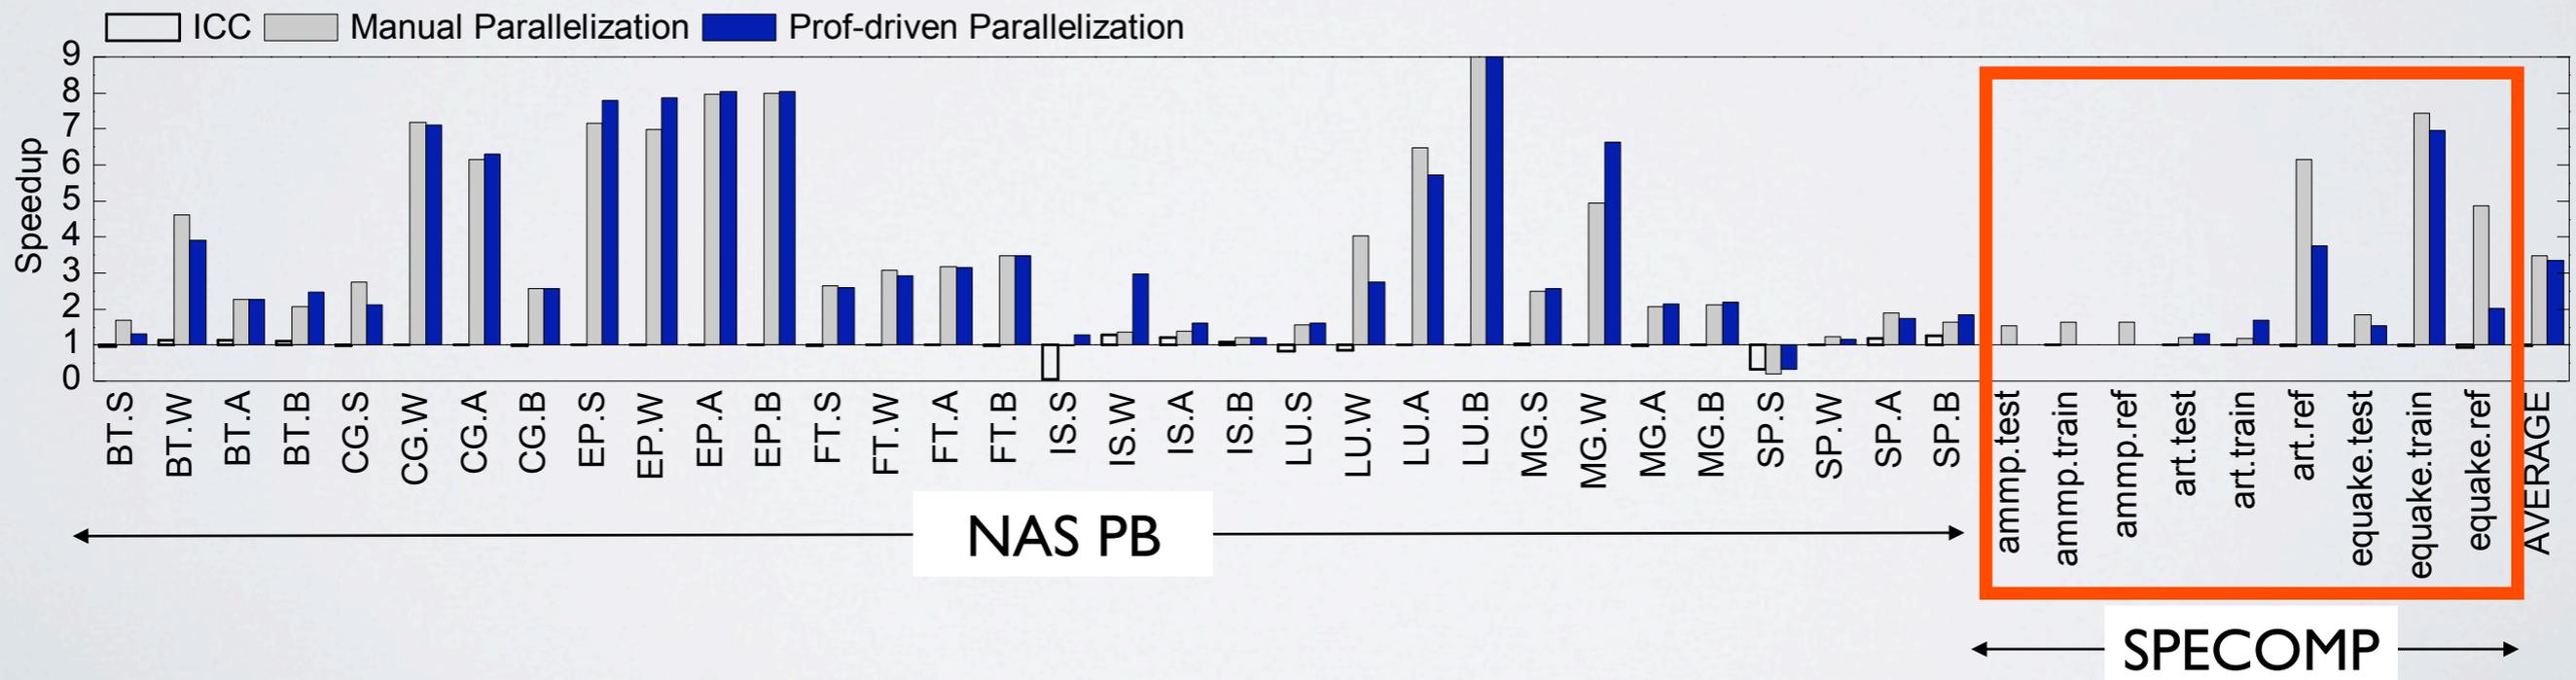
- EP is embarrassingly parallel, still ICC fails completely
- Profile-driven parallelisation detects critical loop





Speedup (Intel Xeon)

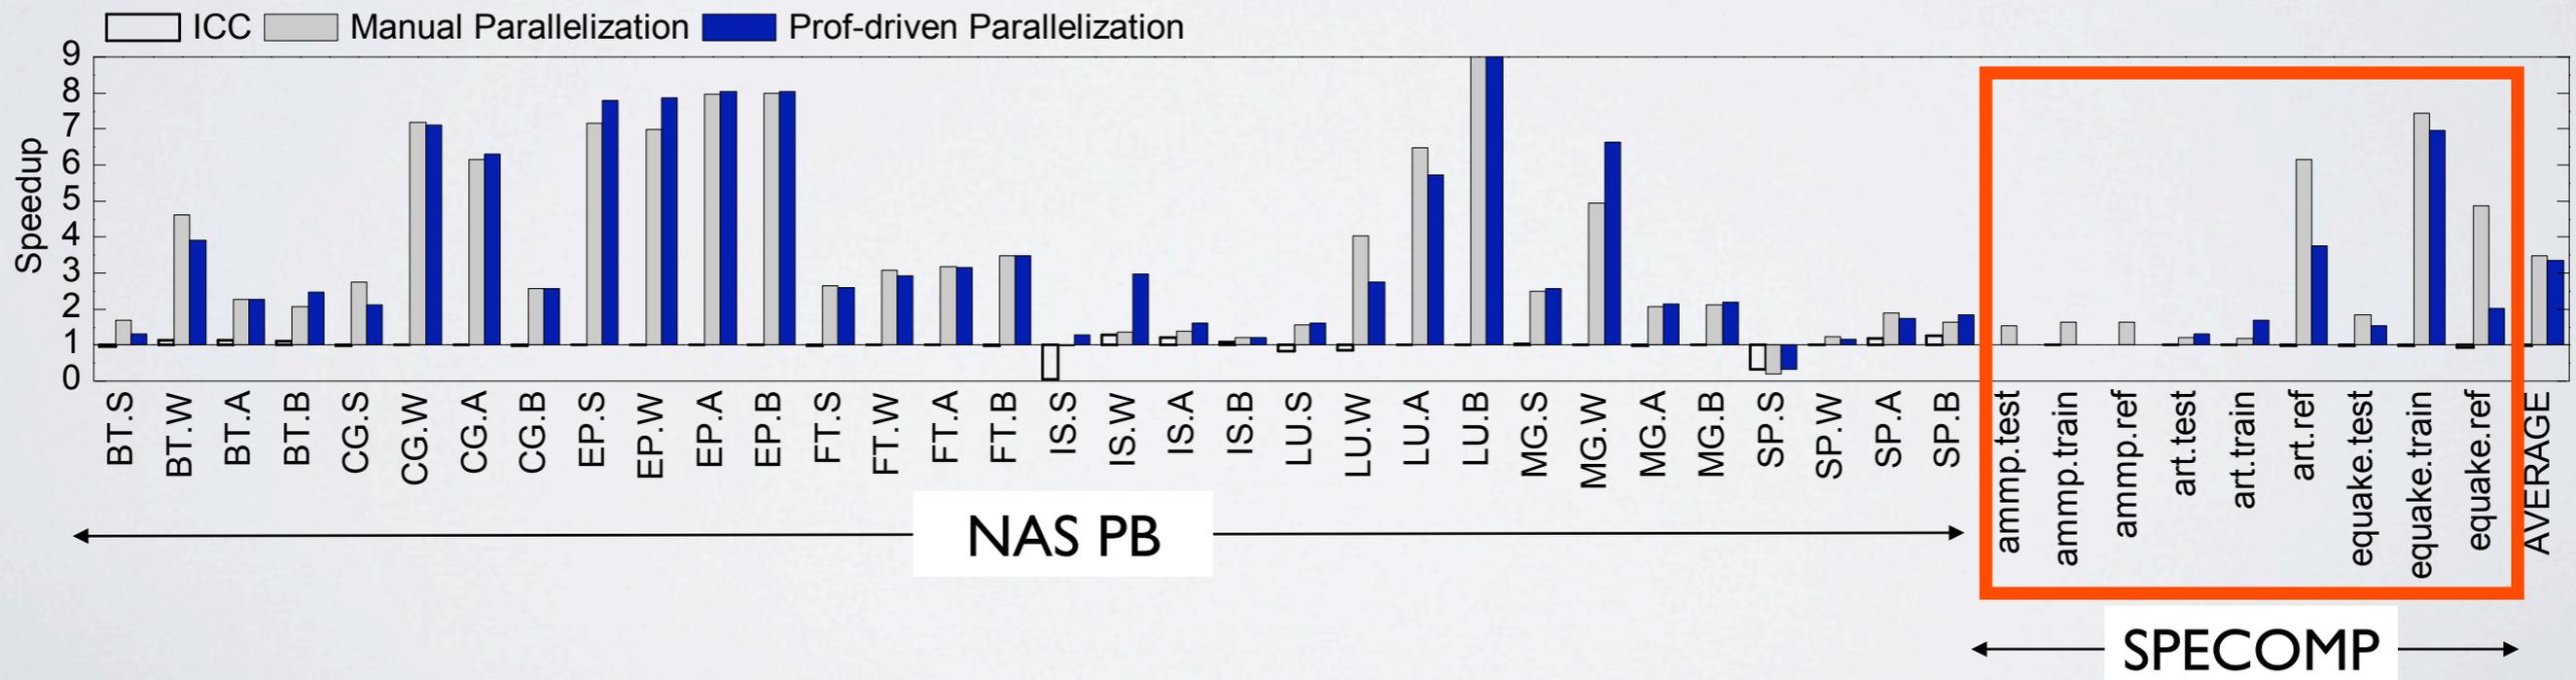
- SPECOMP 2001 benchmarks include additional sequential optimisations besides parallelisation





Speedup (Intel Xeon)

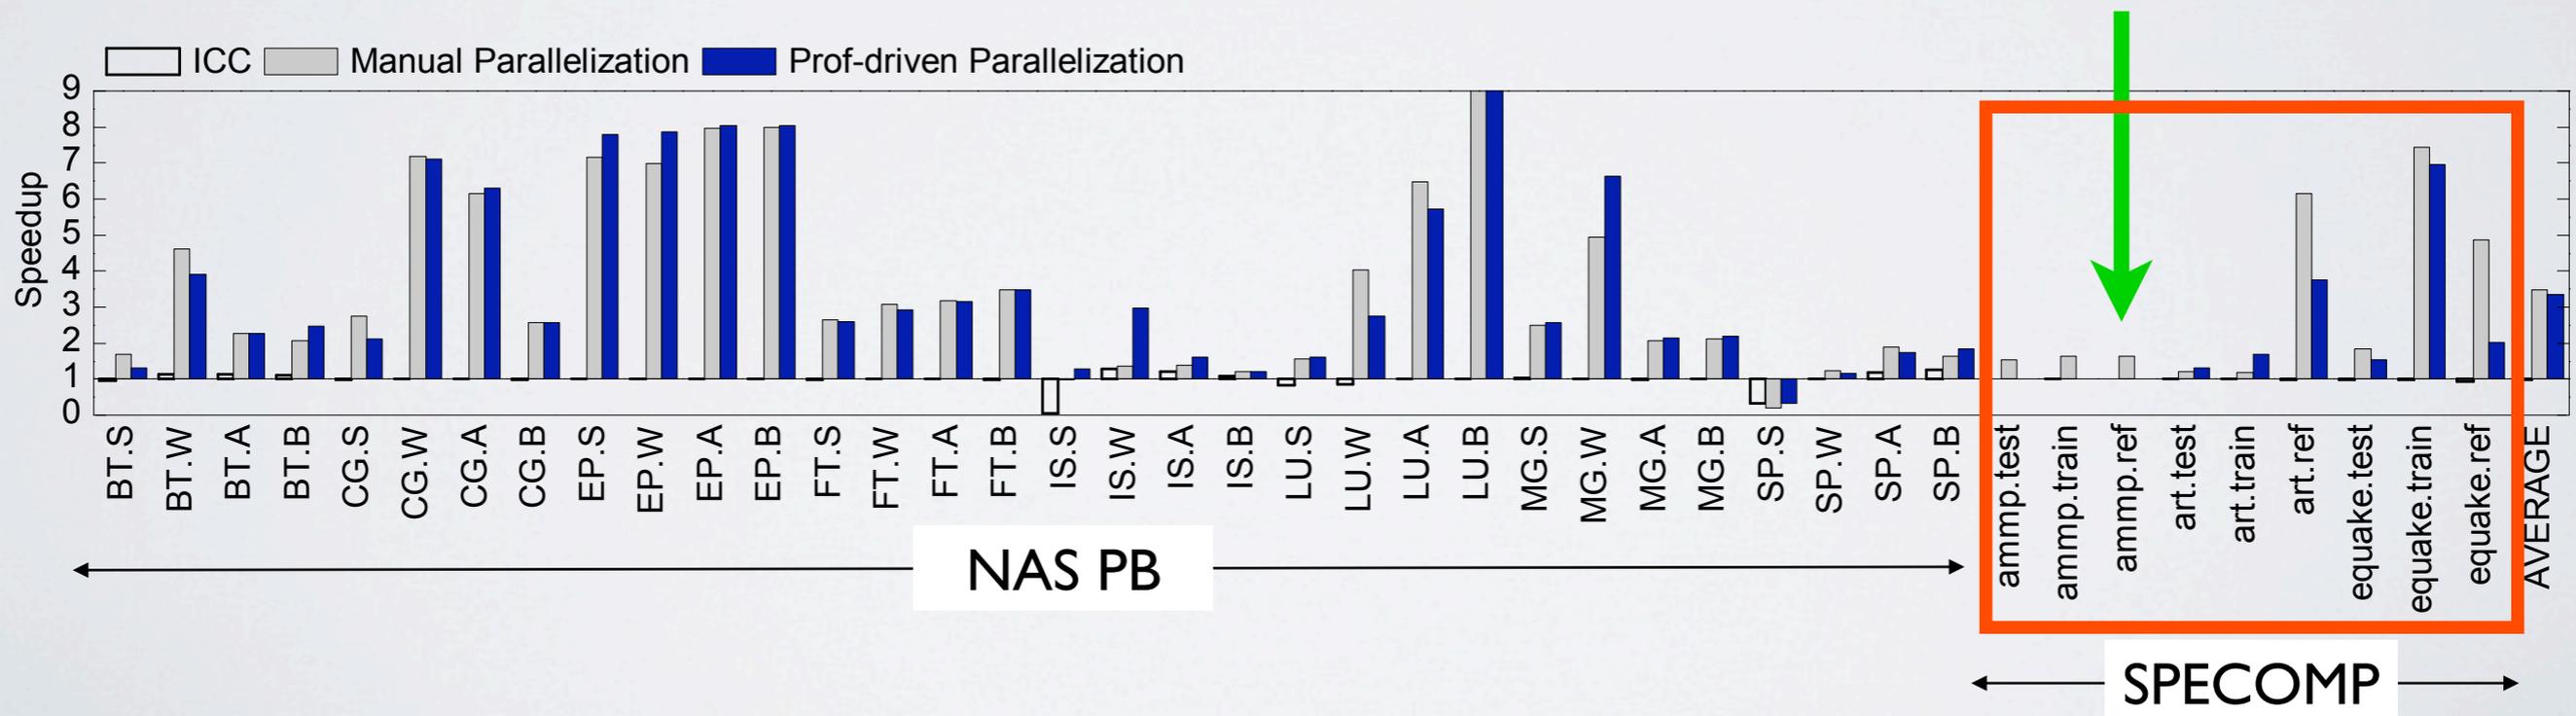
- SPECOMP single-threaded has average speedup of **2x** over SPECFP due to sequential optimisations





Speedup (Intel Xeon)

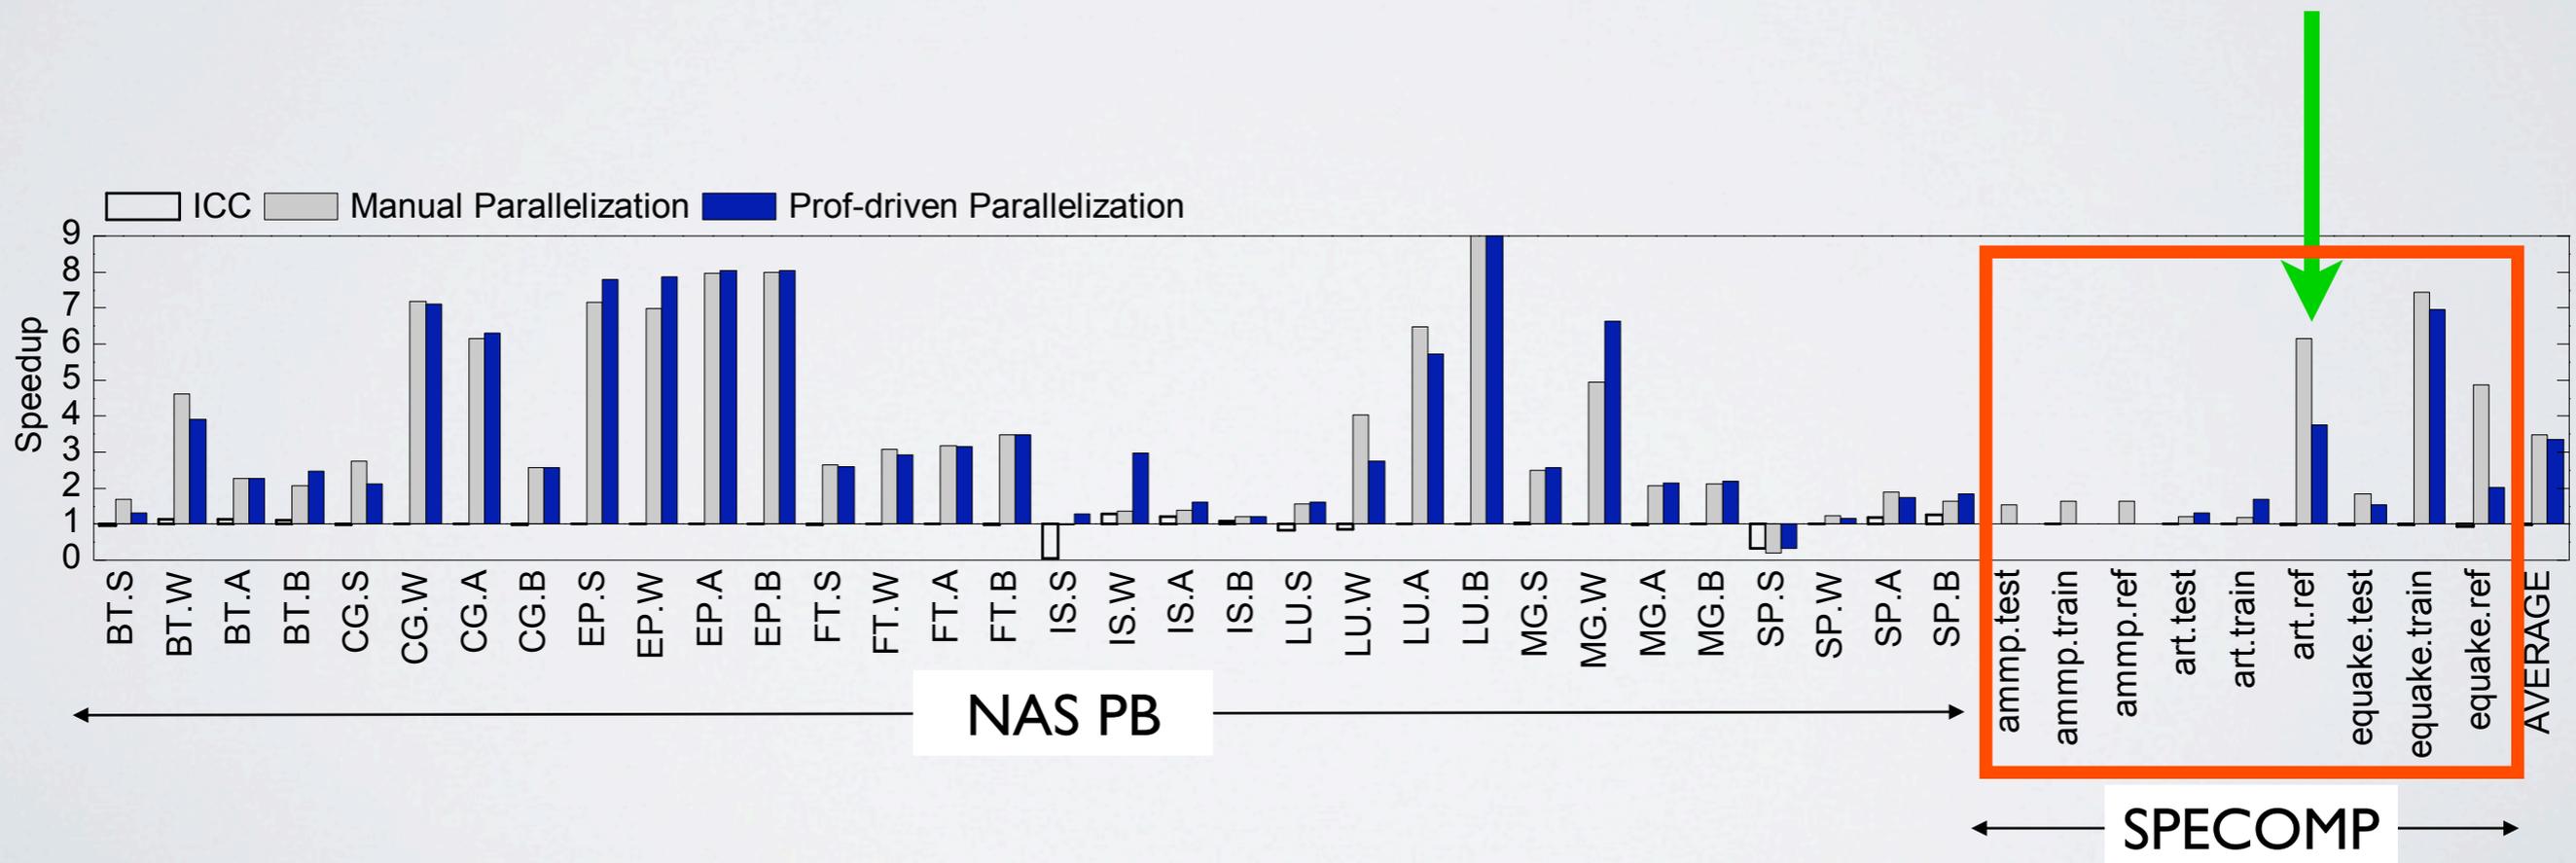
- ammp: critical loop not parallelised by profile-driven technique. Misclassification by ML.
- Manual parallelisation: **1.6x** on 8 cores





Speedup (Intel Xeon)

- SPECOMP.art **3.34x** with **1 thread**
- Profile-driven parallelisation delivers **4x** speedup without sequential optimisation





PART 3: EXTRACTION OF PIPELINE PARALLELISM



Observations

- There is more parallelism available beyond parallel FOR loops
 - Programmers exploit coarse-grained parallelism routinely
 - Auto-parallelising compilers don't!
- Parallel Design Patterns
 - Static: Pipelines, Task Graphs
 - Dynamic: Task Farms, Divide & Conquer, ...
- Serious Programme Restructuring Required
 - Let's Do It!



Motivating Example

EEMBC mp3player Algorithmic components

```
1 while(end) {
2   /* ...input... */
3   decode_info(&bs, &fr_ps);
4   /* ...input... */
5   III_get_side_info(&bs, &III_side_info, &fr_ps);
6
7   for (gr = 0; gr < max_gr; gr++) {
8     for (ch = 0; ch < stereo; ch++) {
9       III_get_scale_factors(gr, ch, ...);
10      III_huffman_decode(gr, ch, ...);
11      III_dequantize_sample(gr, ch, ...);
12    } /* ch */
13
14    III_stereo(gr, ...);
15
16    for (ch = 0; ch < stereo; ch++) {
17      III_reorder(ch, gr, ...);
18      III_antialias(ch, gr, ...);
19
20      for (sb = 0; sb < SBLIMIT; sb++) {
21        III_hybrid(sb, ch, ...);
22      } /* ss */
23
24      for (ss = 0; ss < SSLIMIT; ss++) {
25        for (sb = 0; sb < SBLIMIT; sb++) {
26          if ((ss % 2) && (sb % 2))
27            polyPhaseIn[sb] = -hybridOut[sb][ss];
28          else
29            polyPhaseIn[sb] = hybridOut[sb][ss];
30        } /* sb */
31        clip += SubBandSynthesis(ch, ss, ...);
32      } /* ss */
33    } /* ch */
34  } /* gr */
35
36  out_fifo(*pcm_sample, ...);
37 } /* while */
```

input,
header info **2%**

Huffman
decoding **5%**

dequantize **44%**

stereo **<1%**

reorder **<1%**

antialias **<1%**

hybrid **22%**

subband
synthesis **24%**

output **2%**



Motivating Example

EEMBC mp3player Algorithmic components

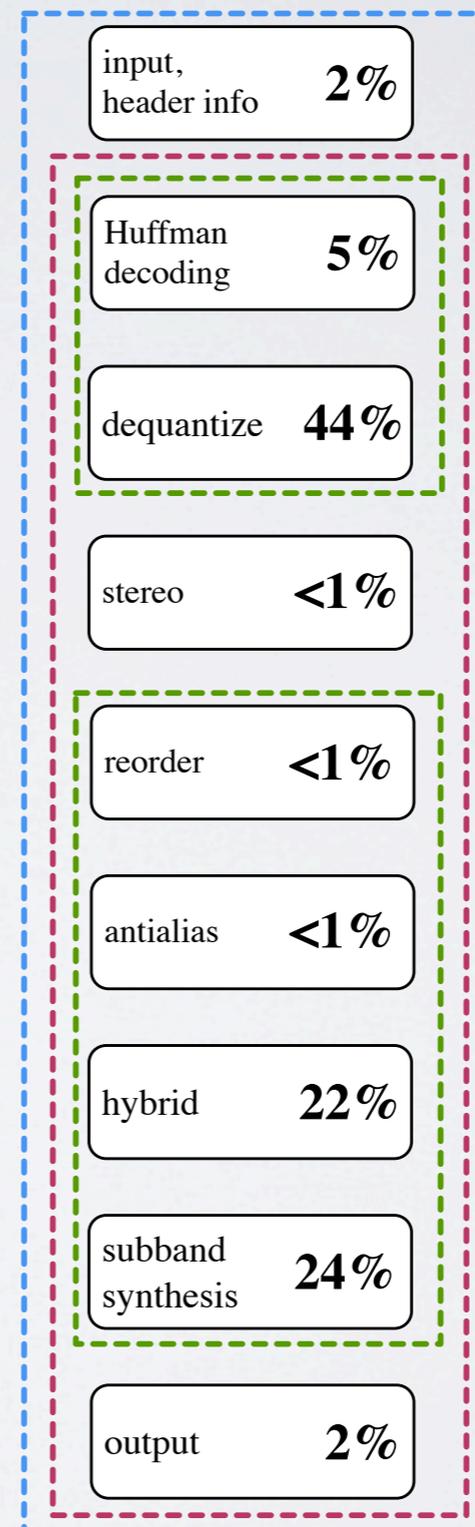
```
1 while(end) {
2   /* ...input... */
3   decode_info(&bs, &fr_ps);
4   /* ...input... */
5   III_get_side_info(&bs, &III_side_info, &fr_ps);
6
7   for (gr = 0; gr < max_gr; gr++) {
8     for (ch = 0; ch < stereo; ch++) {
9       III_get_scale_factors(gr, ch, ...);
10      III_huffman_decode(gr, ch, ...);
11      III_dequantize_sample(gr, ch, ...);
12    } /* ch */
13
14    III_stereo(gr, ...);
15
16    for (ch = 0; ch < stereo; ch++) {
17      III_reorder(ch, gr, ...);
18      III_antialias(ch, gr, ...);
19
20      for (sb = 0; sb < SBLIMIT; sb++) {
21        III_hybrid(sb, ch, ...);
22      } /* ss */
23
24      for (ss = 0; ss < SSLIMIT; ss++) {
25        for (sb = 0; sb < SBLIMIT; sb++) {
26          if ((ss % 2) && (sb % 2))
27            polyPhaseIn[sb] = -hybridOut[sb][ss];
28          else
29            polyPhaseIn[sb] = hybridOut[sb][ss];
30        } /* sb */
31        clip += SubBandSynthesis(ch, ss, ...);
32      } /* ss */
33    } /* ch */
34  } /* gr */
35
36  out_fifo(*pcm_sample, ...);
37 } /* while */
```

Level 1

Level 3

Level 2

Level 3





Motivating Example

EEMBC mp3player

Single-level partitioning

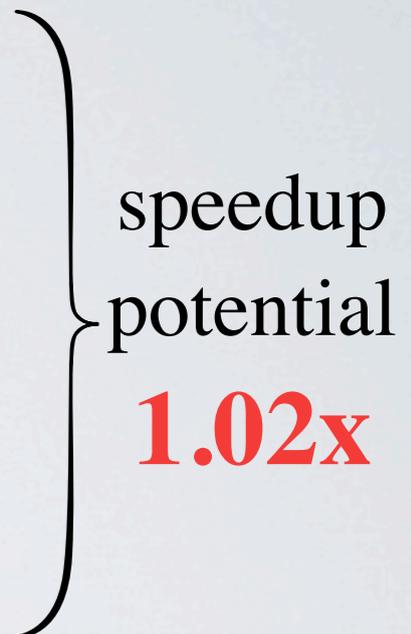
```
1 while(end) {
2   /* ...input... */
3   decode_info(&bs, &fr_ps);
4   /* ...input... */
5   III_get_side_info(&bs, &III_side_info, &fr_ps);
6
7   for (gr = 0; gr < max_gr; gr++) {
8     for (ch = 0; ch < stereo; ch++) {
9       III_get_scale_factors(gr, ch, ...);
10      III_huffman_decode(gr, ch, ...);
11      III_dequantize_sample(gr, ch, ...);
12    } /* ch */
13
14    III_stereo(gr, ...);
15
16    for (ch = 0; ch < stereo; ch++) {
17      III_reorder(ch, gr, ...);
18      III_antialias(ch, gr, ...);
19
20      for (sb = 0; sb < SBLIMIT; sb++) {
21        III_hybrid(sb, ch, ...);
22      } /* sb */
23
24      for (ss = 0; ss < SSLIMIT; ss++) {
25        for (sb = 0; sb < SBLIMIT; sb++) {
26          if ((ss % 2) && (sb % 2))
27            polyPhaseIn[sb] = -hybridOut[sb][ss];
28          else
29            polyPhaseIn[sb] = hybridOut[sb][ss];
30        } /* sb */
31        clip += SubBandSynthesis(ch, ss, ...);
32      } /* ss */
33    } /* ch */
34  } /* gr */
35
36  out_fifo(*pcm_sample, ...);
37 } /* while */
```

Level 1

Level 3

Level 2

Level 3



Single-level pipeline is inefficient!



Motivating Example

EEMBC mp3player

Multi-level partitioning

```

1 while(end) {
2   /* ...input... */
3   decode_info(&bs, &fr_ps);
4   /* ...input... */
5   III_get_side_info(&bs, &III_side_info, &fr_ps);
6
7   for (gr = 0; gr < max_gr; gr++) {
8     for (ch = 0; ch < stereo; ch++) {
9       III_get_scale_factors(gr, ch, ...);
10      III_huffman_decode(gr, ch, ...);
11      III_dequantize_sample(gr, ch, ...);
12    } /* ch */
13
14    III_stereo(gr, ...);
15
16    for (ch = 0; ch < stereo; ch++) {
17      III_reorder(ch, gr, ...);
18      III_antialias(ch, gr, ...);
19
20      for (sb = 0; sb < SBLIMIT; sb++) {
21        III_hybrid(sb, ch, ...);
22      } /* ss */
23
24      for (ss = 0; ss < SSLIMIT; ss++) {
25        for (sb = 0; sb < SBLIMIT; sb++) {
26          if ((ss % 2) && (sb % 2))
27            polyPhaseIn[sb] = -hybridOut[sb][ss];
28          else
29            polyPhaseIn[sb] = hybridOut[sb][ss];
30        } /* sb */
31        clip += SubBandSynthesis(ch, ss, ...);
32      } /* ss */
33    } /* ch */
34  } /* gr */
35
36  out_fifo(*pcm_sample, ...);
37 } /* while */

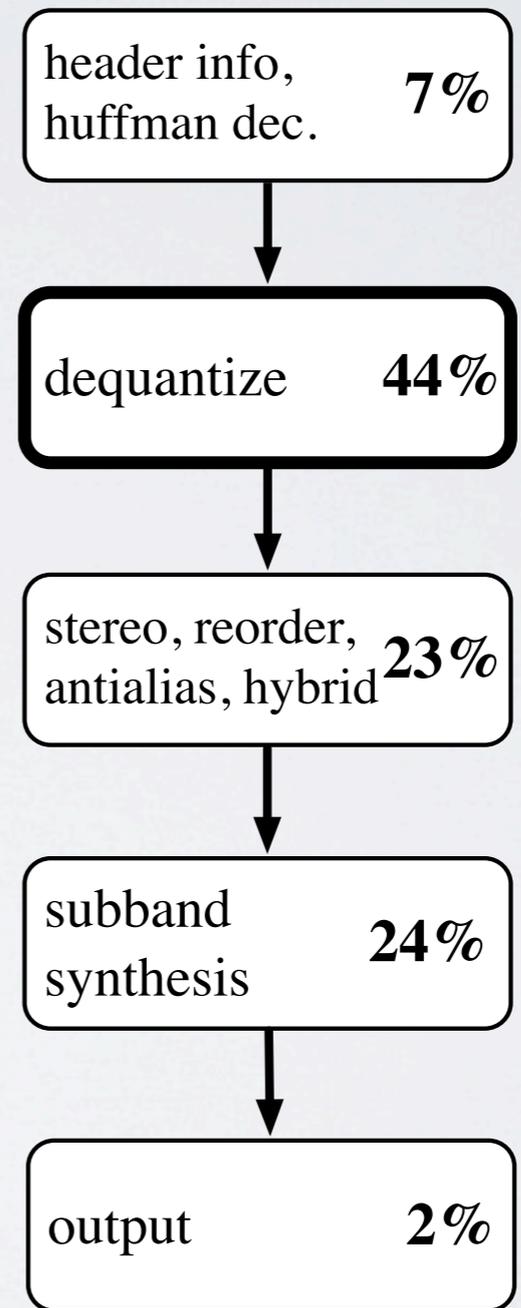
```

Level 1

Level 3

Level 2

Level 3



speedup potential
2.27x



Motivating Example

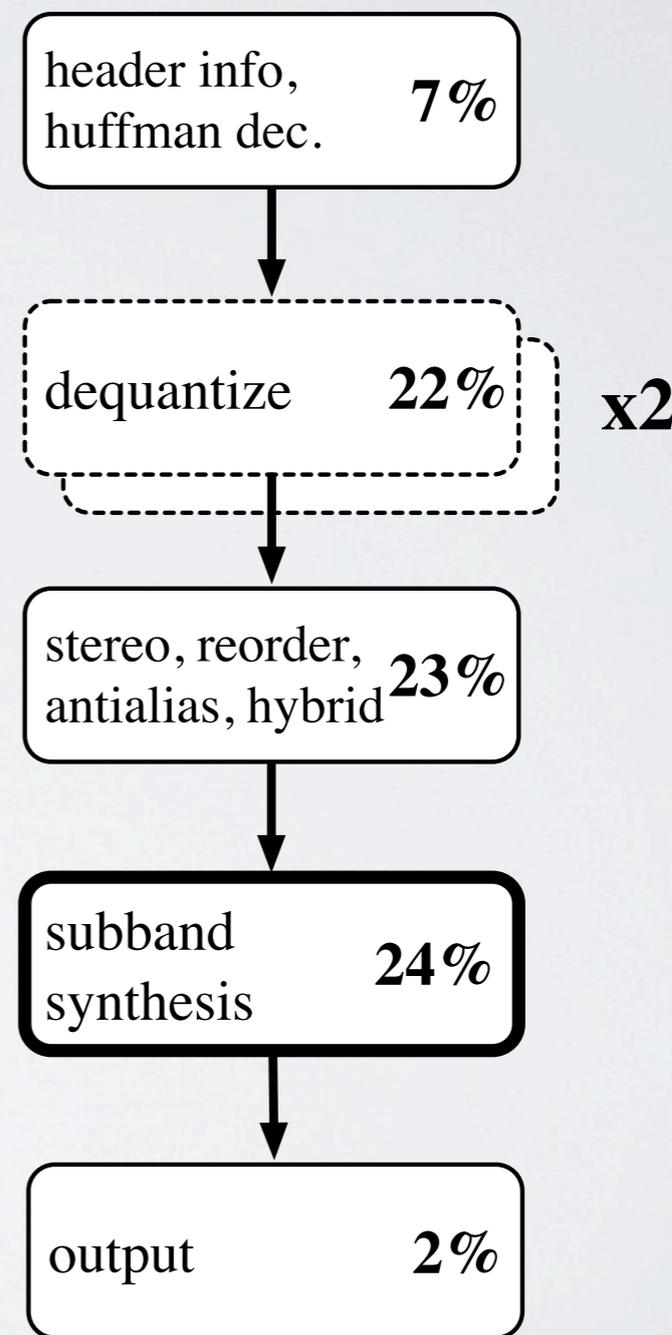
EEMBC mp3player

Multi-level + Replication Partitioning

```

1 while(end) {
2   /* ...input... */
3   decode_info(&bs, &fr_ps);
4   /* ...input... */
5   III_get_side_info(&bs, &III_side_info, &fr_ps);
6
7   for (gr = 0; gr < max_gr; gr++) {
8     for (ch = 0; ch < stereo; ch++) {
9       III_get_scale_factors(gr, ch, ...);
10      III_huffman_decode(gr, ch, ...);
11      III_dequantize_sample(gr, ch, ...);
12    } /* ch */
13
14    III_stereo(gr, ...);
15
16    for (ch = 0; ch < stereo; ch++) {
17      III_reorder(ch, gr, ...);
18      III_antialias(ch, gr, ...);
19
20      for (sb = 0; sb < SBLIMIT; sb++) {
21        III_hybrid(sb, ch, ...);
22      } /* ss */
23
24      for (ss = 0; ss < SSLIMIT; ss++) {
25        for (sb = 0; sb < SBLIMIT; sb++) {
26          if ((ss % 2) && (sb % 2))
27            polyPhaseIn[sb] = -hybridOut[sb][ss];
28          else
29            polyPhaseIn[sb] = hybridOut[sb][ss];
30        } /* sb */
31        clip += SubBandSynthesis(ch, ss, ...);
32      } /* ss */
33    } /* ch */
34  } /* gr */
35
36  out_fifo(*pcm_sample, ...);
37 } /* while */

```



speedup potential **4.16x**



Motivating Example

EEMBC mp3player

Multi-level + Replication
Partitioning

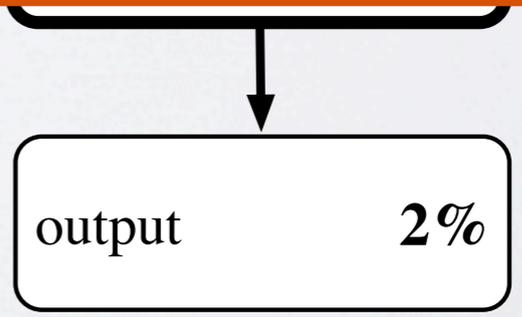
```
1 while(end) {
2   /* ...input... */
3   decode_info(&bs, &fr-ps);
4   /* ...input... */
5   III
6
7   for
8   for
9
10
11 }
12
13 II
14
15 for
16 for
17
18
19
20
21
22
23
24
25
26
27
28   else
29     polyPhaseIn[sb] = hybridOut[sb][ss];
30   } /* sb */
31   clip += SubBandSynthesis(ch, ss, ...);
32   } /* ss */
33   } /* ch */
34 } /* gr */
35
36 out_fifo(*pcm_sample, ...);
37 } /* while */
```

Level 1

header info

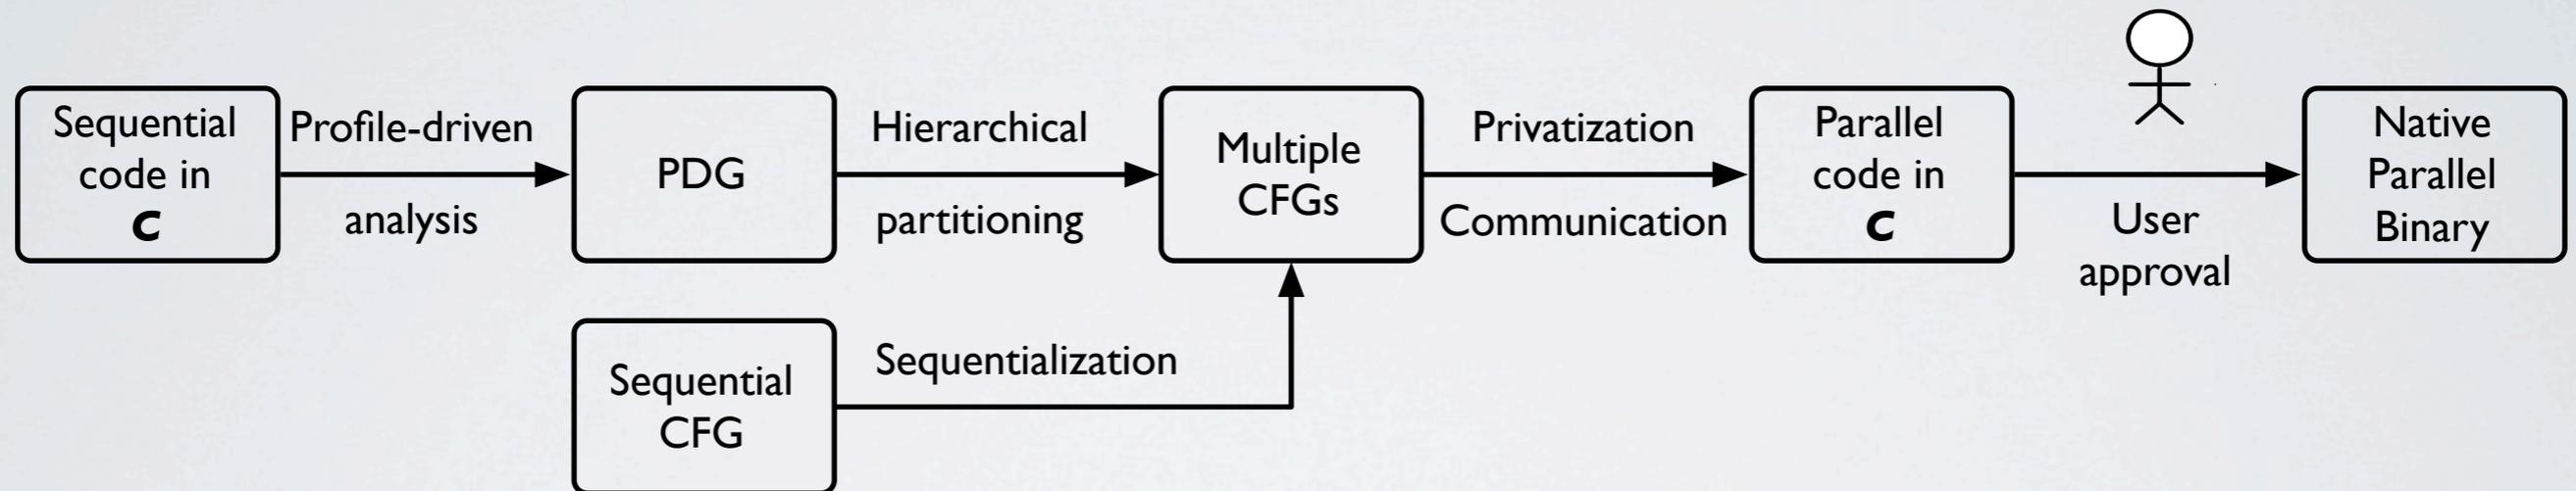
- **Hierarchical** pipelines increase efficiency
- **Replication** of pipeline stages exposes additional parallelism
- **Orthogonal** to traditional parallelisation approaches (parallel loops inside pipeline stages)

speedup potential
4.16x





Approach





Partitioning Strategy

- Pipeline performance determined by the **slowest stage**
- Apply **code transformations** only to the slowest stage to uncover further parallelism



Partitioning Algorithm

- Top-down approach: loops and functions *folded*
- Preprocess PDG of the loop:
 - Form Strongly Connected Components
 - Focus on slowest component:
 - If data-parallel (i) greedily augment it, and (ii) replicate until another component is the slowest
 - If not data-parallel try to reduce the execution time by *unfolding* loop/function nodes in the component
- Partition pipeline using the load of the slowest component as threshold

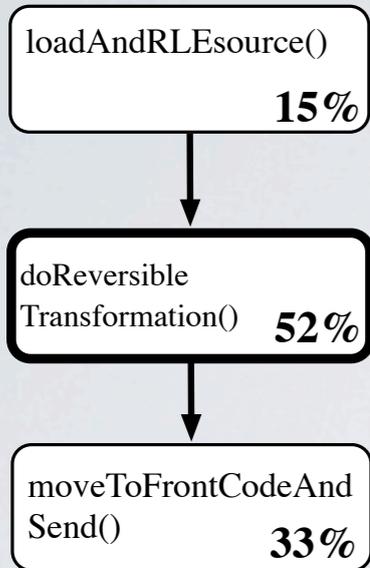


Partitioning Operations

- Loop/Function unfolding
 - “Opening up” loop/function for hierarchical partitioning
- Replication
 - Duplication of partitioning unit for parallel execution
- Split function
 - Insert pipeline stage boundary within function body
- Augment block
 - Merge separate blocks into single pipeline stage

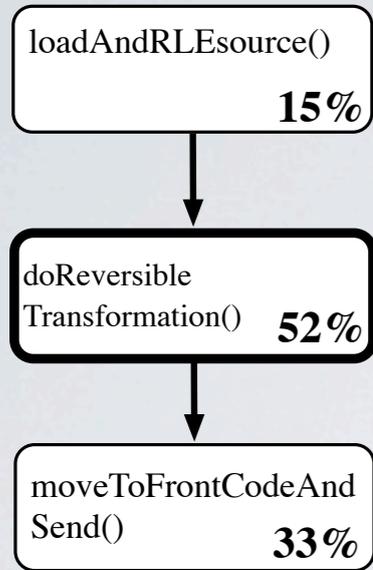


Example

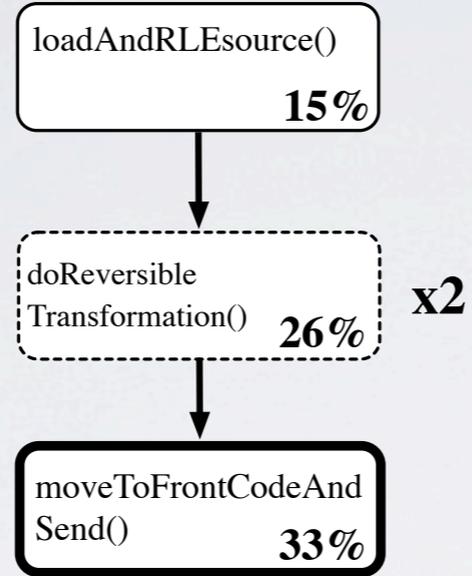




Example

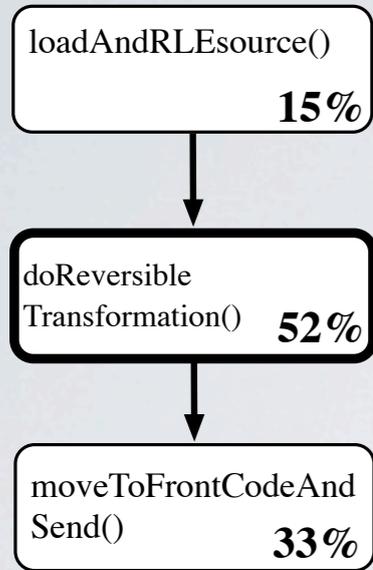


replicate
x2

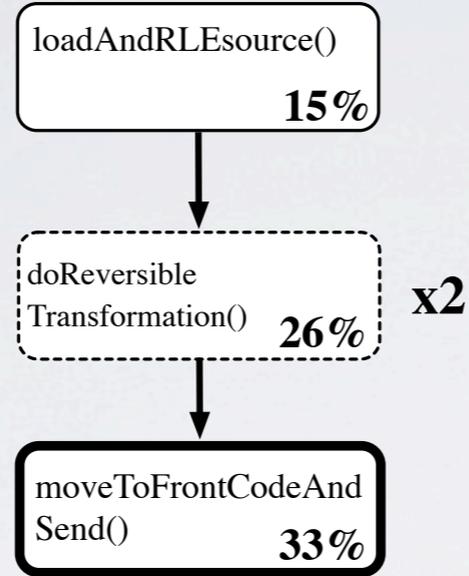




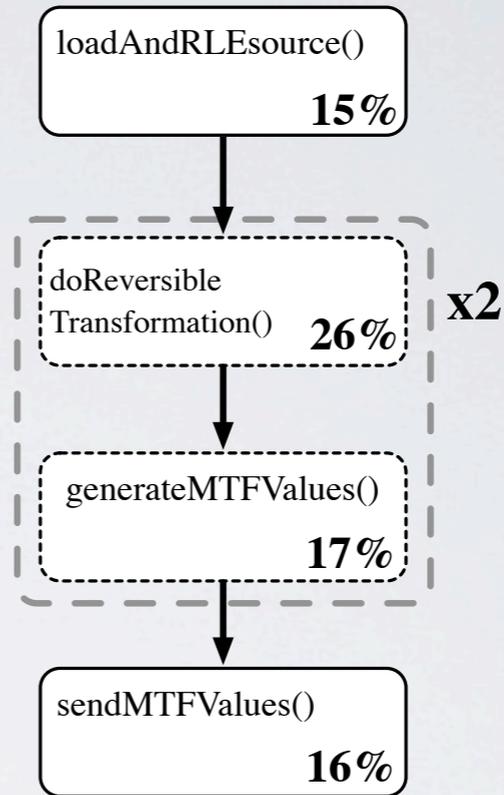
Example



→
replicate
x2

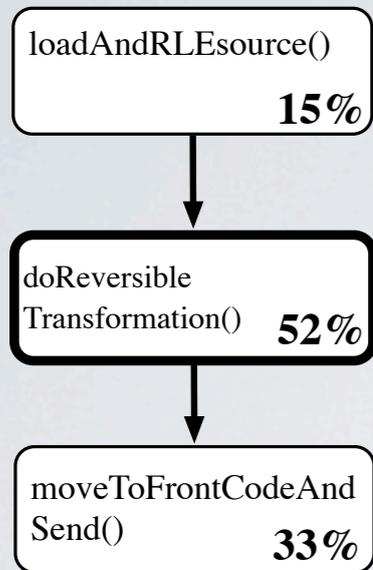


→
split
function

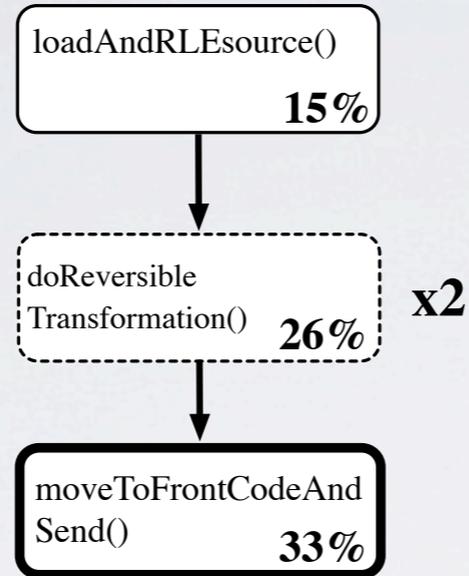




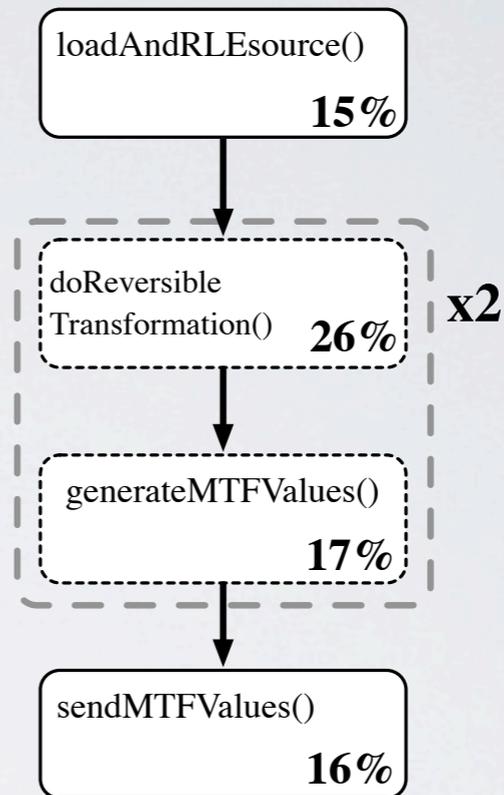
Example



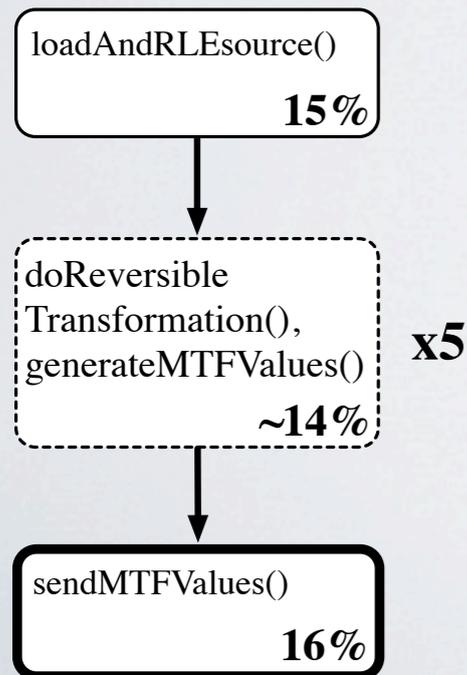
replicate x2



split function

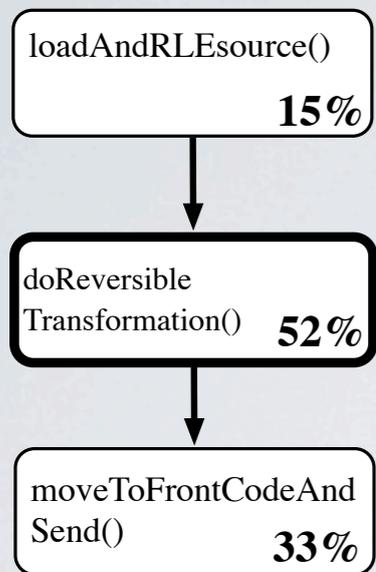


augment & replicate x5

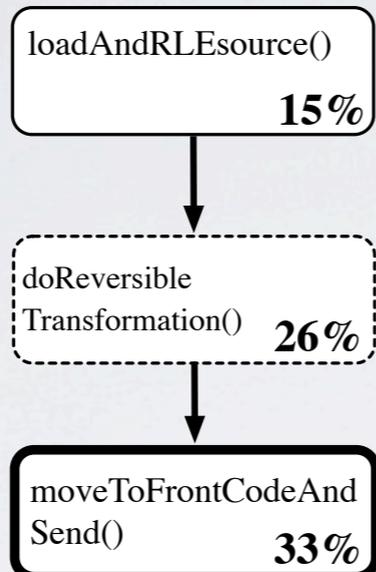




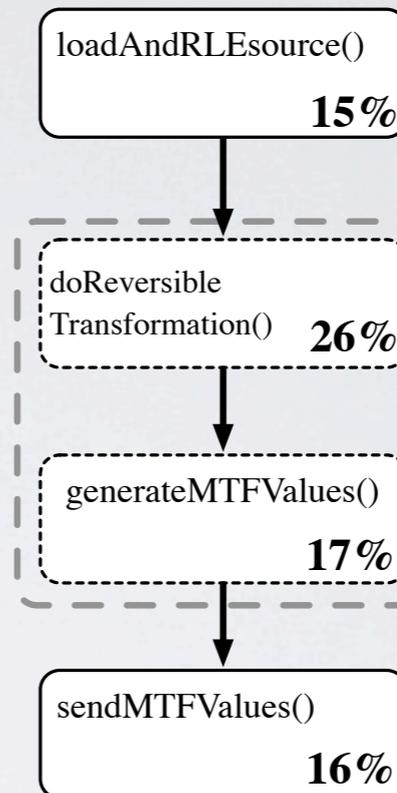
Example



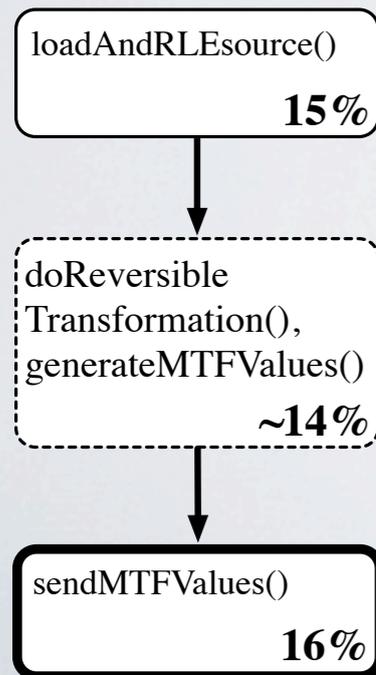
*replicate
x2*



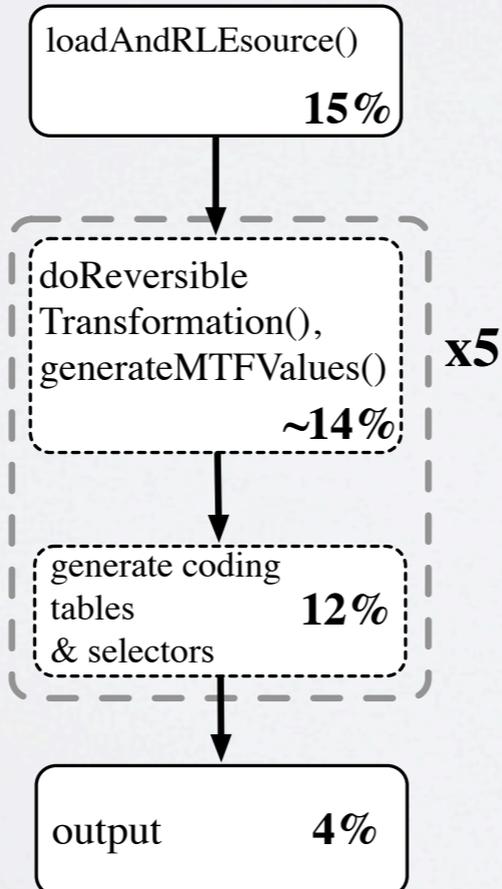
*split
function*



*augment
&
replicate
x5*

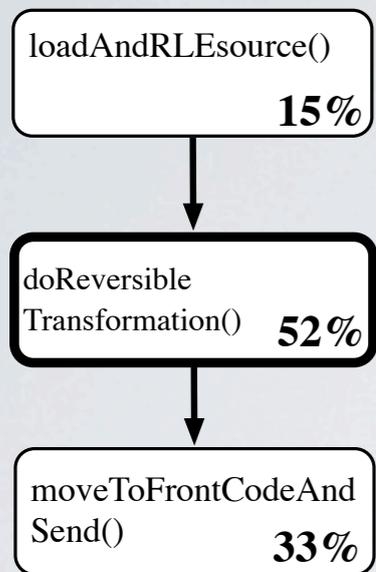


*split
function*

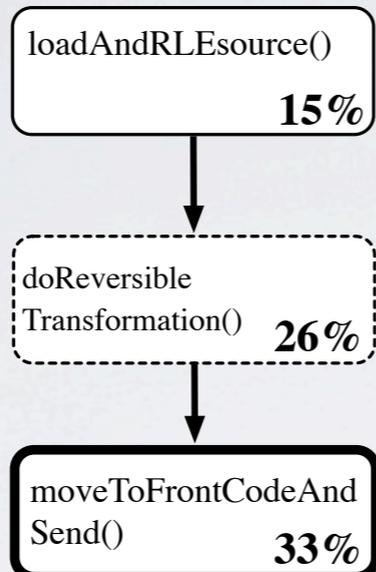




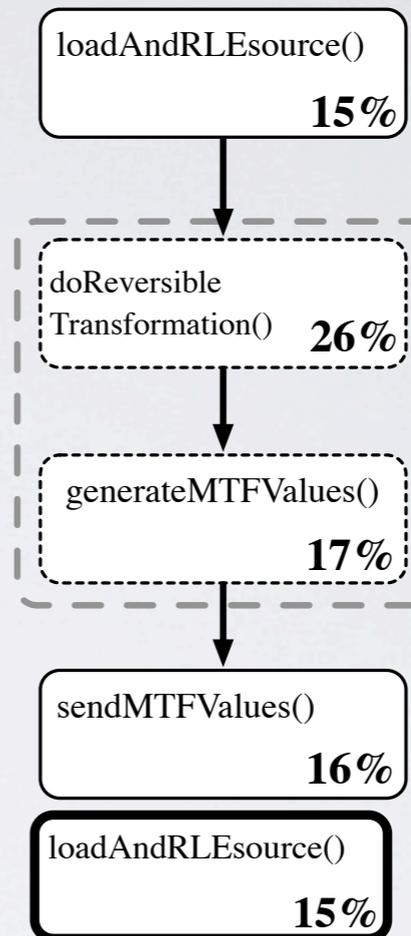
Example



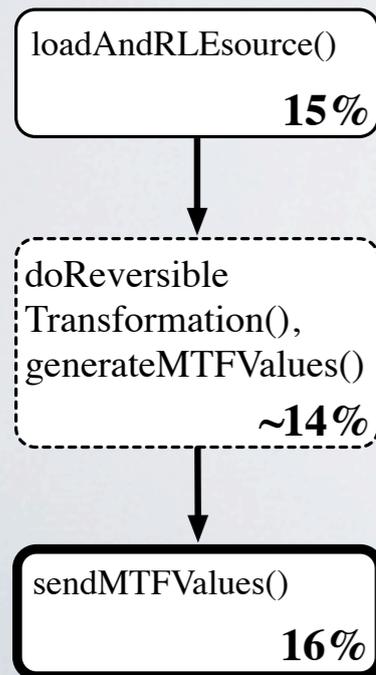
replicate x2



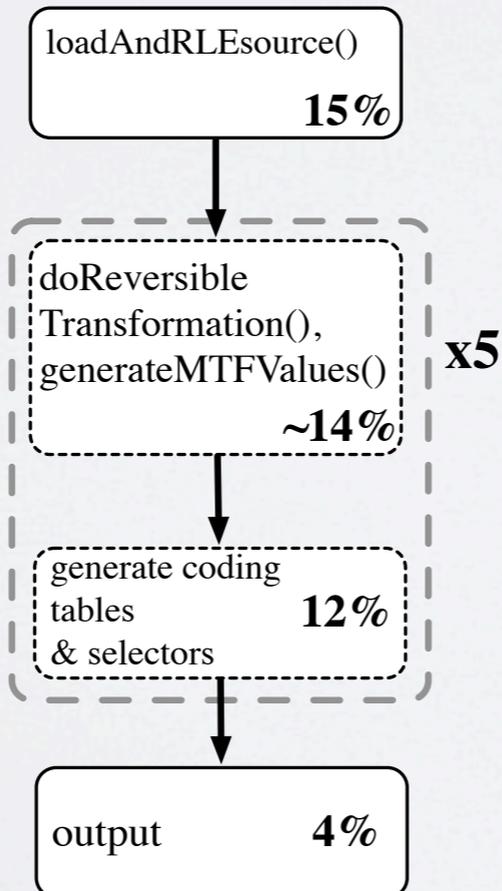
split function



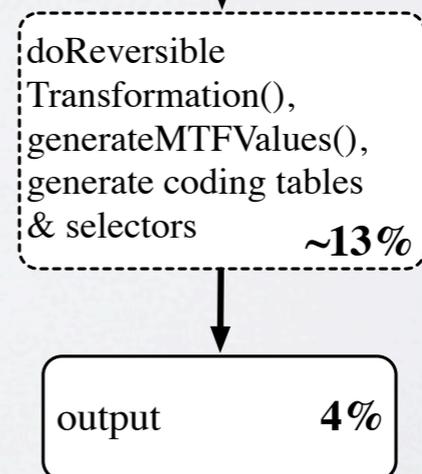
augment & replicate x5



split function



augment & replicate x6





Experimental Evaluation

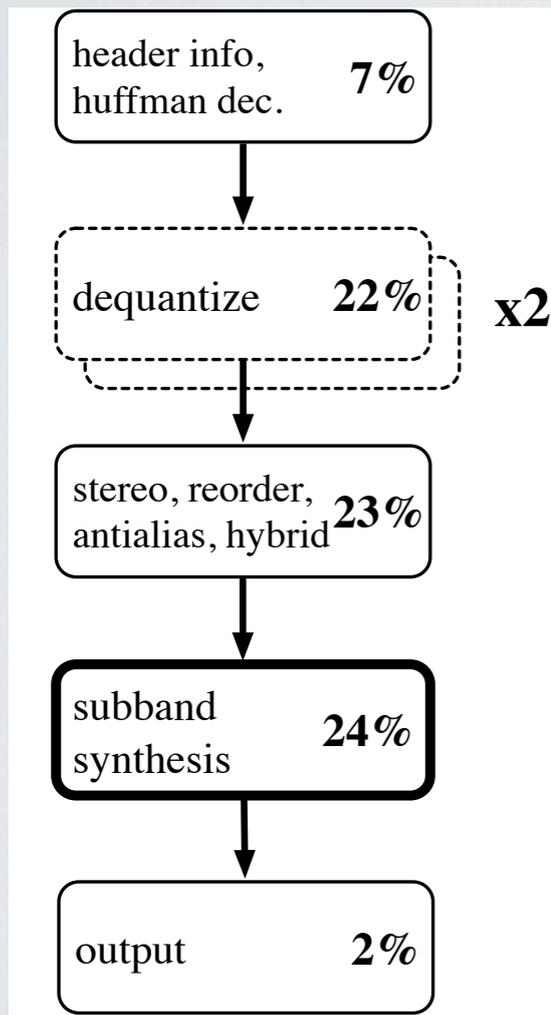
Application	source	lines
MP3 decode	EEMBC 2.0	20K
MPEG-2 decode	EEMBC 2.0	23K
JPEG encode	EEMBC 2.0	22K
bzip2 compress	SPEC CPU2000	5K

Evaluation platform	
Hardware	Dual Socket, Intel Xeon X5450 @ 3.00GHz 2 Quad-cores, 8 cores in total SSE2, SSE3 and SSE4.1 extensions 6Mb L2-cache shared/2 cores (12Mb/chip) 16Gb DDR2 SDRAM
O.S.	64-bit Scientific Linux kernel 2.6.9-55 x86_64
Compiler	GNU GCC 4.4.1 -O3 -march=core2

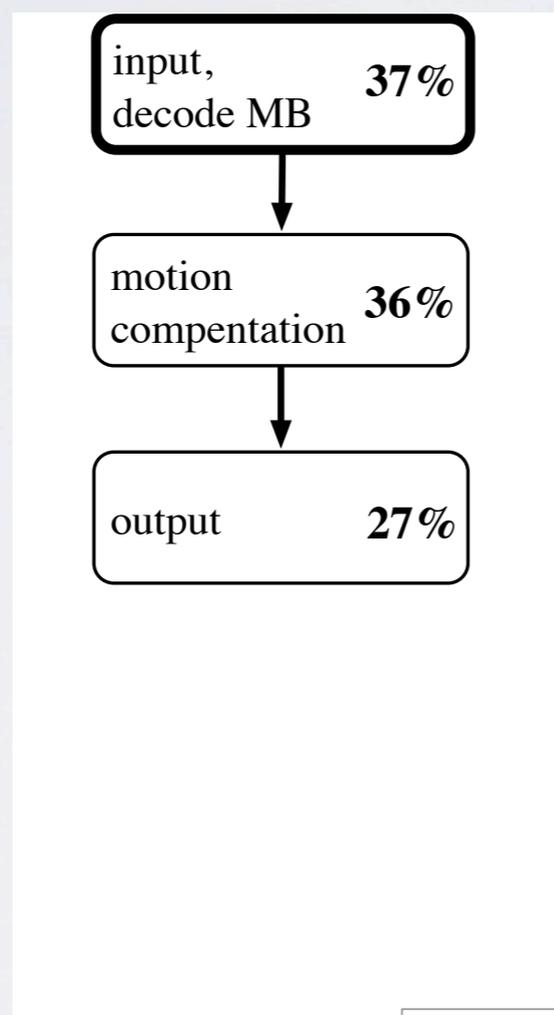


Extracted Pipelines

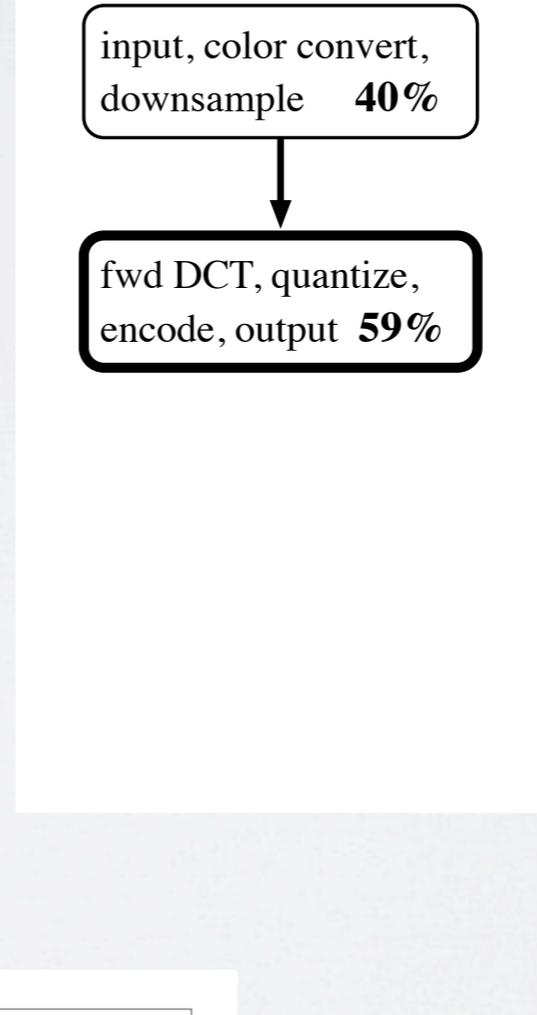
MP3 decoding



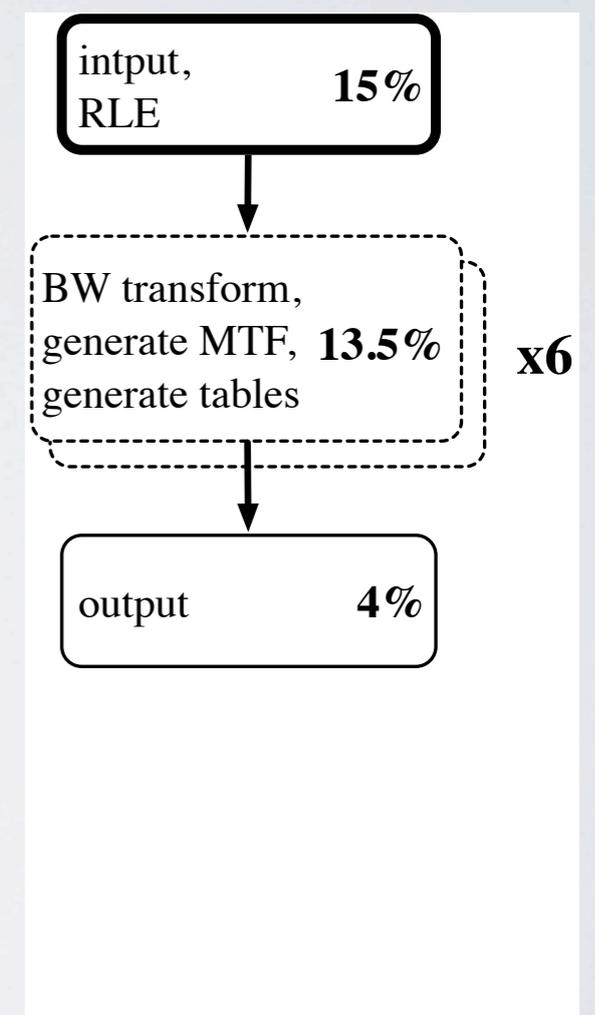
MPEG-2 decoding



JPEG encoding



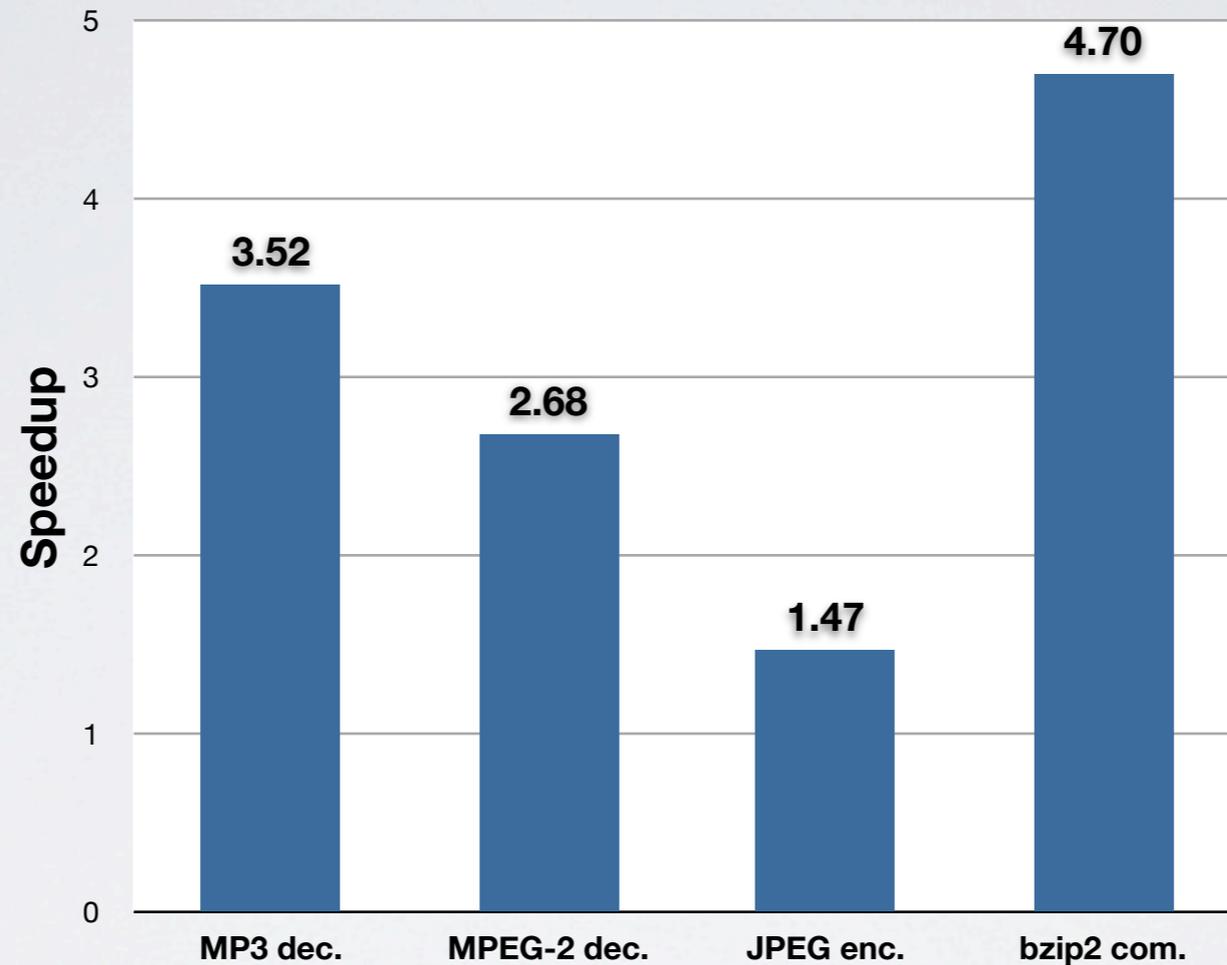
bzip2 compression



▭ : bottleneck stage
▭ (dashed) : replicable stage



Performance Results



Application	replication	multi-loop	func. split	# cores	speedup
MP3 dec.	✓	✓	–	7	3.52x
MPEG-2 dec.	–	✓	✓	3	2.68x
JPEG enc.	–	✓	✓	2	1.47x
bzip2 com.	✓	–	✓	8	4.70x



Further Details

- Sequentialisation of the PDG
- Data privatisation
- Inter-thread communication
- Dynamic memory disambiguation
- Pipeline runtime system



Thanks

- Georgios Tournavitis (now with Intel Research)
- Zheng Wang
- Tobias Edler von Koch
- Igor Böhm
- Damon Fenacci
- Alastair Murray
- Daniel Powell
- Stephen Kyle
- Harry Wagstaff
- Miles Gould
- and my colleagues Mike O'Boyle & Nigel Topham



Summary

- Serious demand for parallelisation tool support
- Static analysis are too conservative
 - Profile driven analyses detect more parallelism, but require additional manual checking
- Mapping of SW parallelism to HW parallelism is non-intuitive and depends on target platform
 - Successful application of machine learning
- More scope for parallelisation beyond FOR loops
 - Start exploiting parallel design patterns



Other Interests

- Everything Parallel
- Code Generation for Embedded Processors
- Fast Instruction Set Simulation
- Parallel JIT Compilation
- Statistical Performance Modelling
- Detection of Parallel Design Patterns
- Mapping for Heterogeneous Multi-Cores



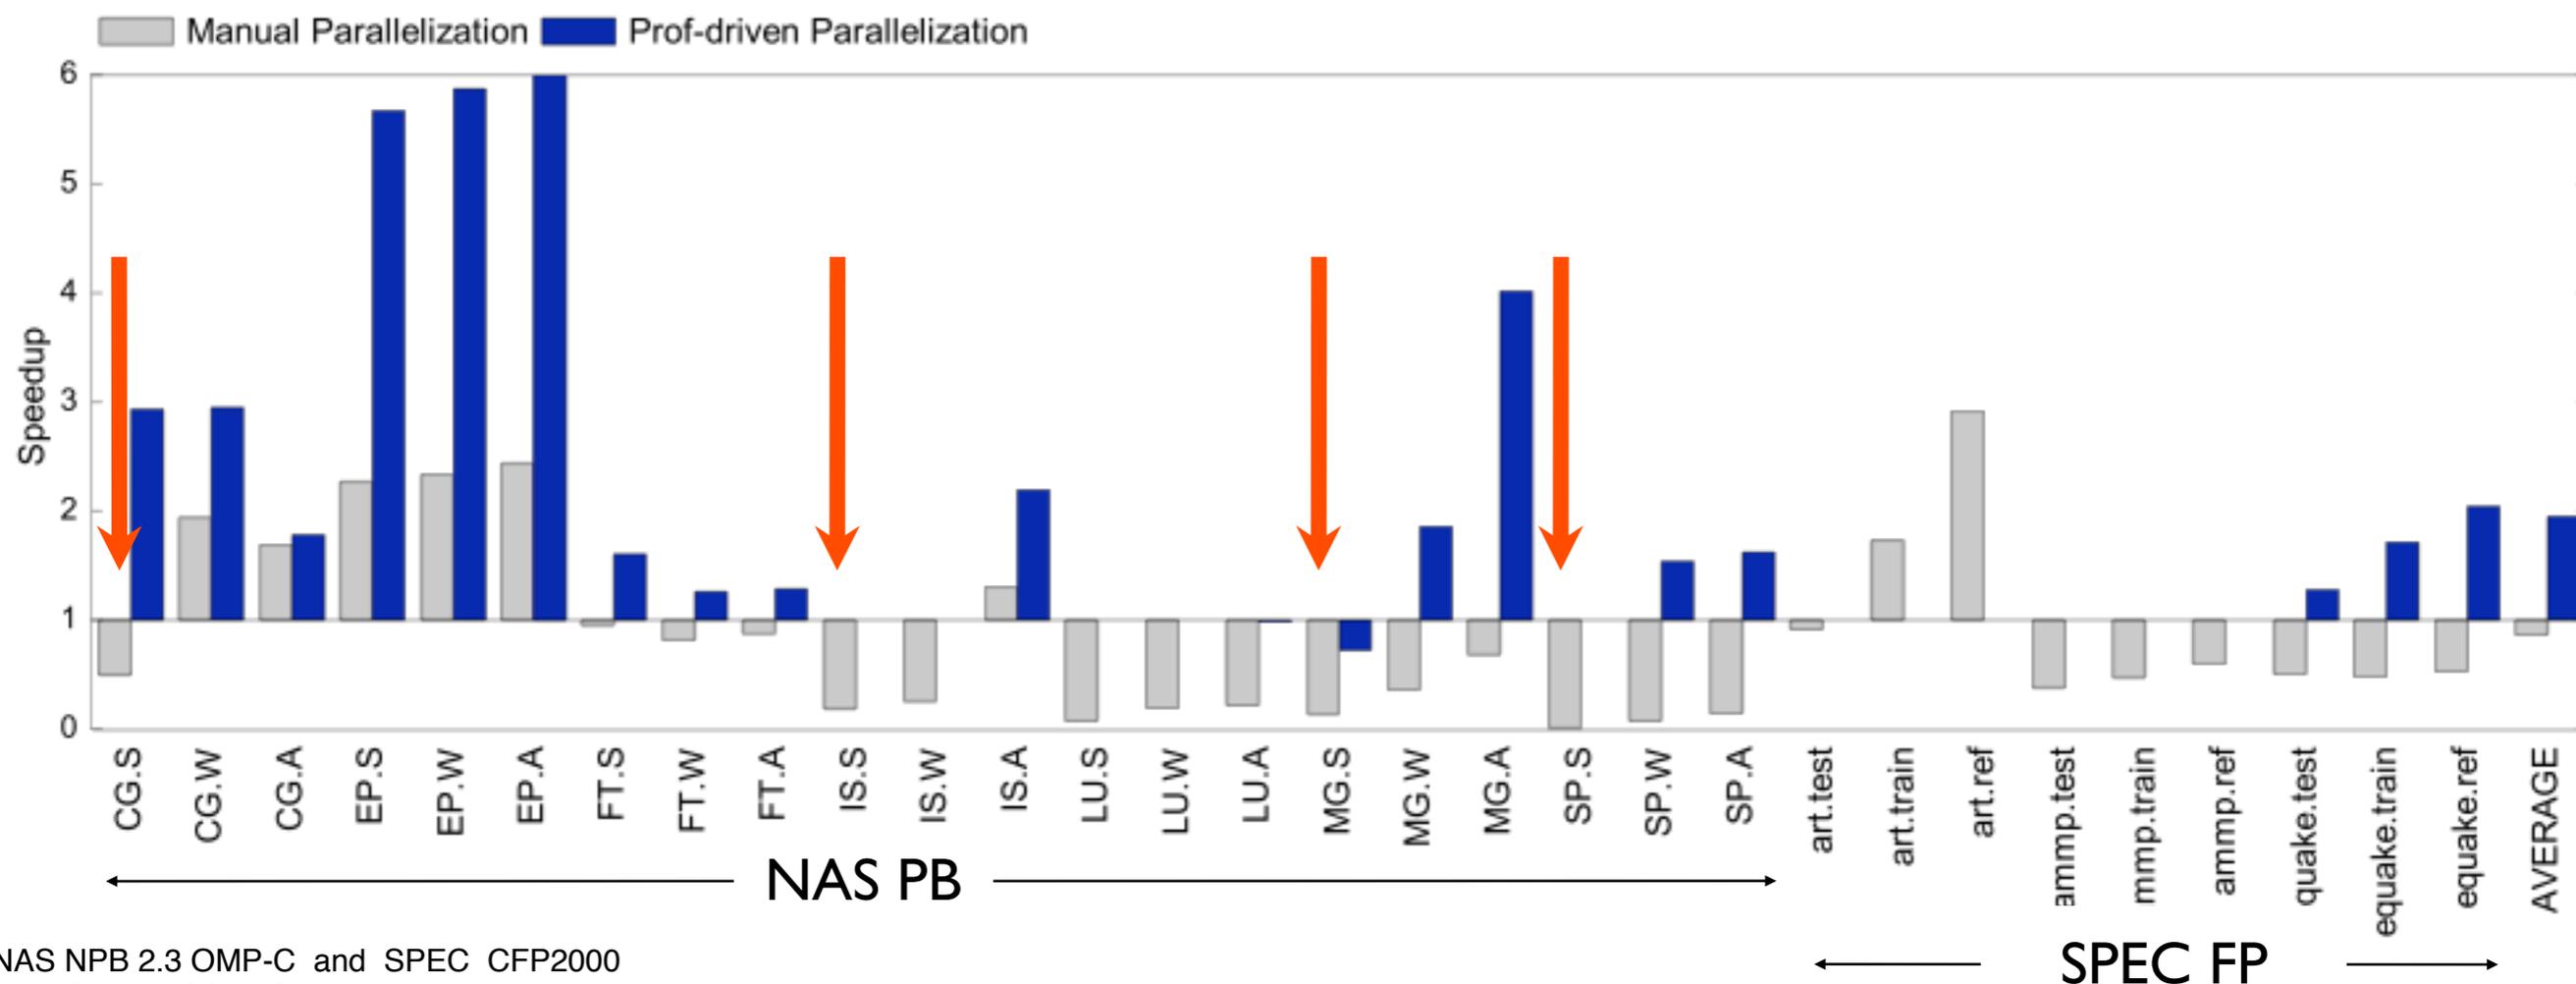
Questions?

BACKUP SLIDES

RESULTS FOR CELL

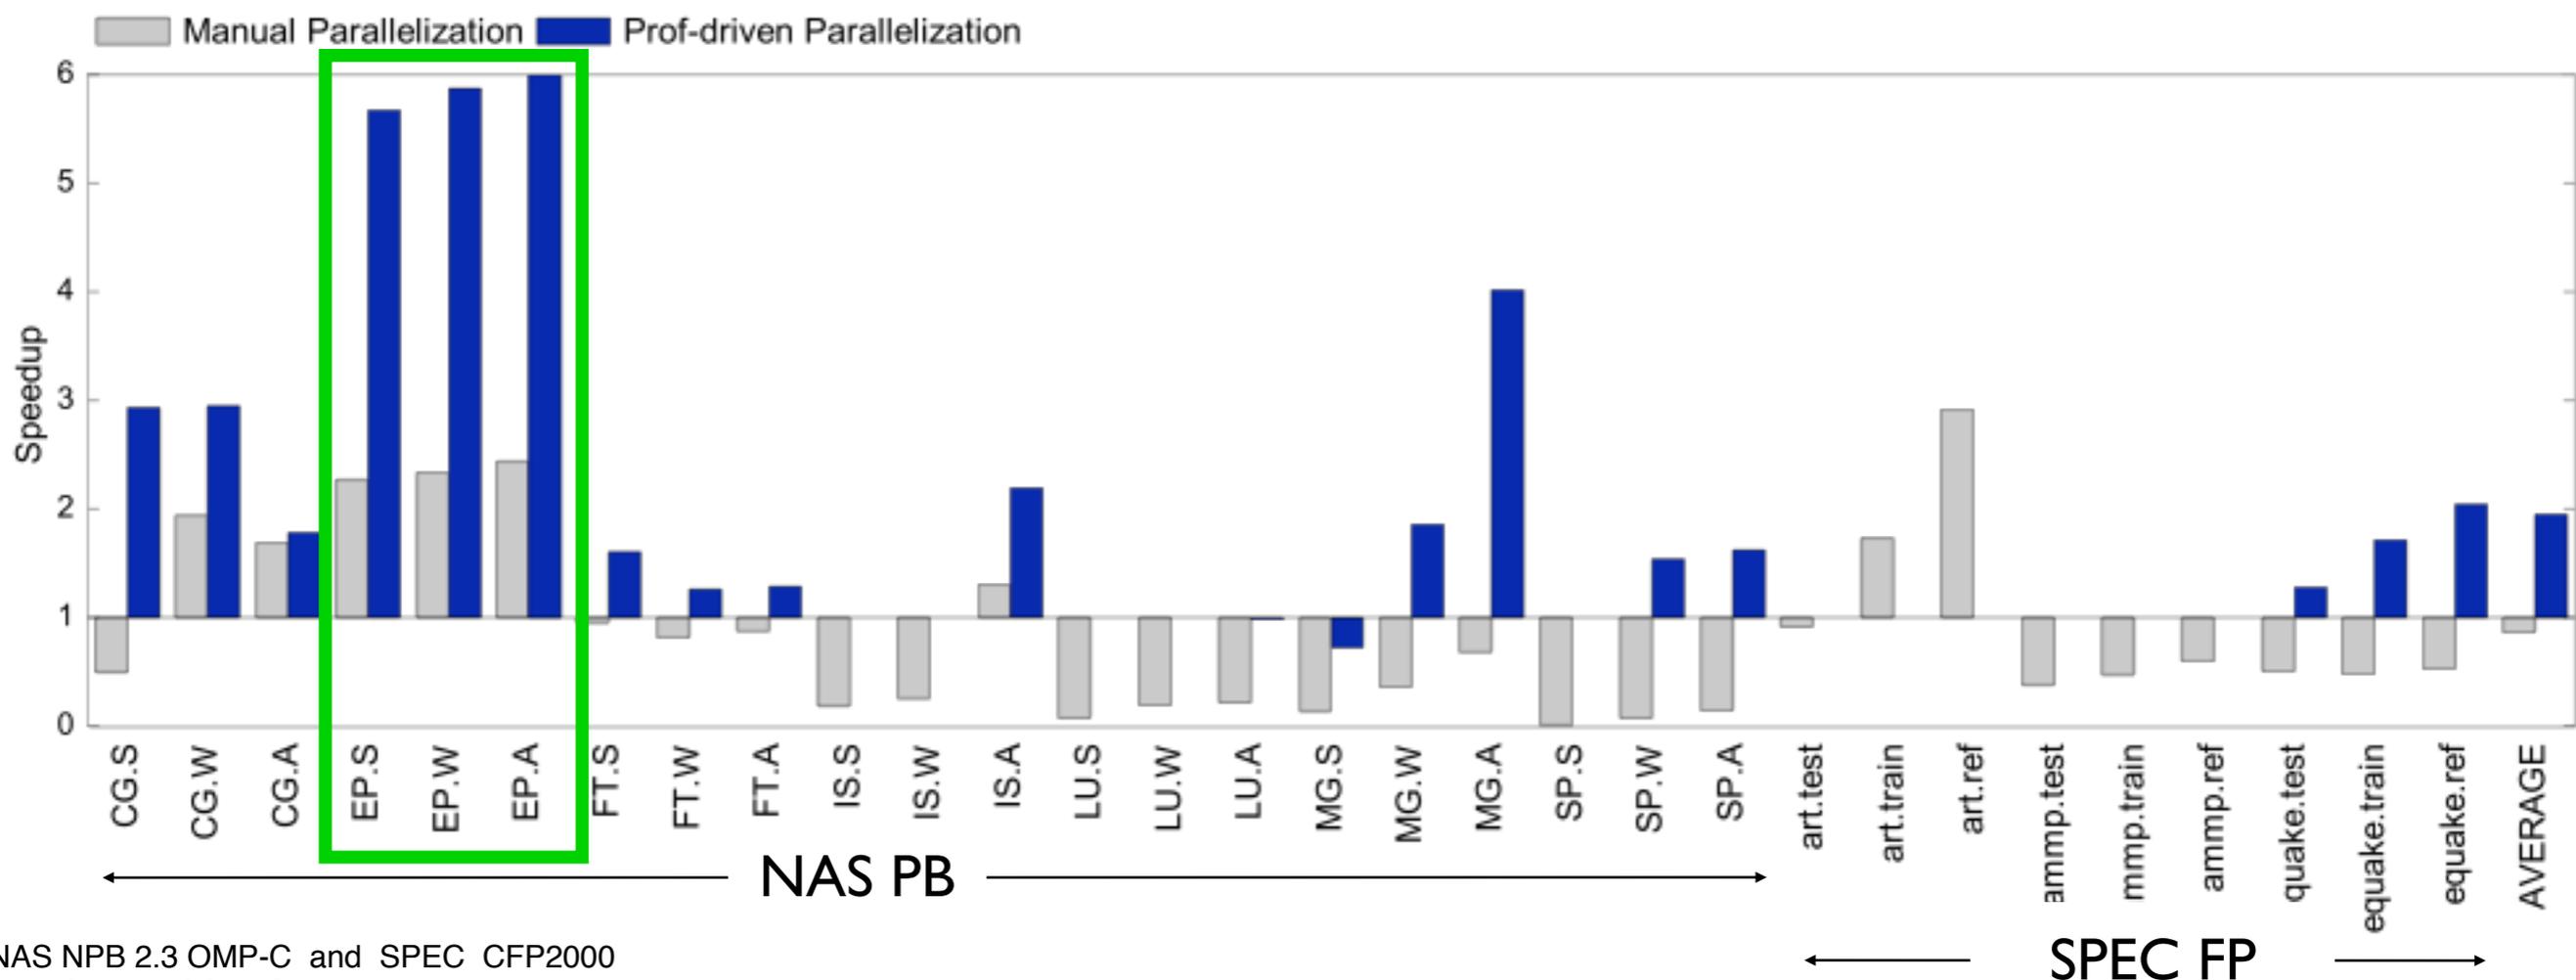
Performance on Cell

- Overhead is more obvious on small datasets



Performance on Cell

- EP gets significant speedup
 - No synchronization
 - Not memory bound



PIPELINES: CODE GENERATION

Parallel code generation



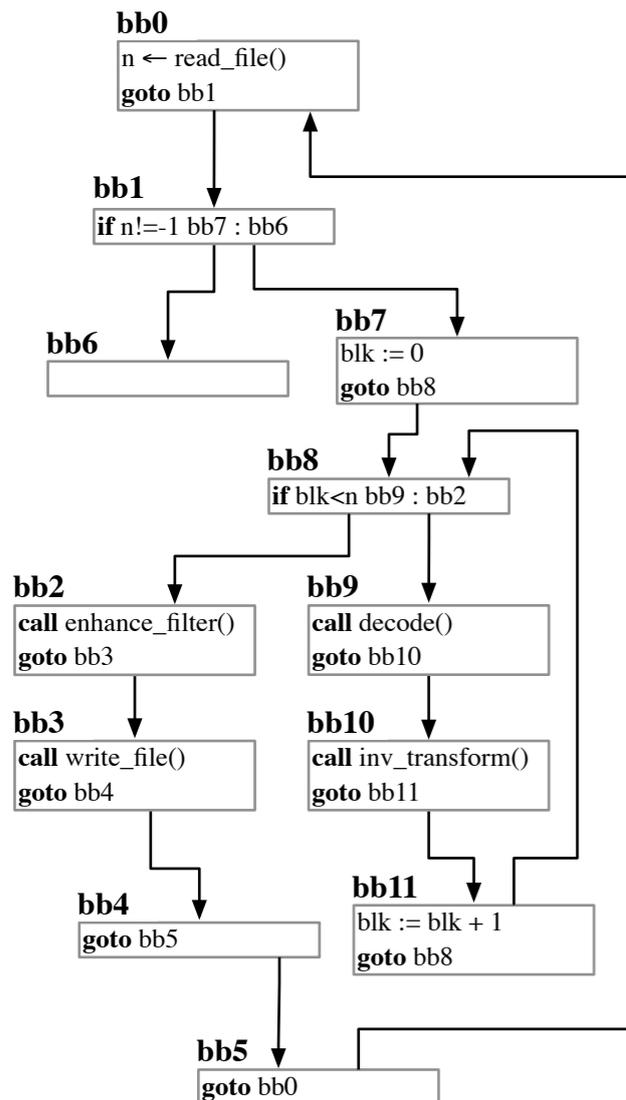
```
5% | 1 while((n = read_file(inf, data)) != EOF) {  
    | 2   for (blk=0; blk<n; blk++) {  
20% | 3     coef[blk] = decode(data, blk);  
50% | 4     raw_data[blk] = inv_transform(coef, blk);  
    | 5   }  
20% | 6   out_data = enhance_filter(raw_data);  
5%  | 7   write_file(outf, out_data);  
    | 8 } /* while */
```

Parallel code generation



5%	1	while ((n = read_file(inf, data)) != EOF) {
	2	for (blk=0; blk<n; blk++) {
20%	3	coef[blk] = decode(data, blk);
50%	4	raw_data[blk] = inv_transform(coef, blk);
	5	}
20%	6	out_data = enhance_filter(raw_data);
5%	7	write_file(outf, out_data);
	8	} /* while */

Sequential CFG



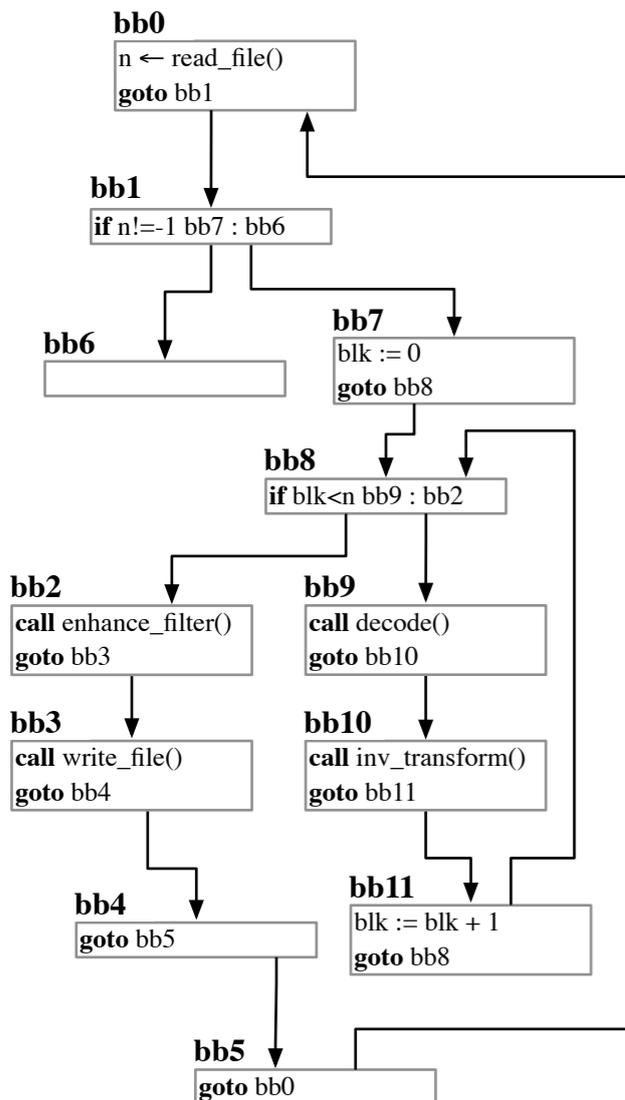
Parallel code generation



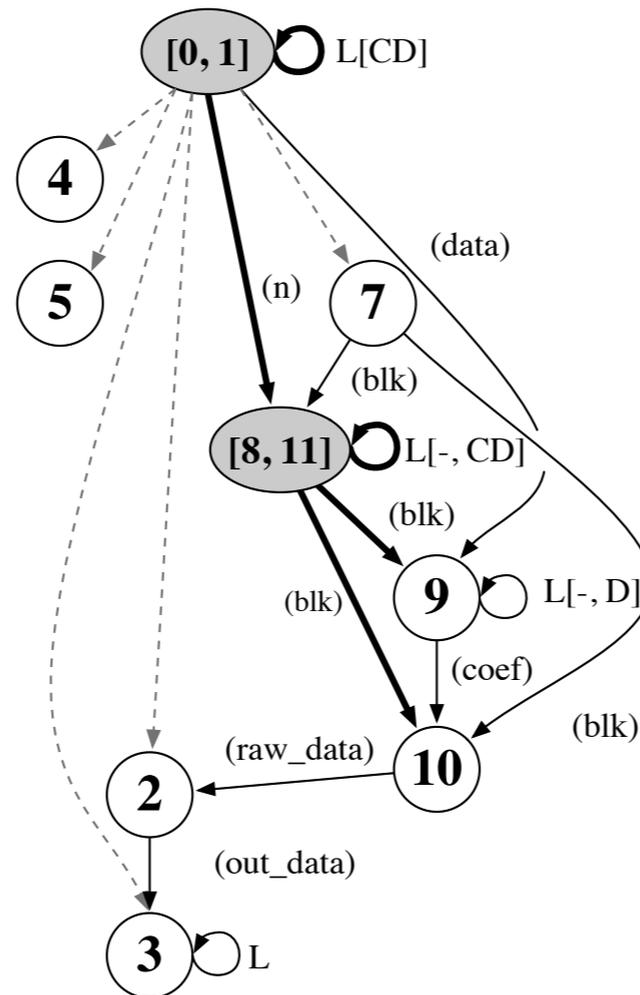
```

5% | 1 while((n = read_file(inf, data)) != EOF) {
    | 2   for (blk=0; blk<n; blk++) {
20% | 3     coef[blk] = decode(data, blk);
50% | 4     raw_data[blk] = inv_transform(coef, blk);
    | 5   }
20% | 6   out_data = enhance_filter(raw_data);
5%  | 7   write_file(outf, out_data);
    | 8 } /* while */
  
```

Sequential CFG



PDG



- control dep.
- data dep.
- control & data dep.

Parallel code generation

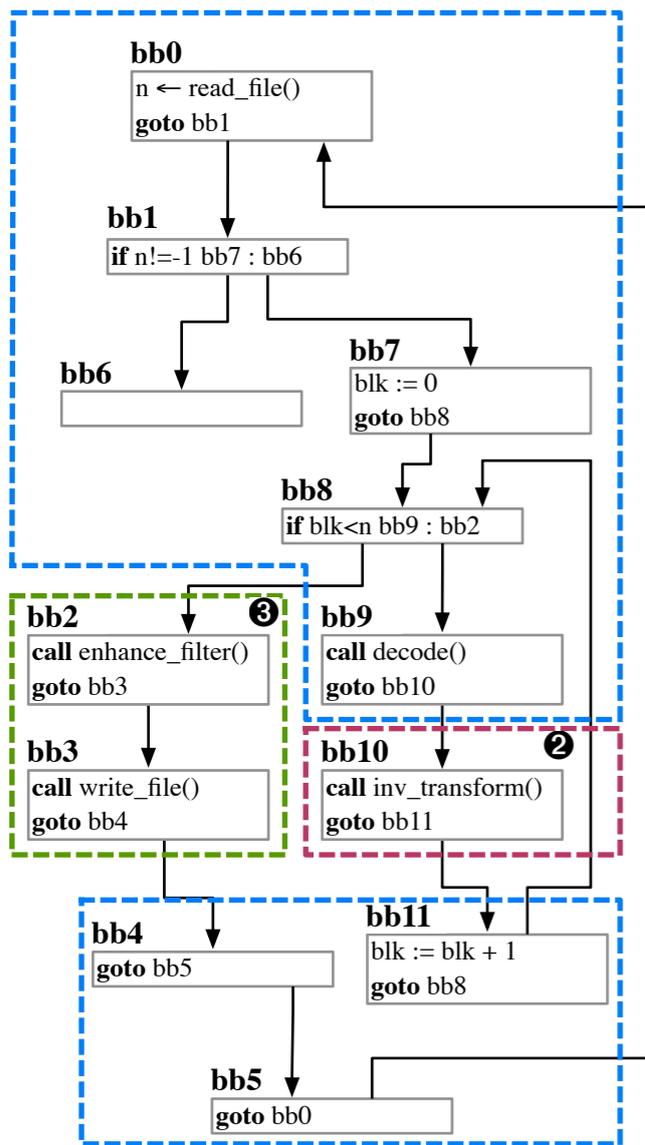
```

5% 1 while((n = read_file(inf, data)) != EOF) {
20% 2   for (blk=0; blk<n; blk++) {
50% 3     coef[blk] = decode(data, blk);
4   raw_data[blk] = inv_transform(coef, blk);
5   }
20% 6   out_data = enhance_filter(raw_data);
5% 7   write_file(outf, out_data);
8 } /* while */

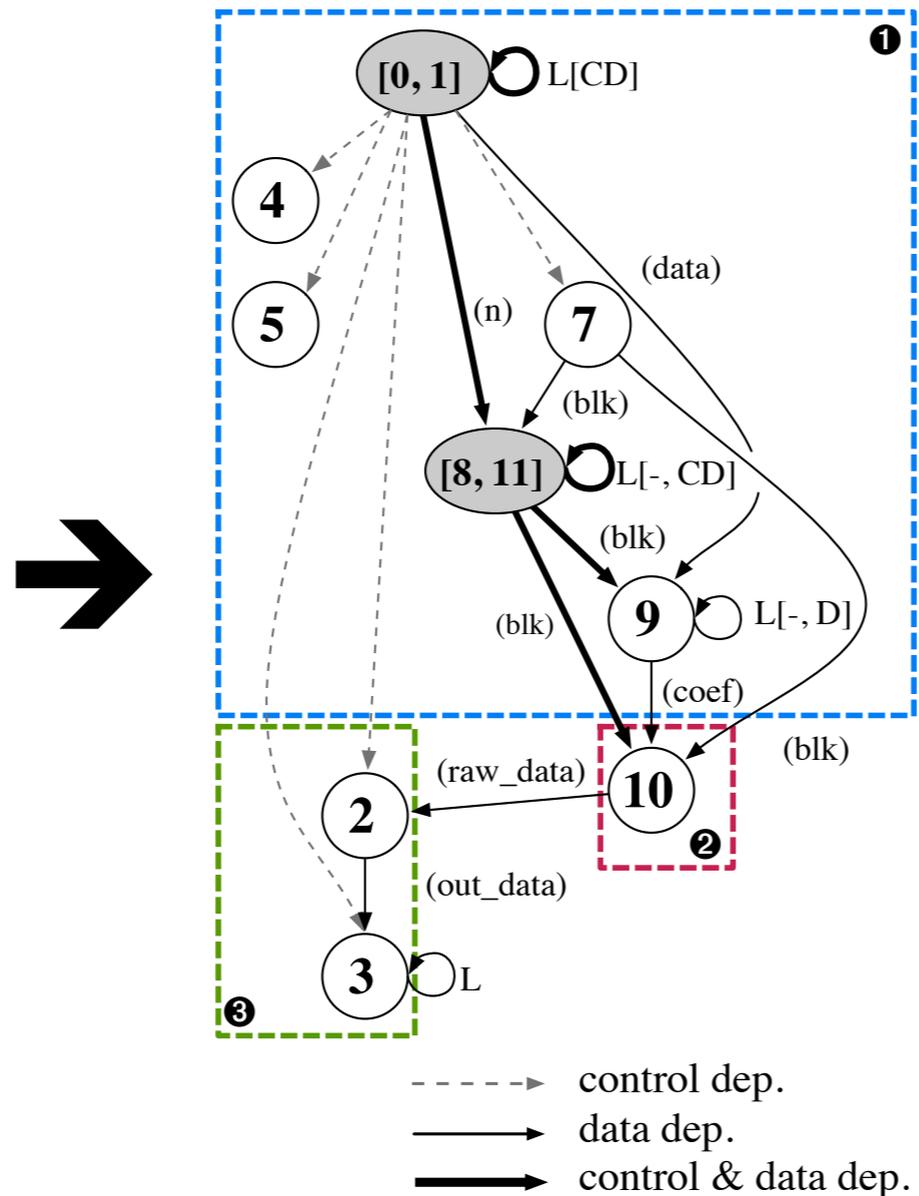
```

Stage 1 (blue dashed box) covers lines 1-2.
Stage 2 (red dashed box) covers lines 3-4.
Stage 3 (green dashed box) covers lines 6-7.

Sequential CFG



PDG Partitioning



Parallel code generation

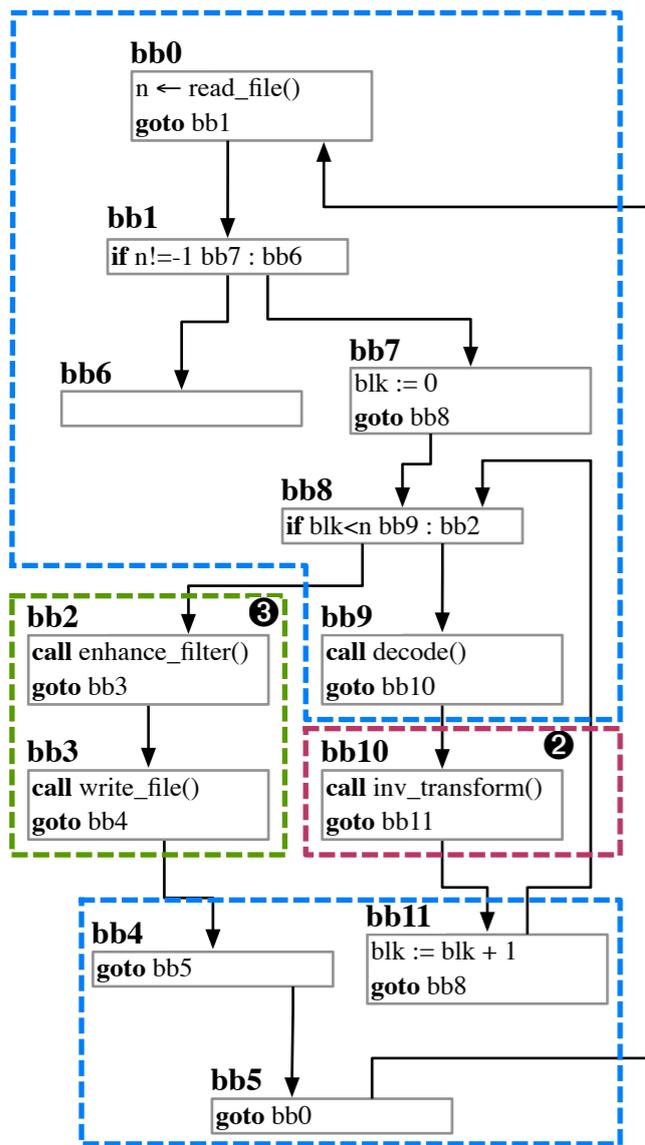
```

1  while((n = read_file(inf, data)) != EOF) {
2    for (blk=0; blk<n; blk++) {
3      coef[blk] = decode(data, blk);
4      raw_data[blk] = inv_transform(coef, blk);
5    }
6    out_data = enhance_filter(raw_data);
7    write_file(outf, out_data);
8  } /* while */

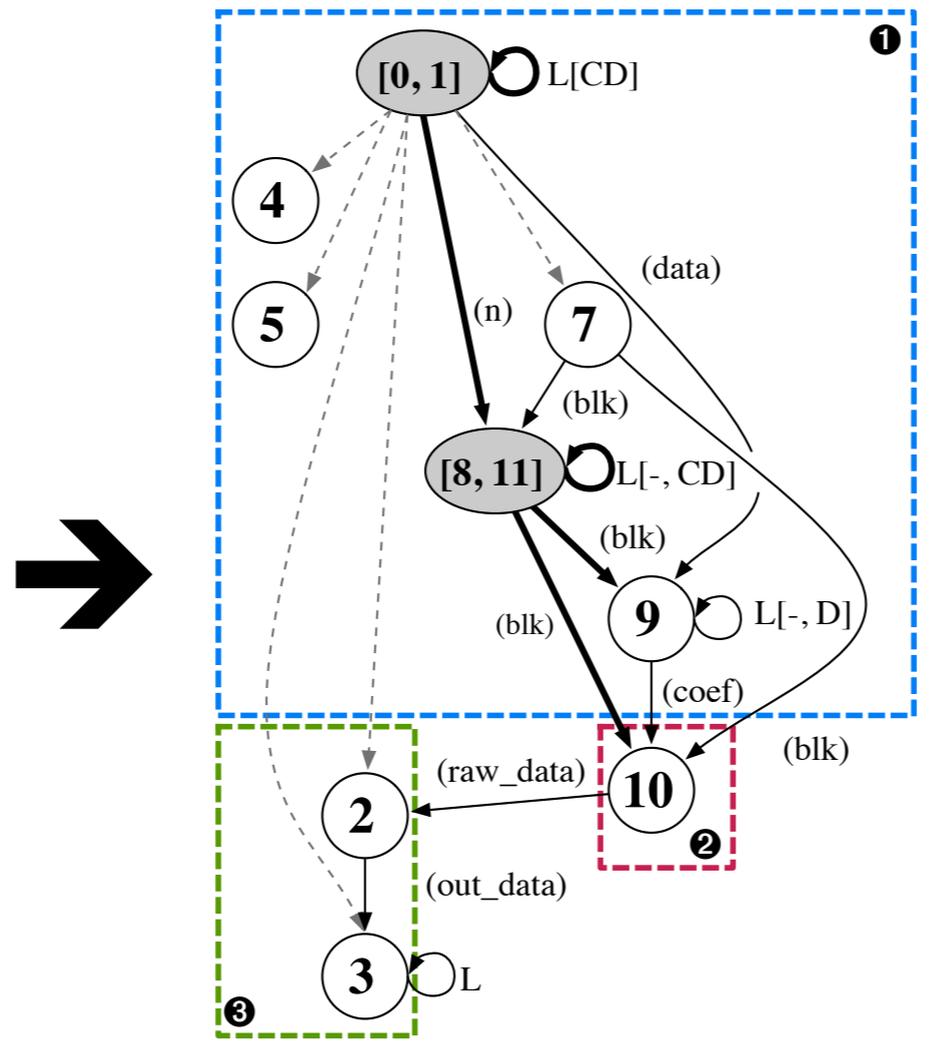
```

5% Stage 1
 20% Stage 2
 50% Stage 3
 20%
 5%

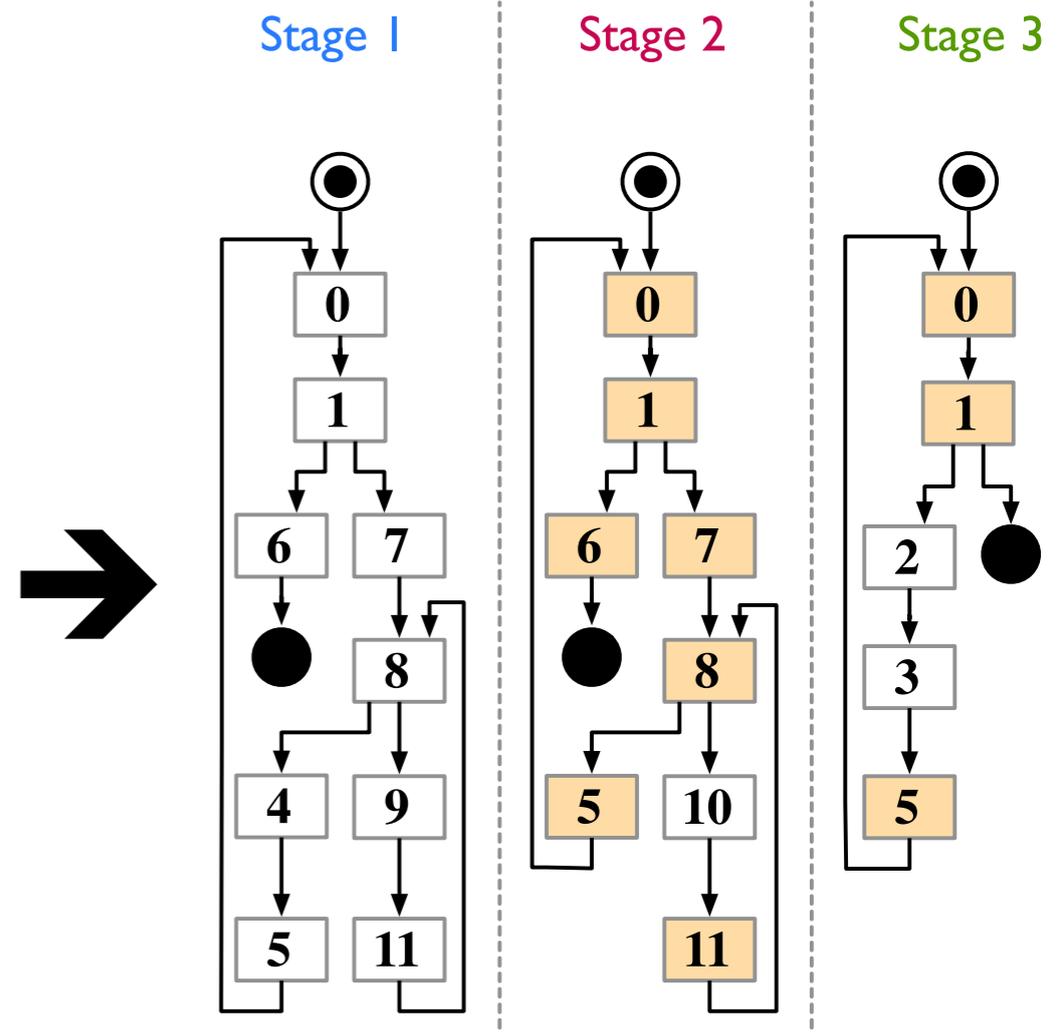
Sequential CFG



PDG Partitioning



Sequentialization



- - - - - control dep.
 ——— data dep.
 ——— control & data dep.

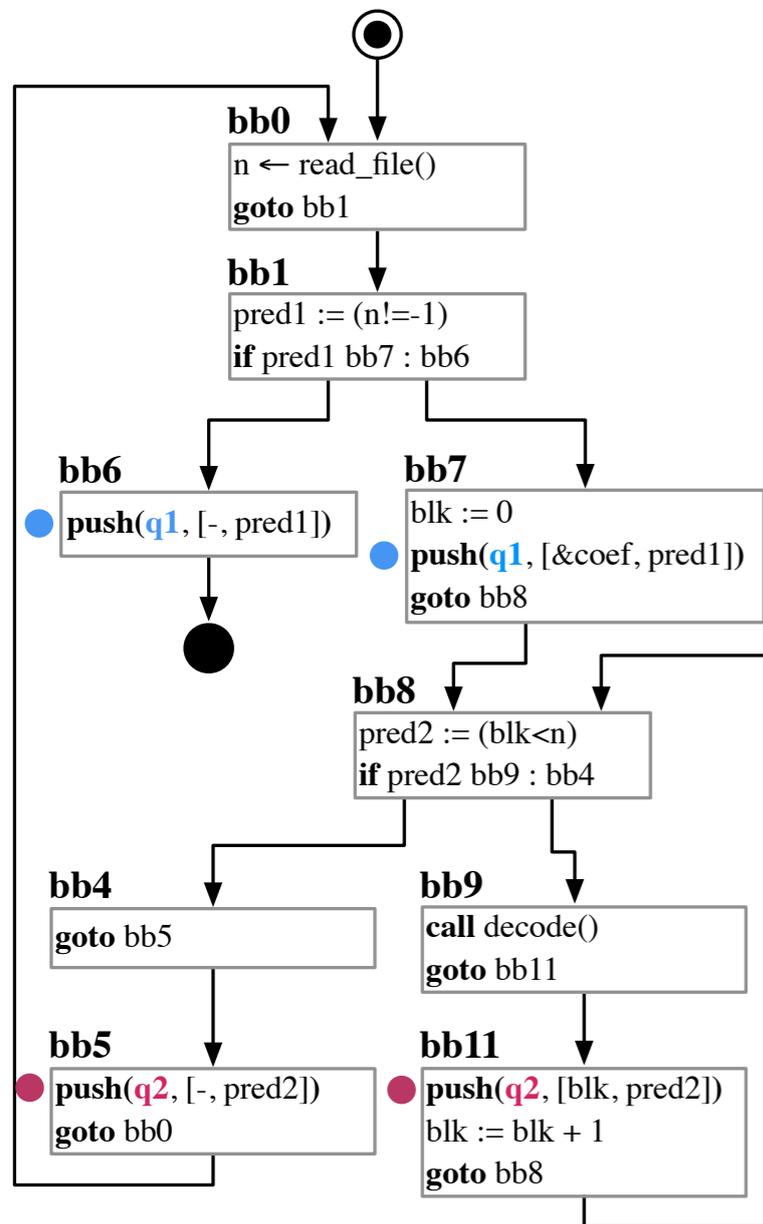
■ replicated BB
 □ normal BB

PIPELINES: COMMUNICATION

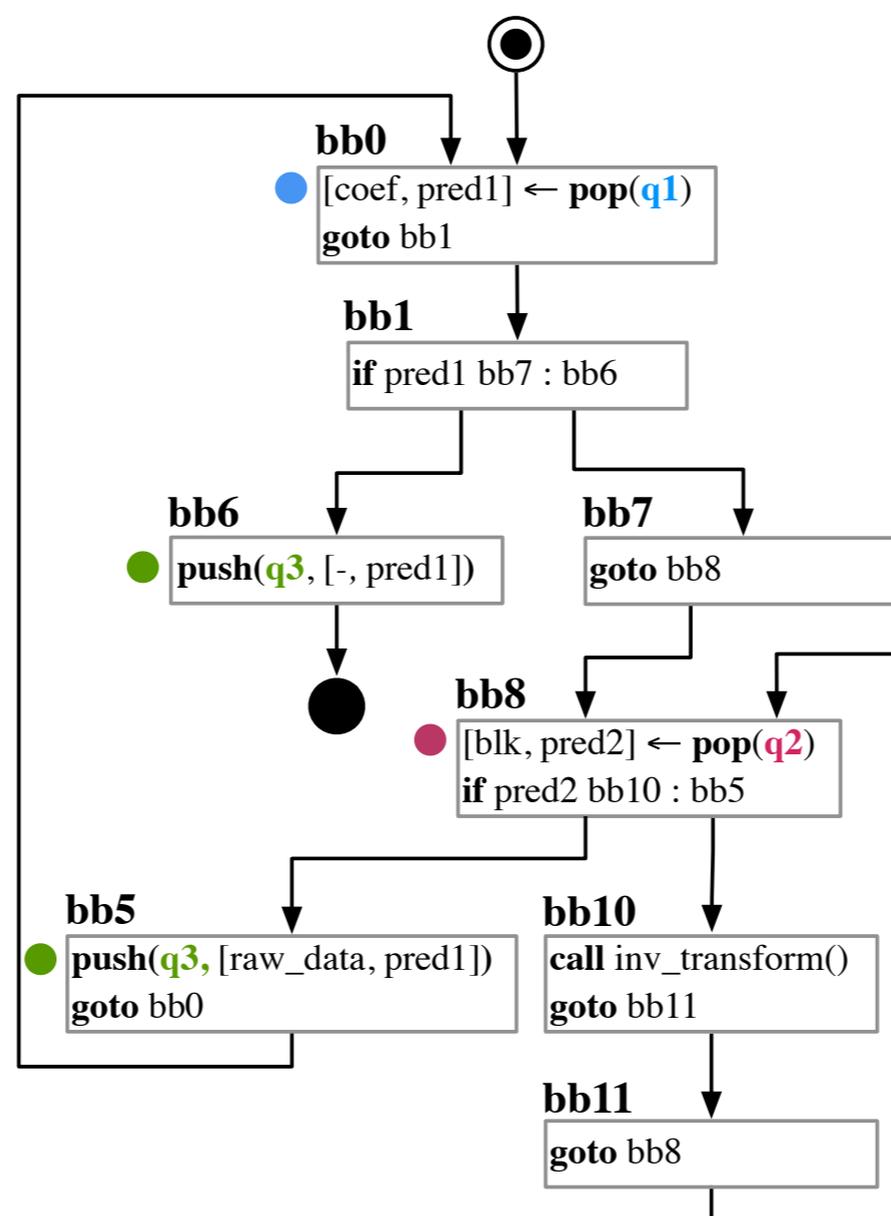
Communication



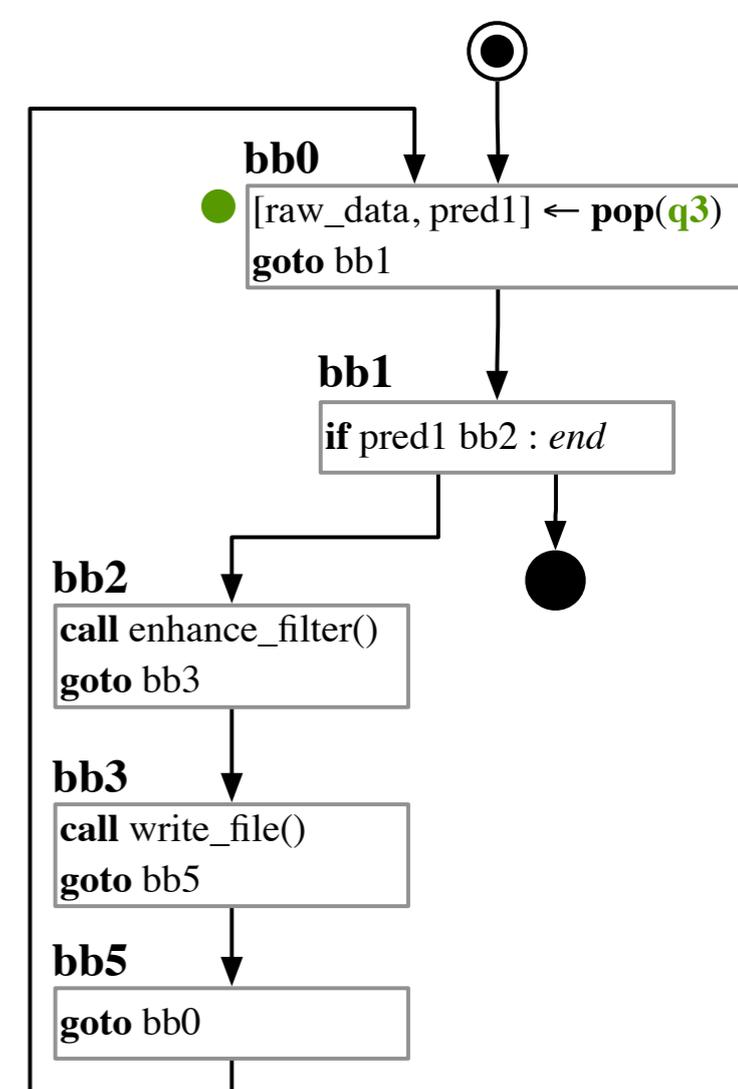
Stage 1



Stage 2



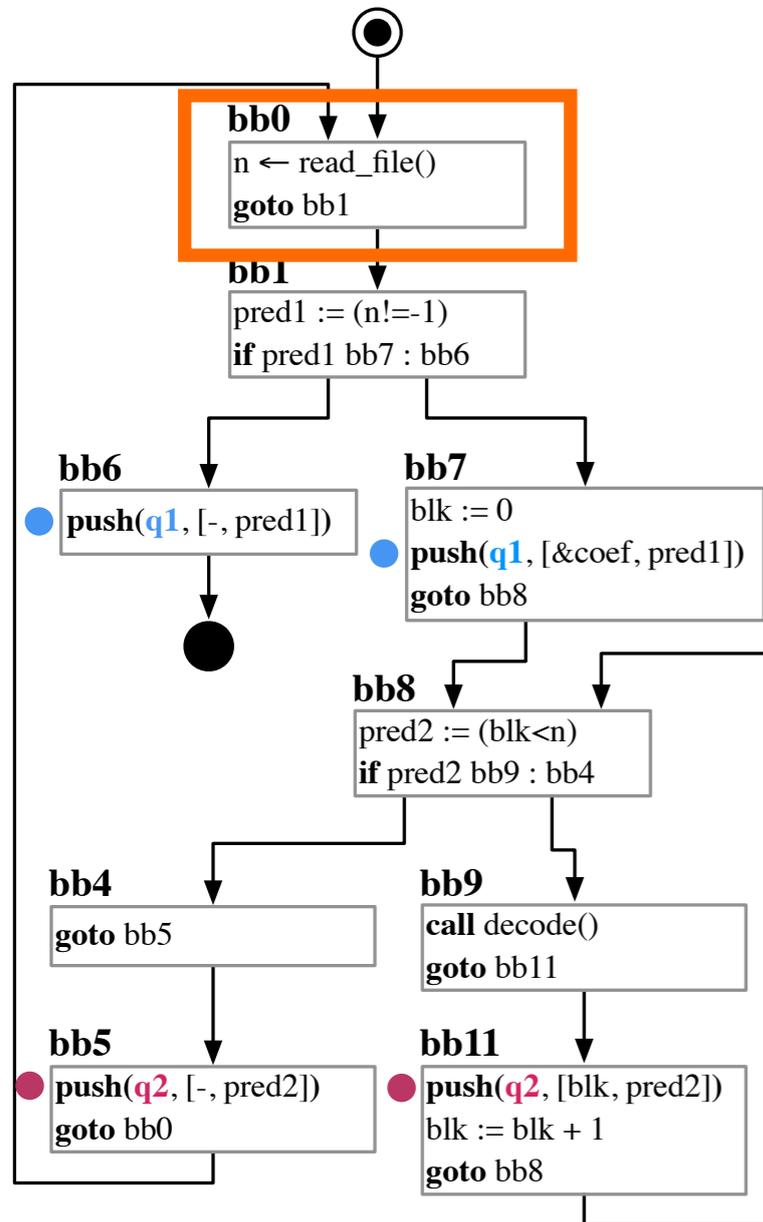
Stage 3



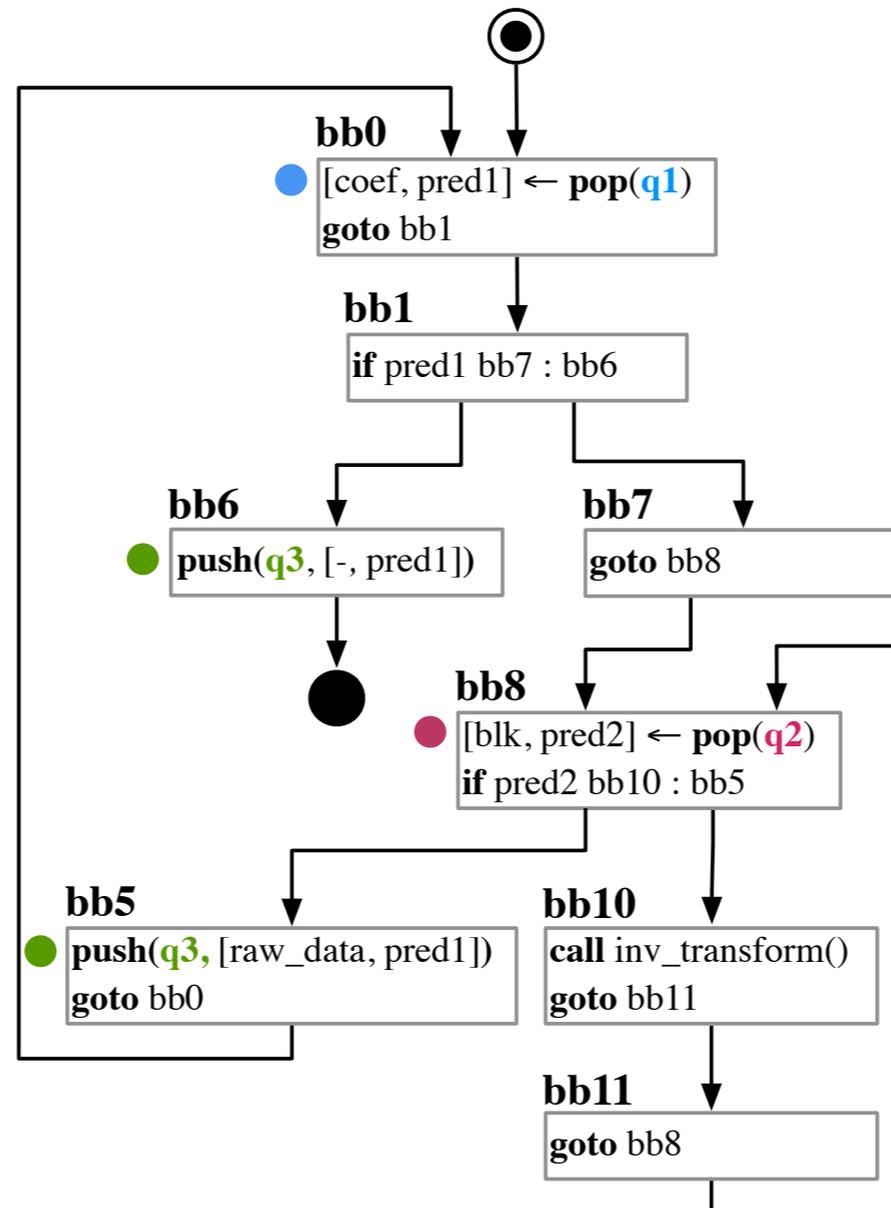
Communication



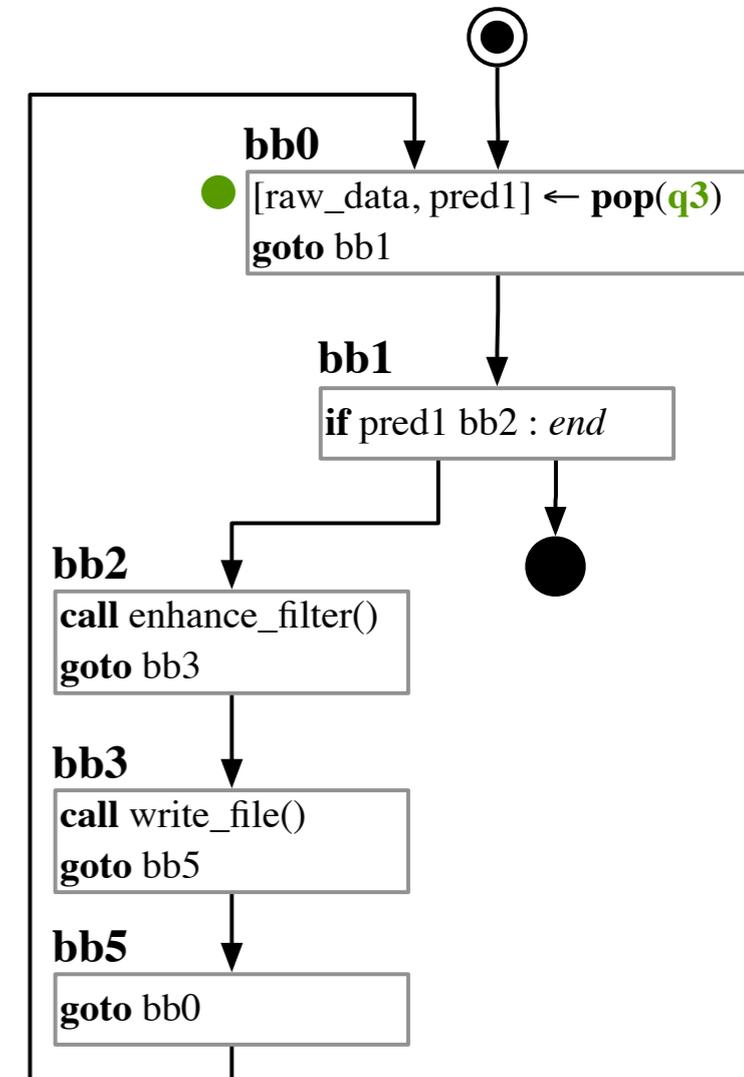
Stage 1



Stage 2



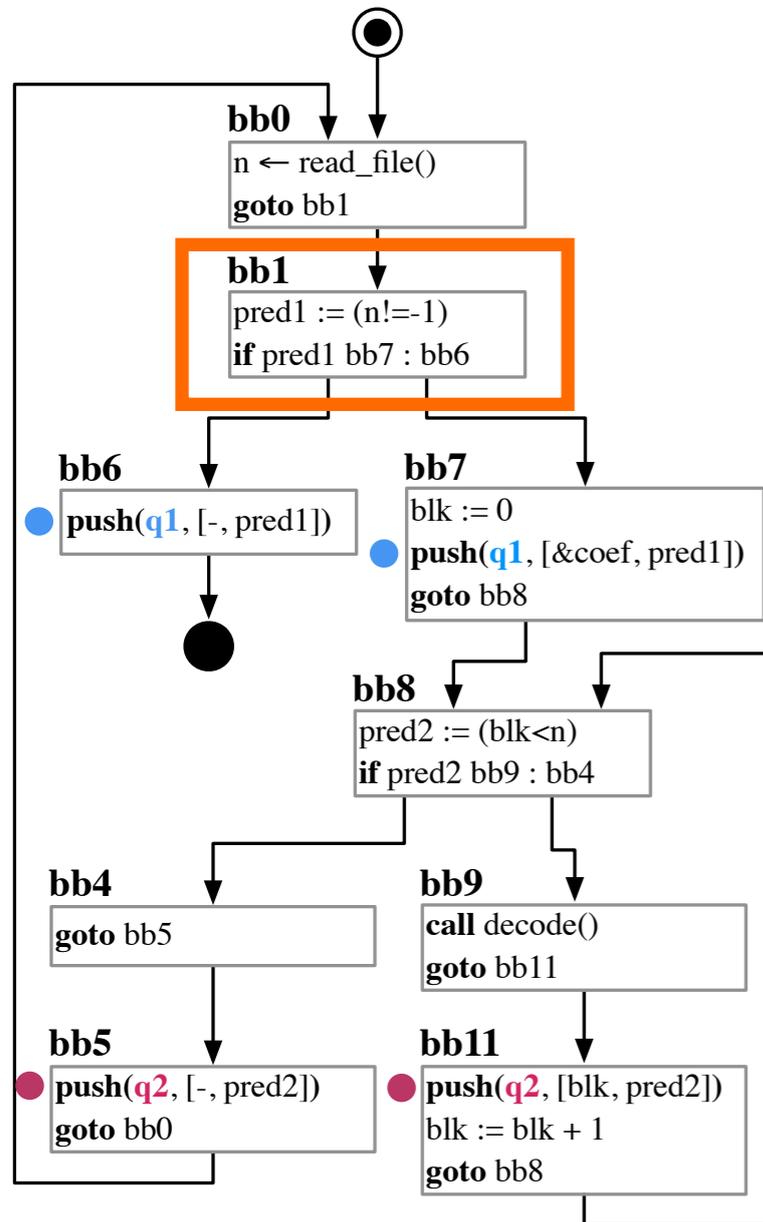
Stage 3



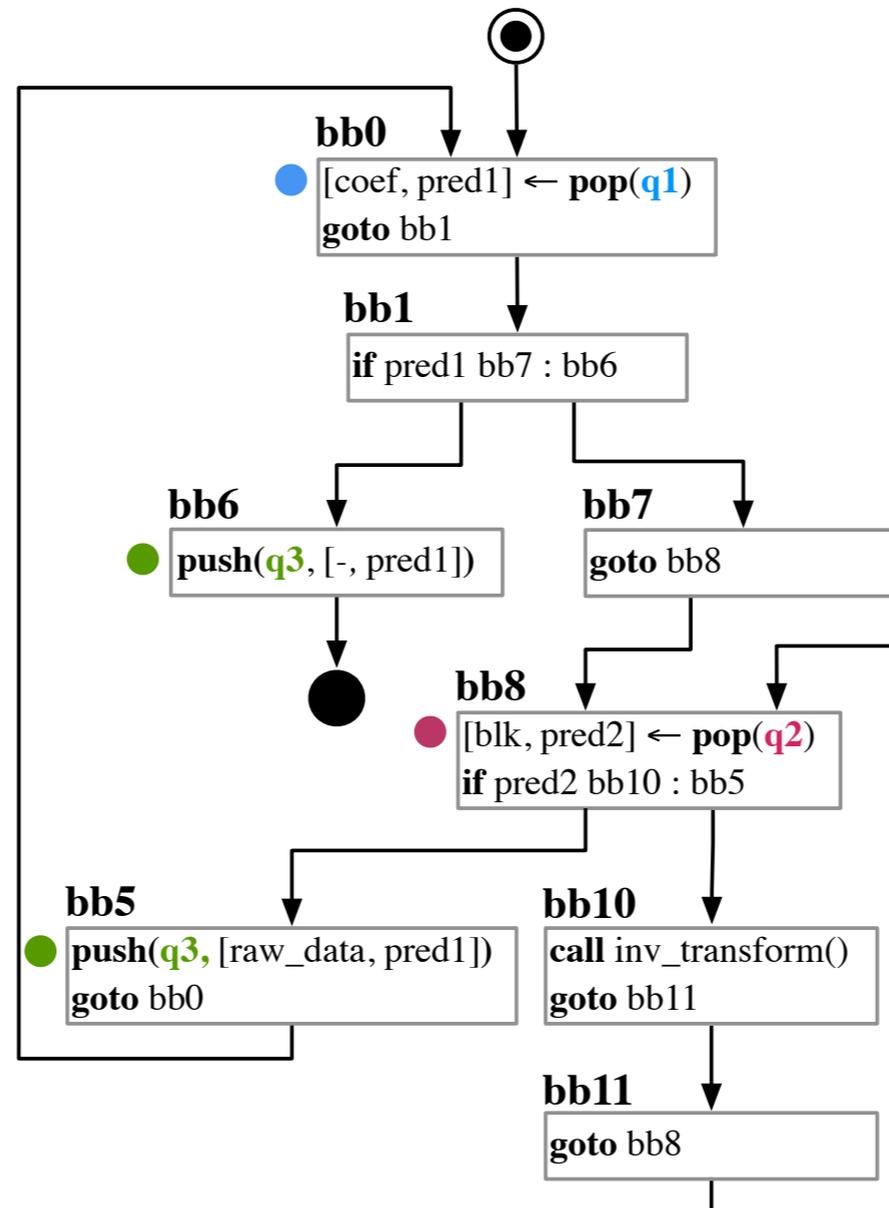
Communication



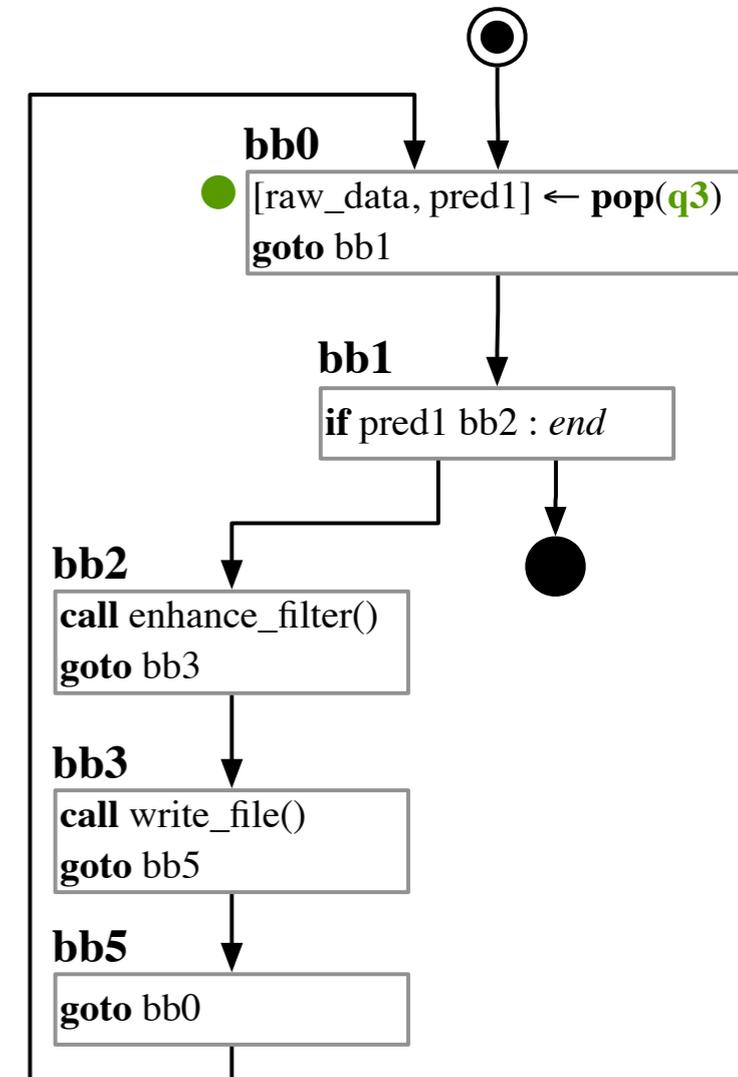
Stage 1



Stage 2



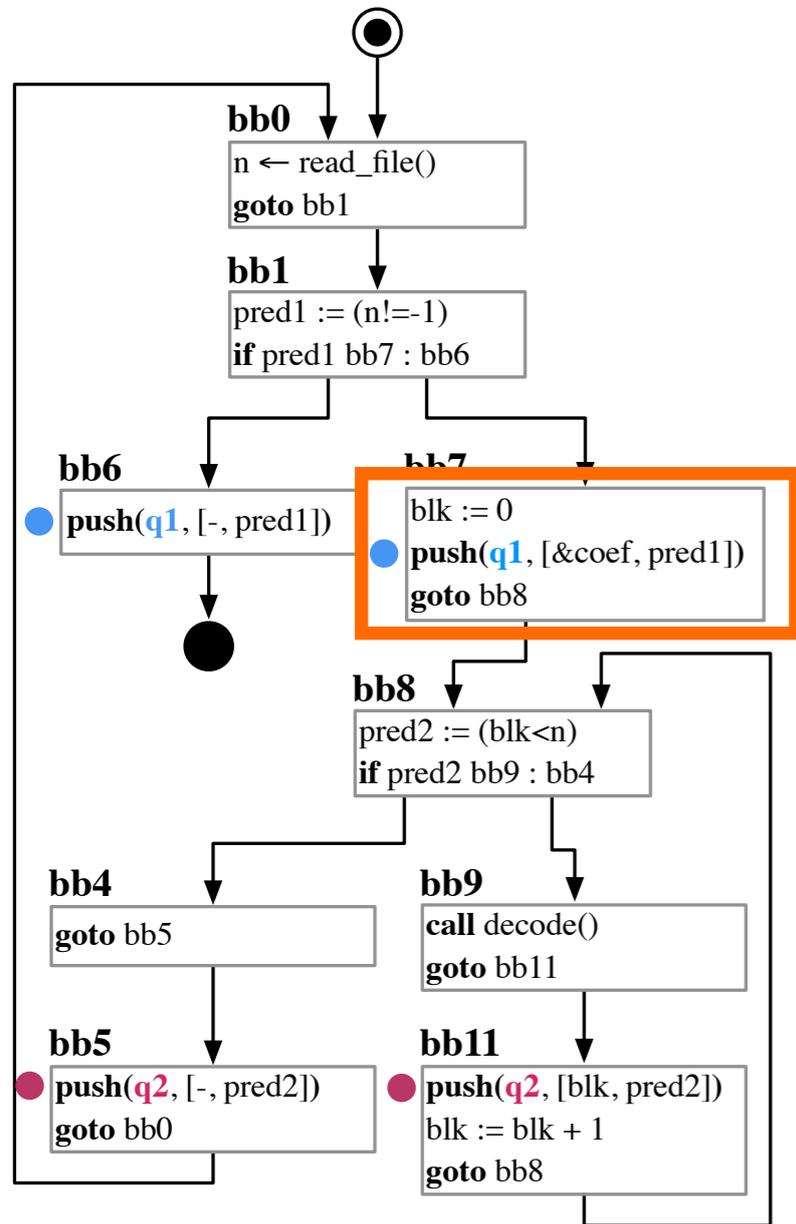
Stage 3



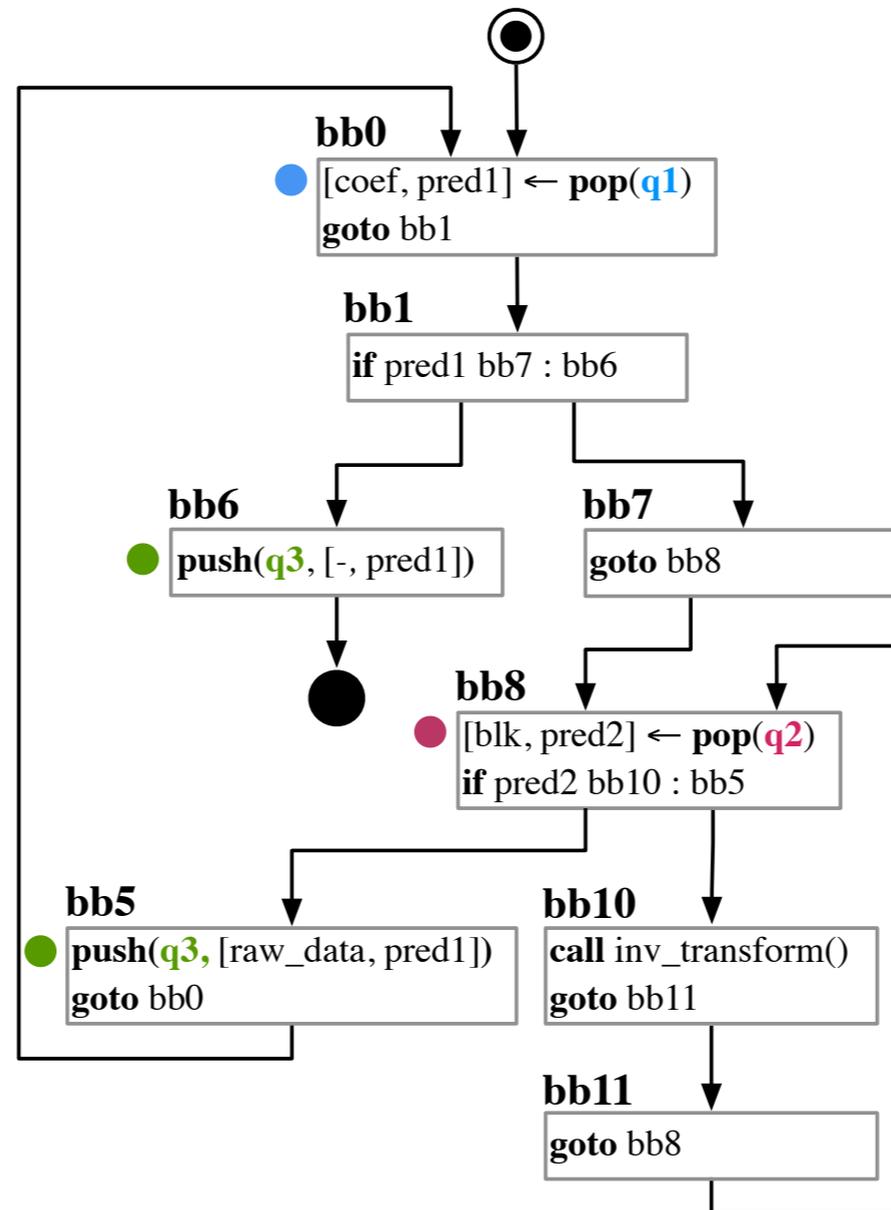
Communication



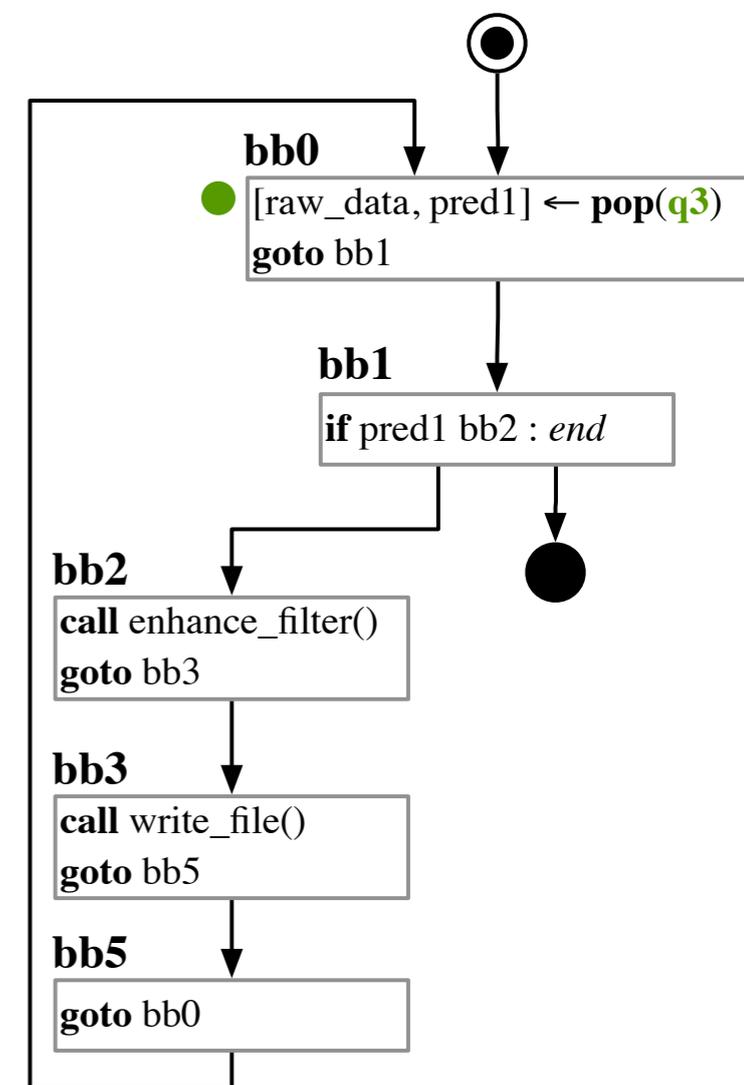
Stage 1



Stage 2



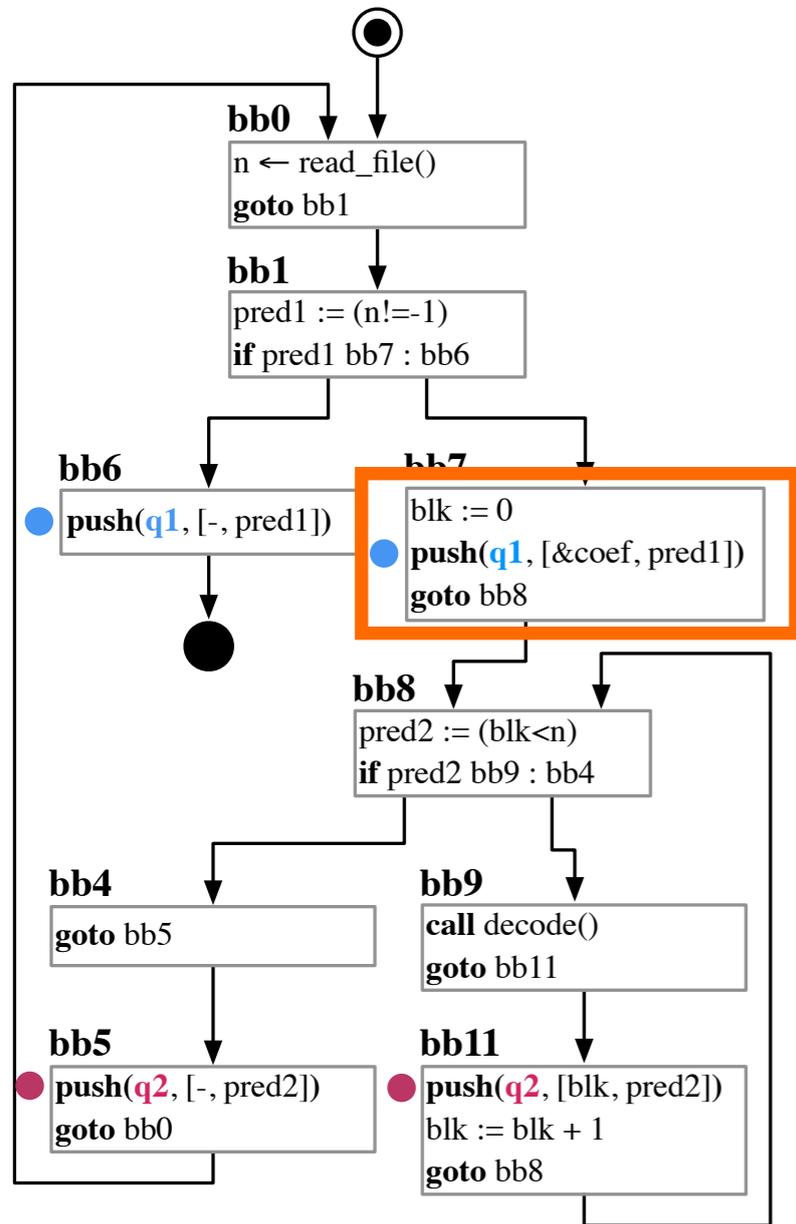
Stage 3



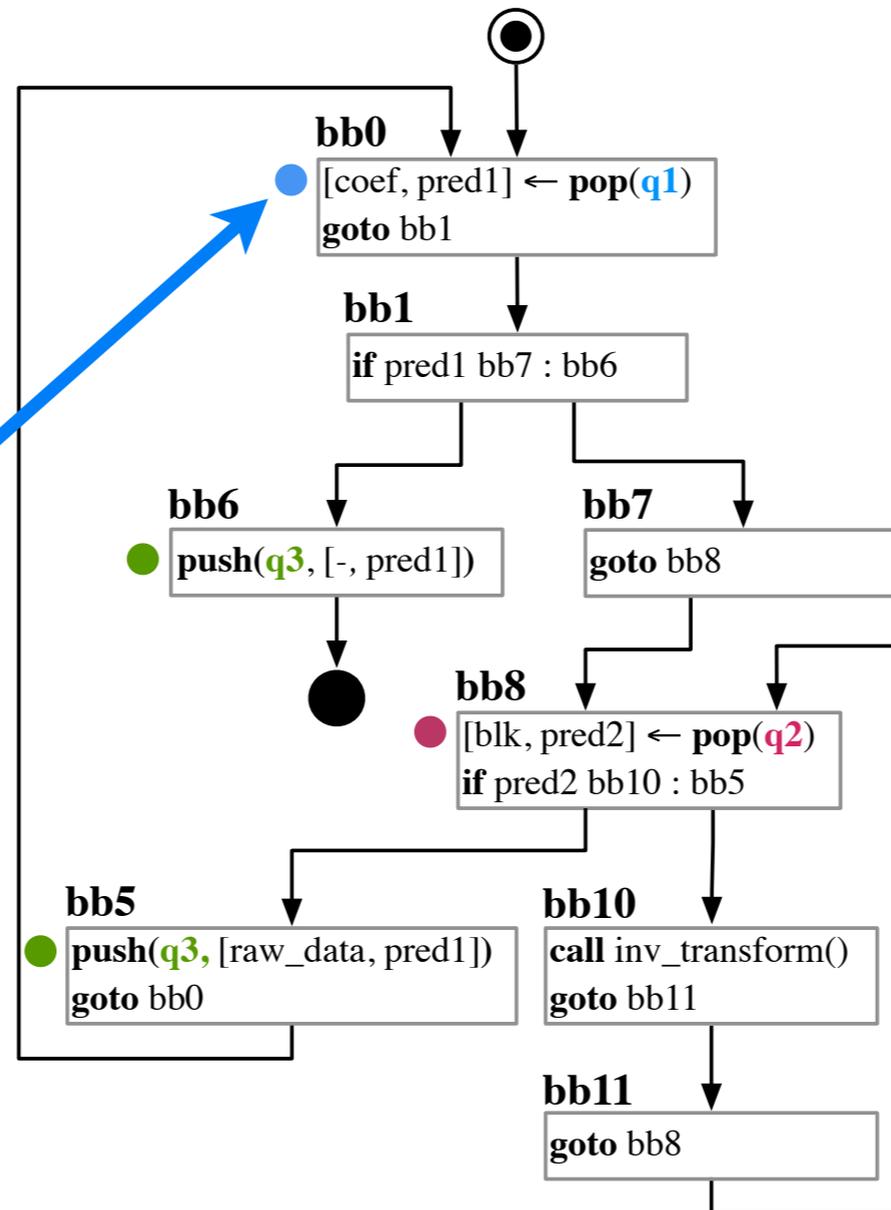
Communication



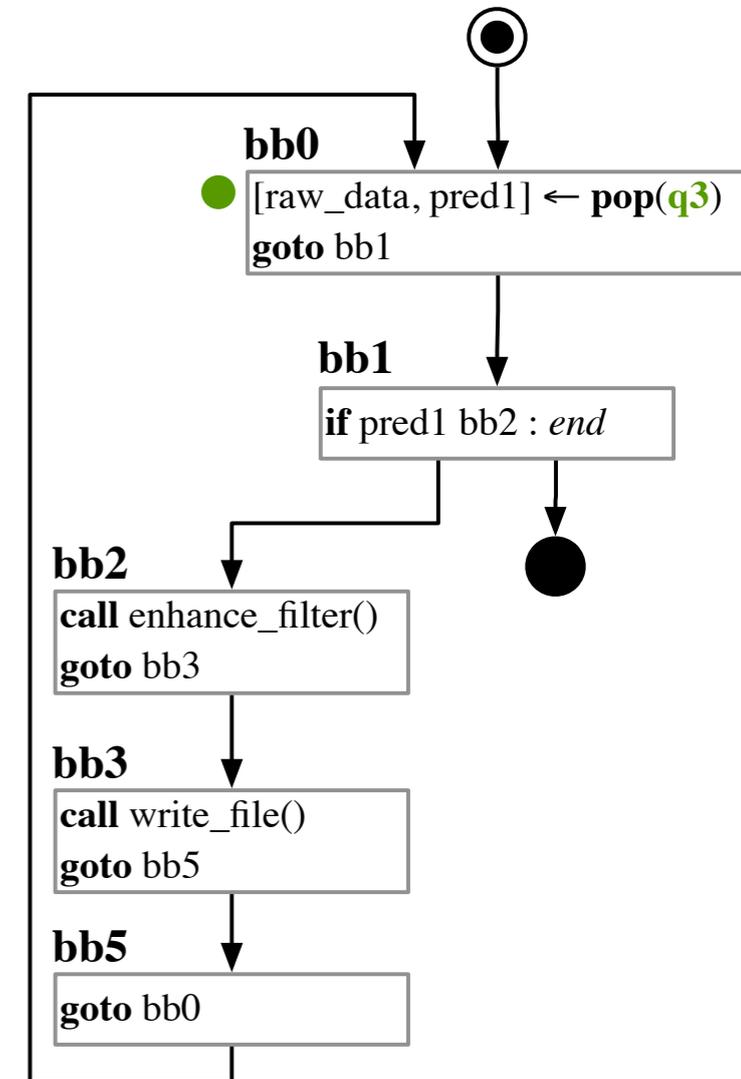
Stage 1



Stage 2



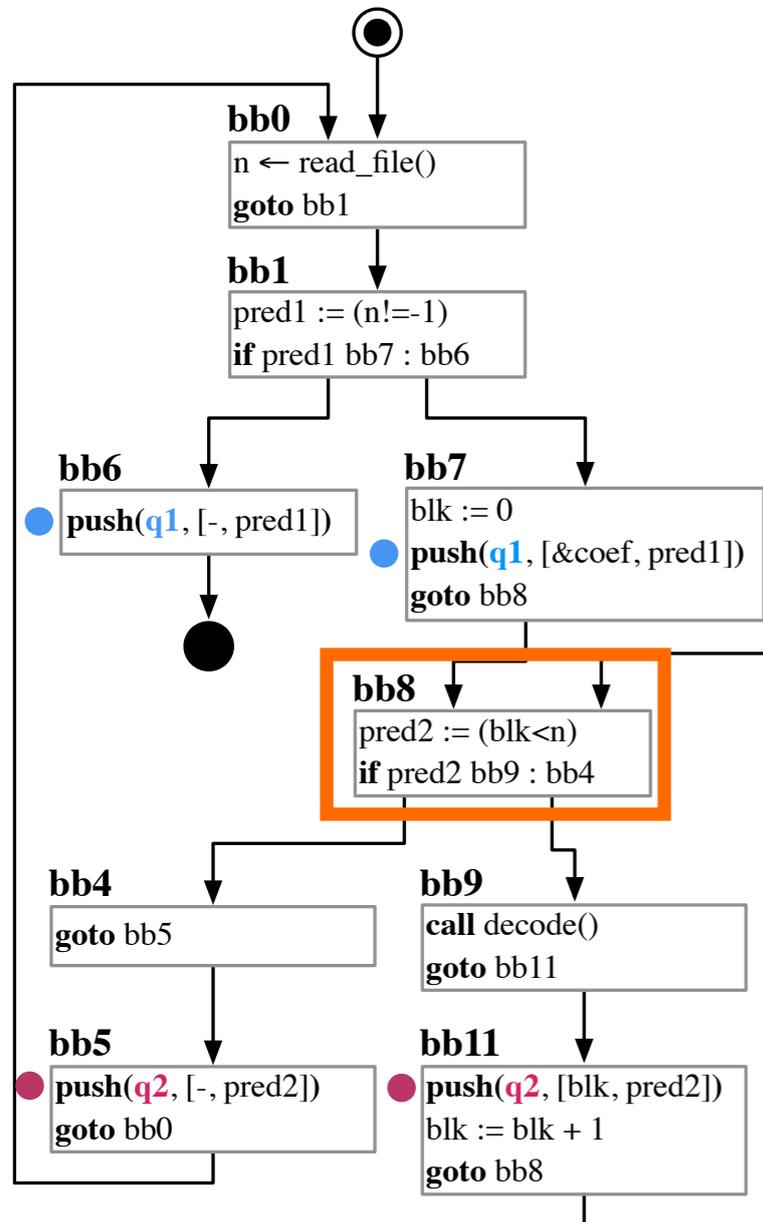
Stage 3



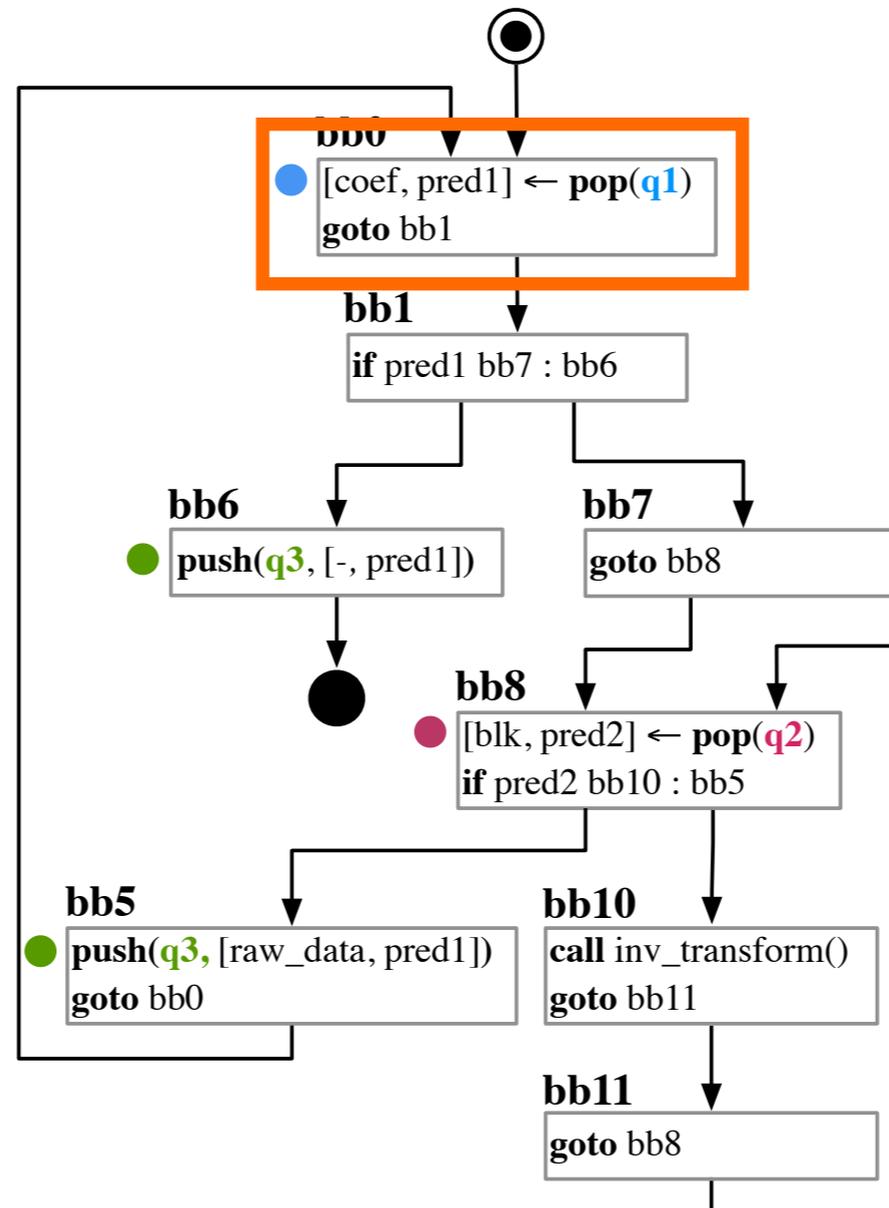
Communication



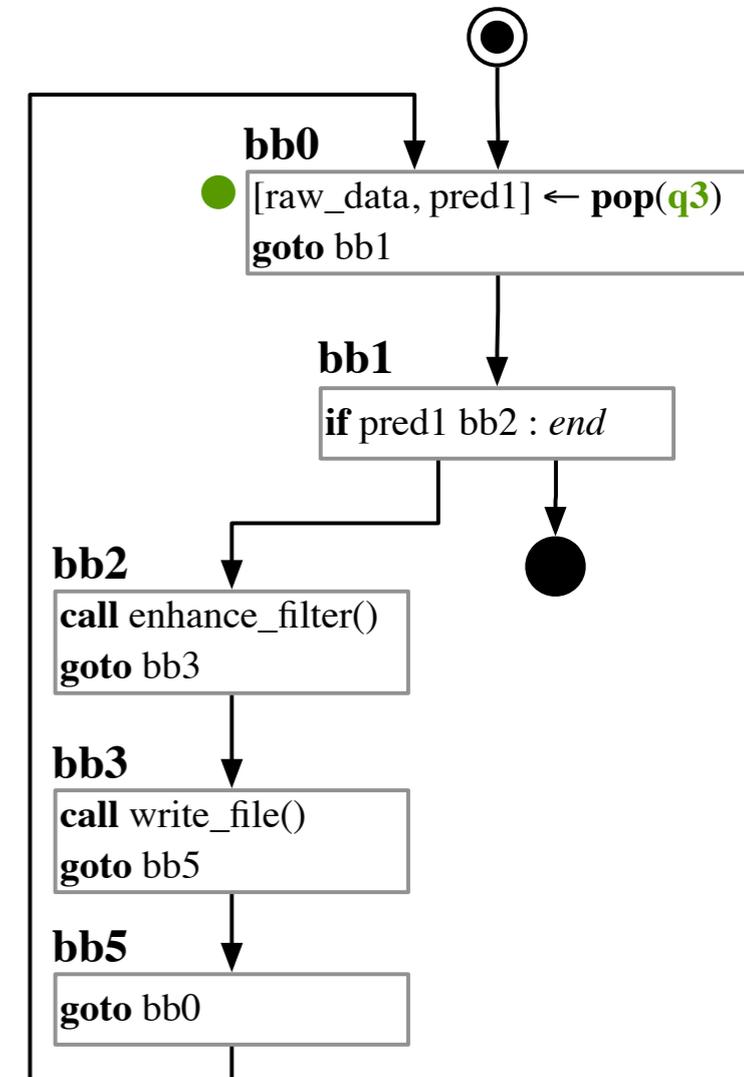
Stage 1



Stage 2



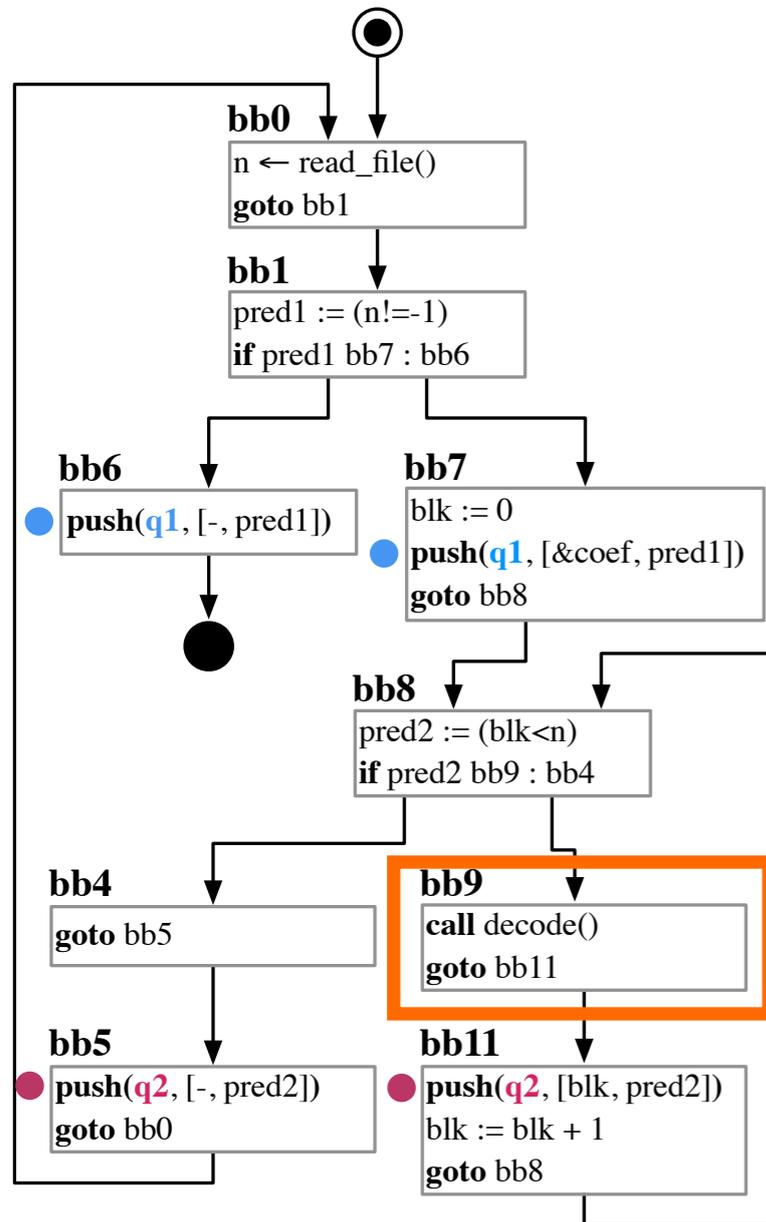
Stage 3



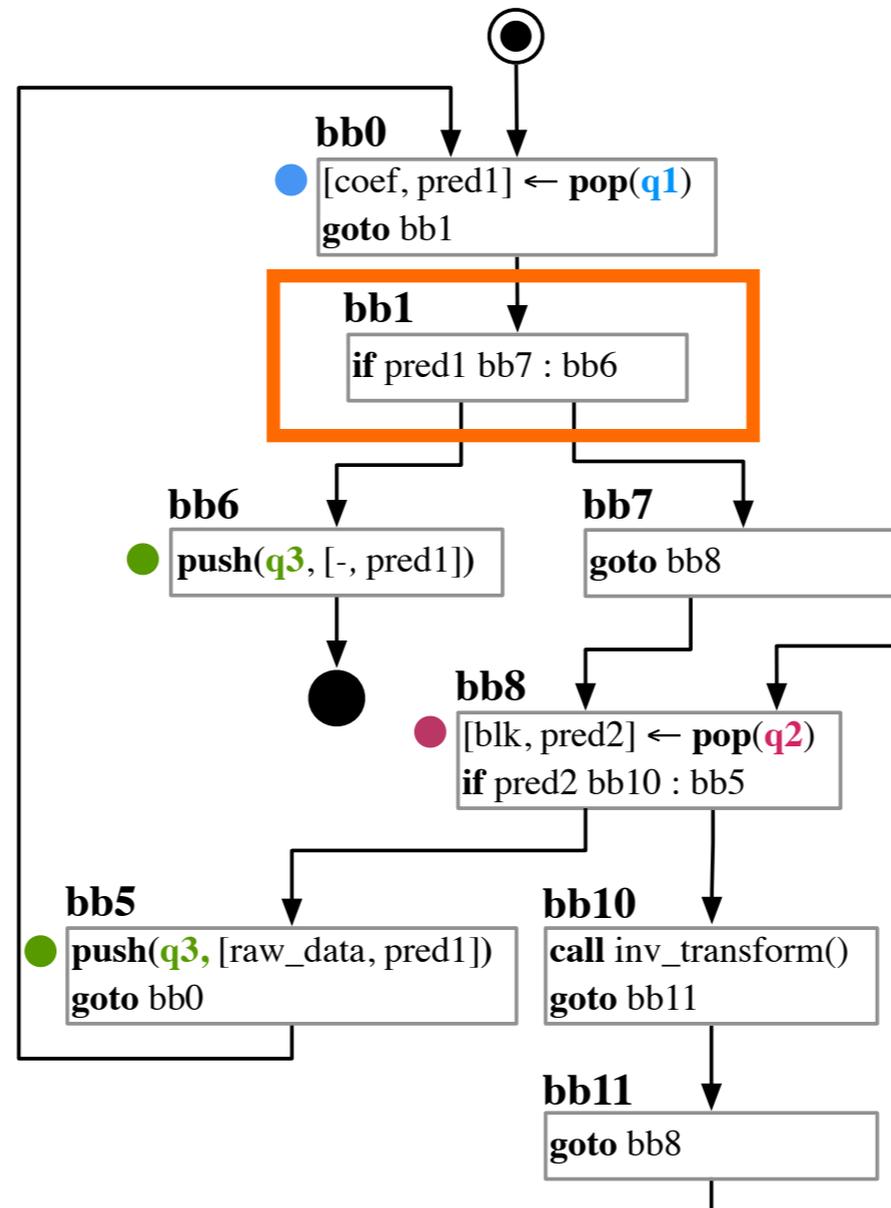
Communication



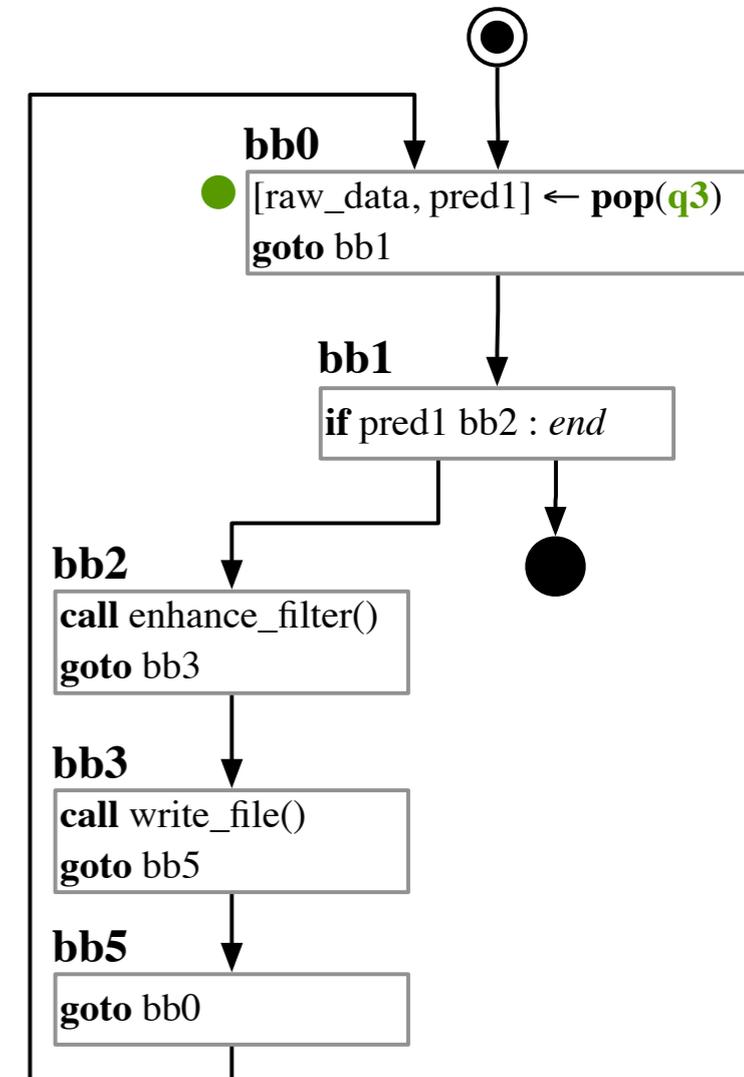
Stage 1



Stage 2



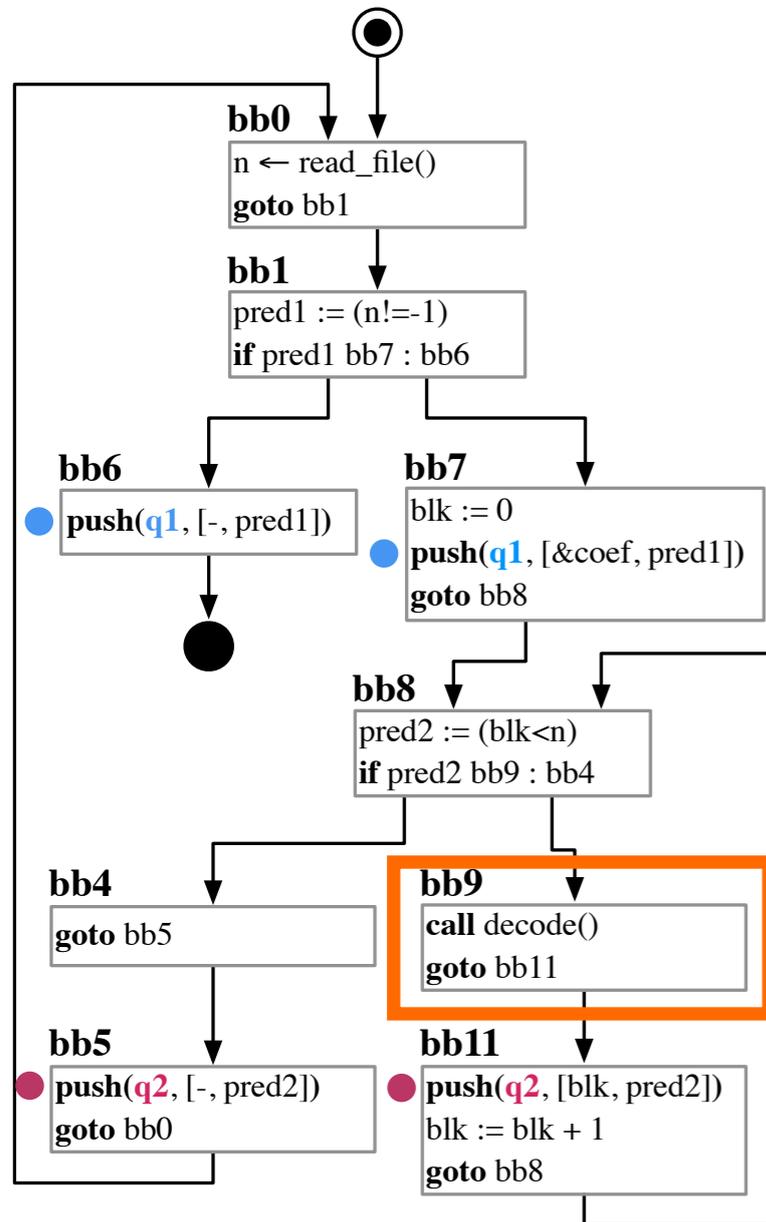
Stage 3



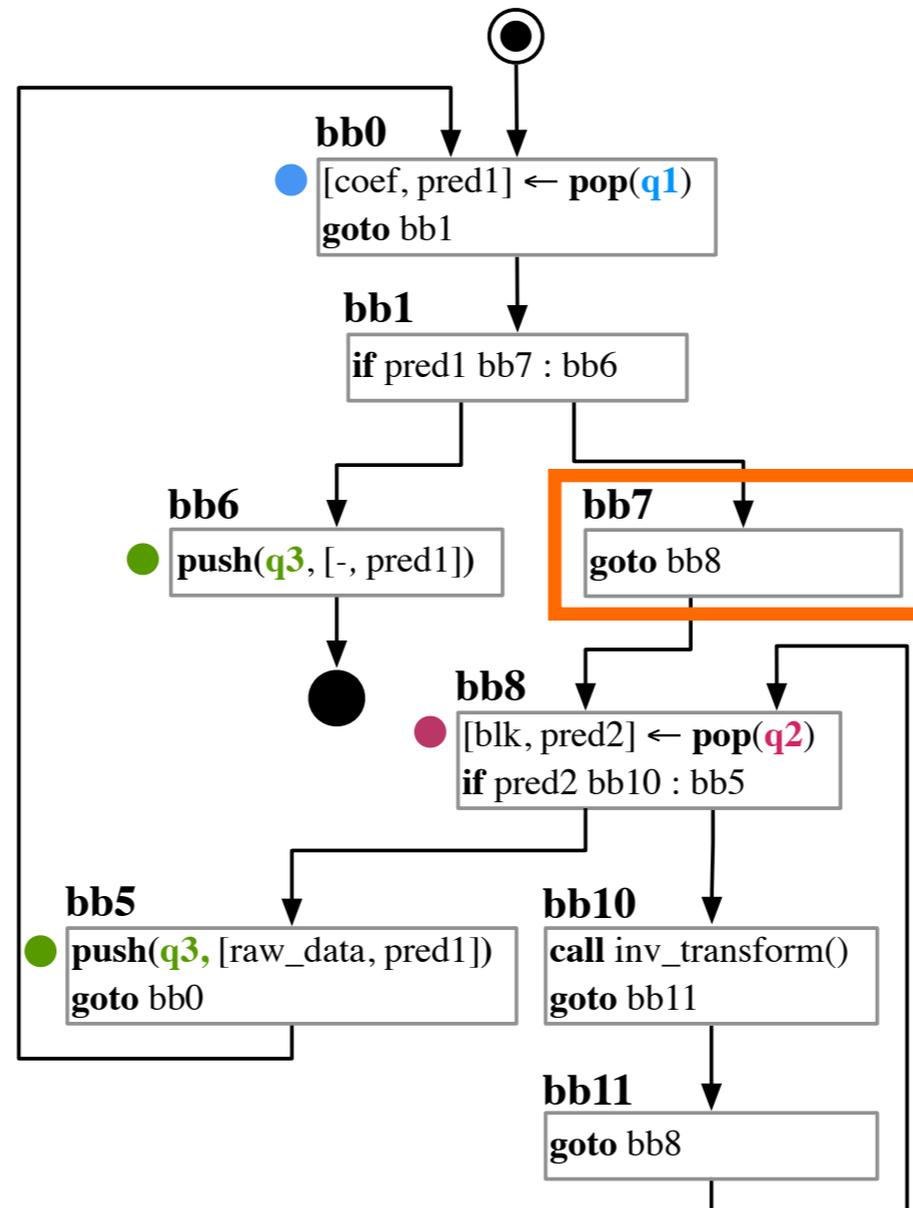
Communication



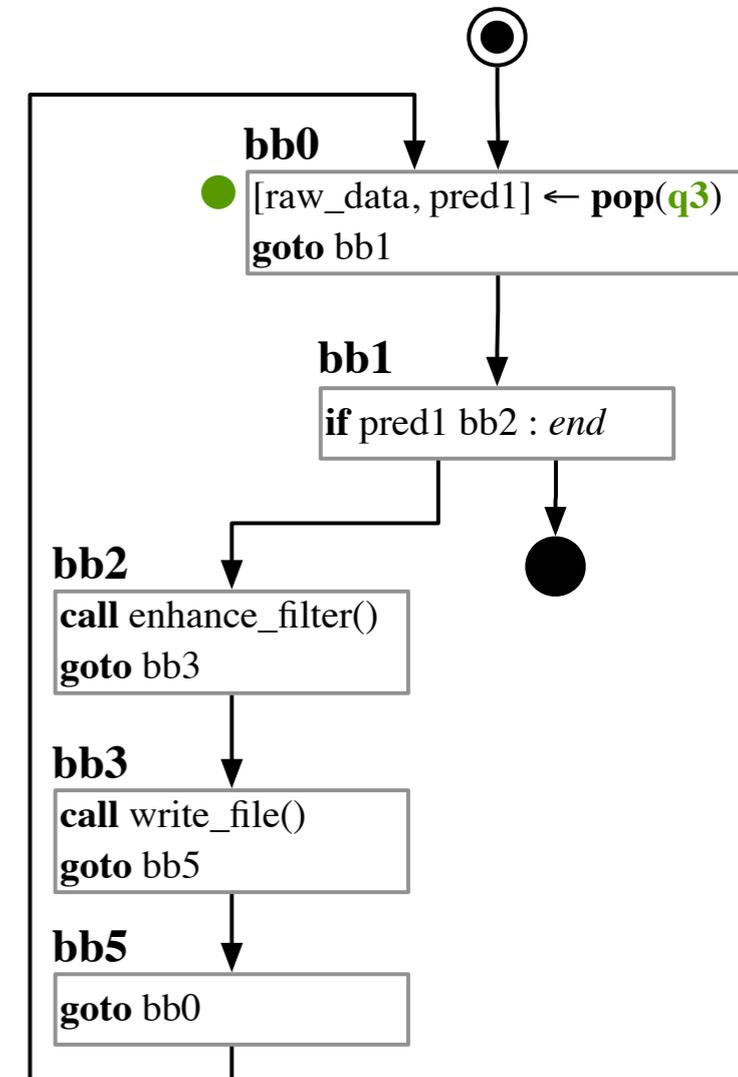
Stage 1



Stage 2



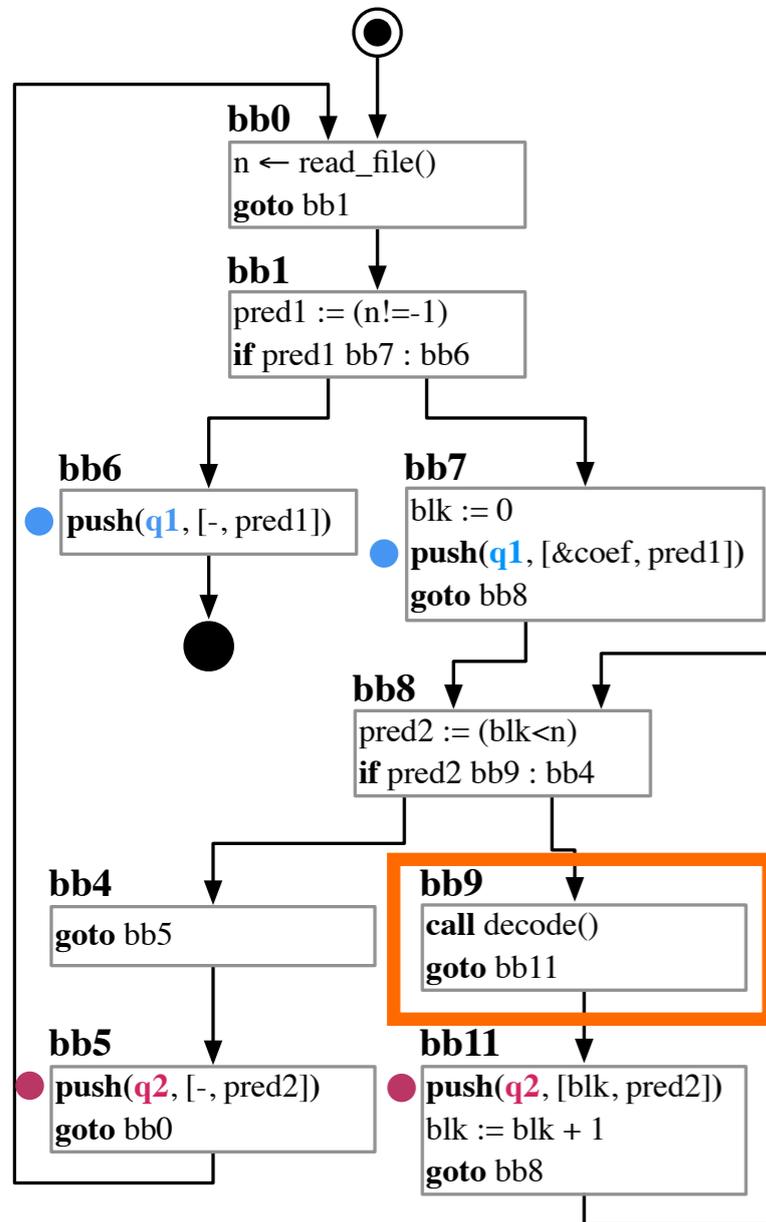
Stage 3



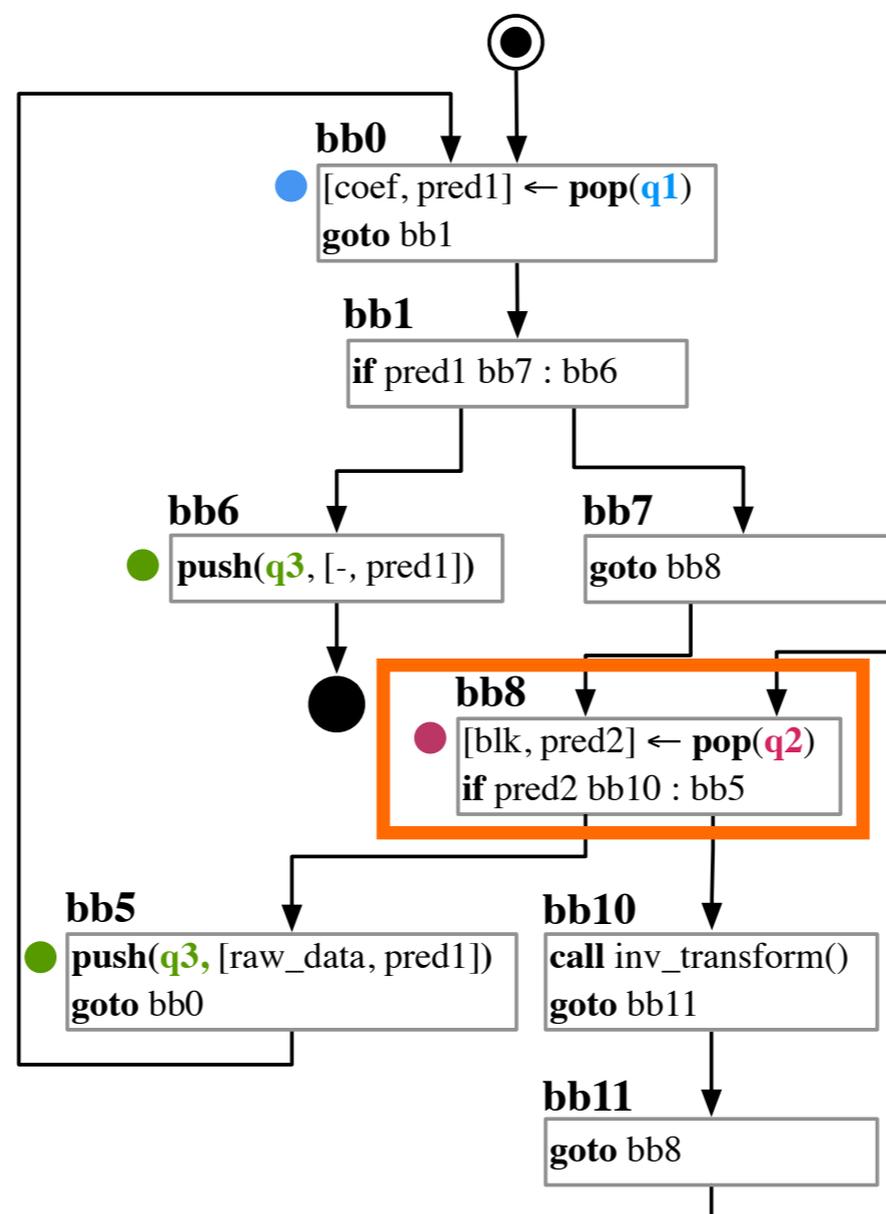
Communication



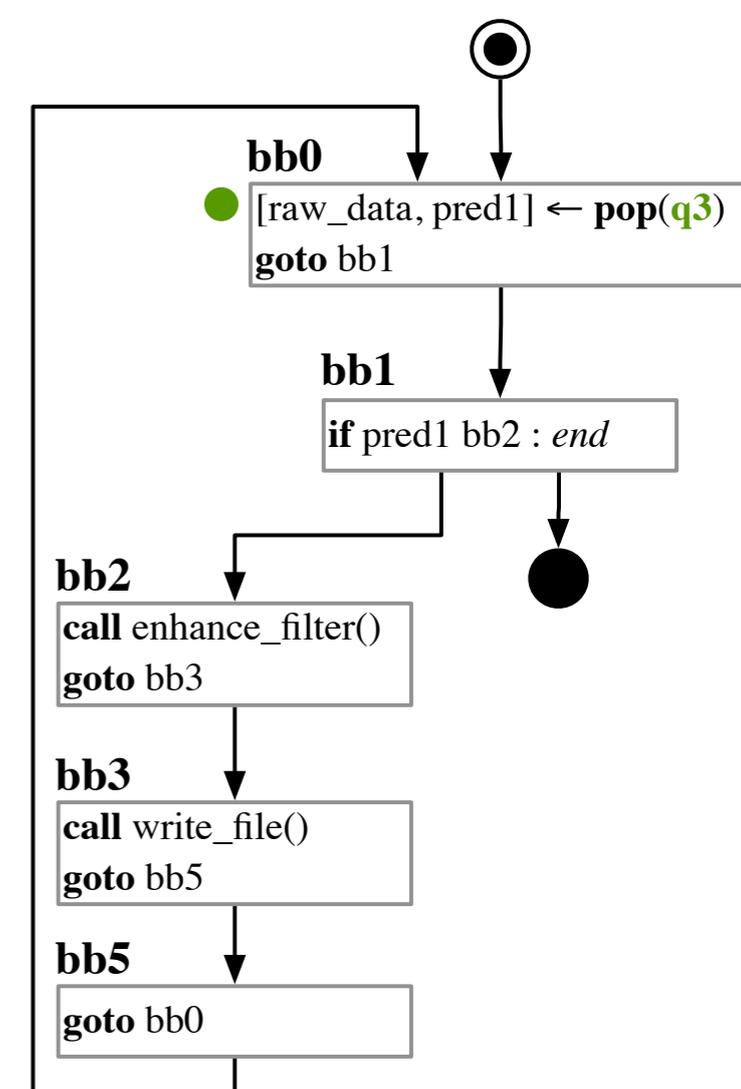
Stage 1



Stage 2



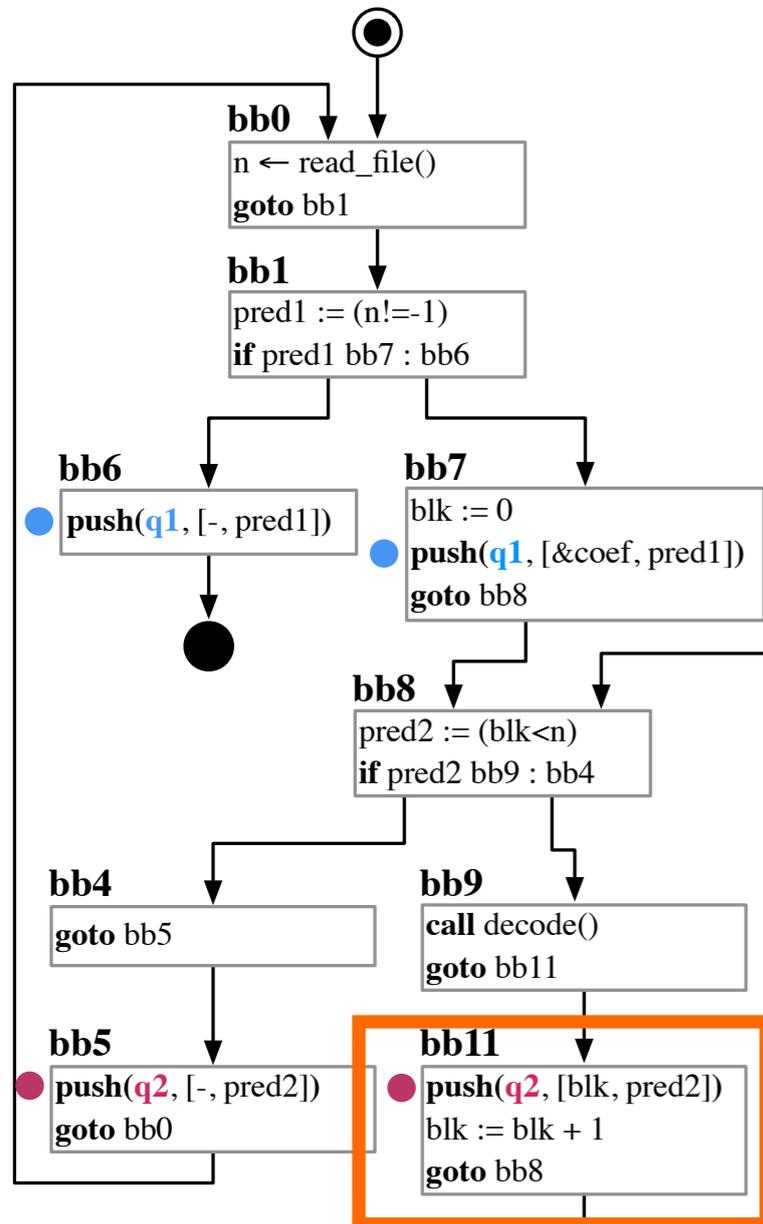
Stage 3



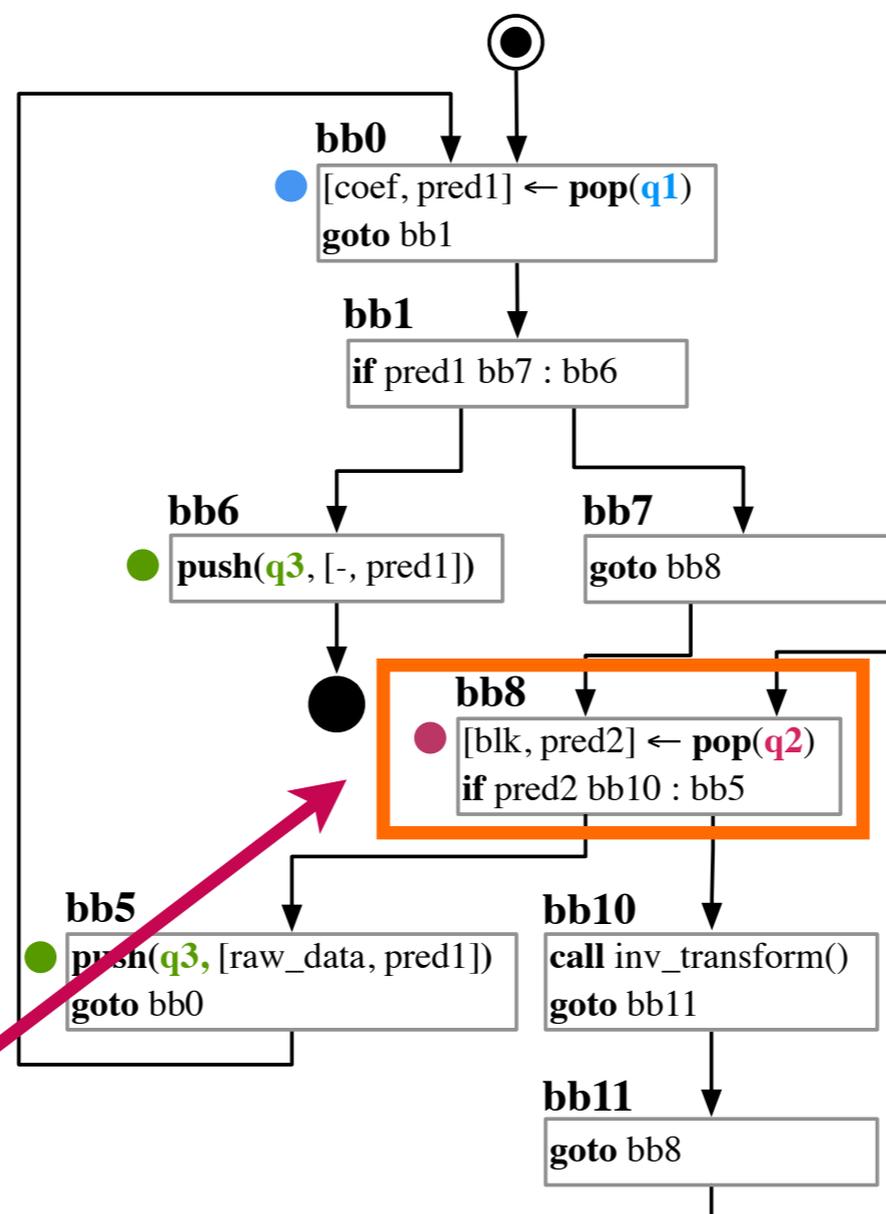
Communication



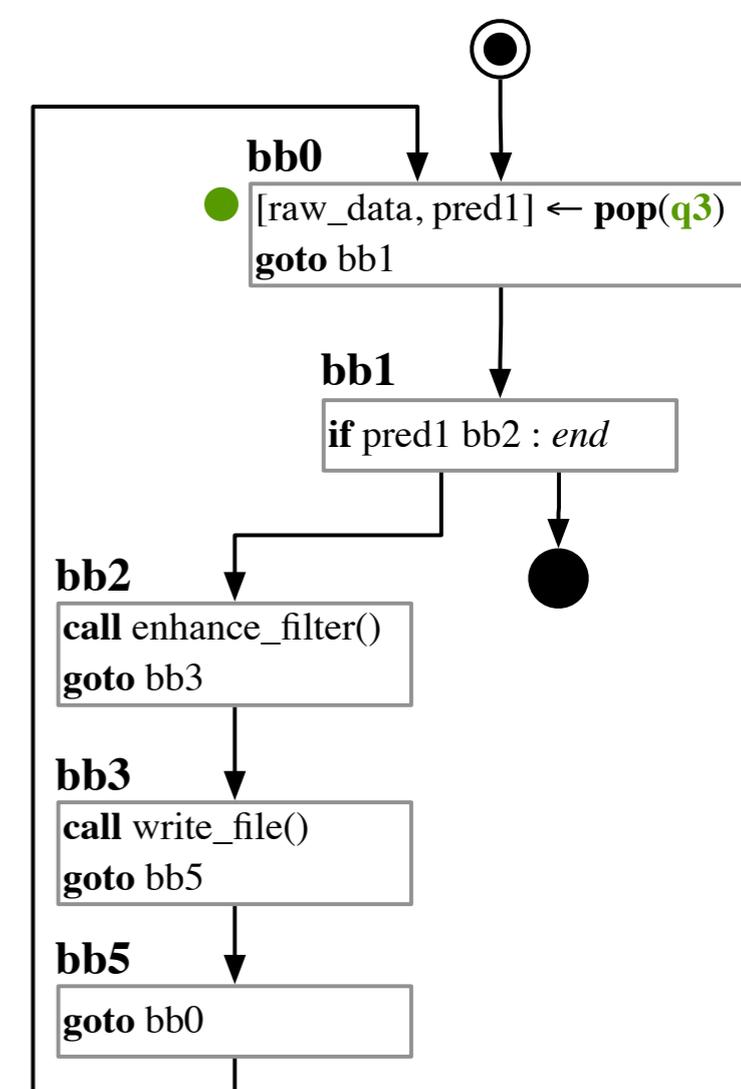
Stage 1



Stage 2



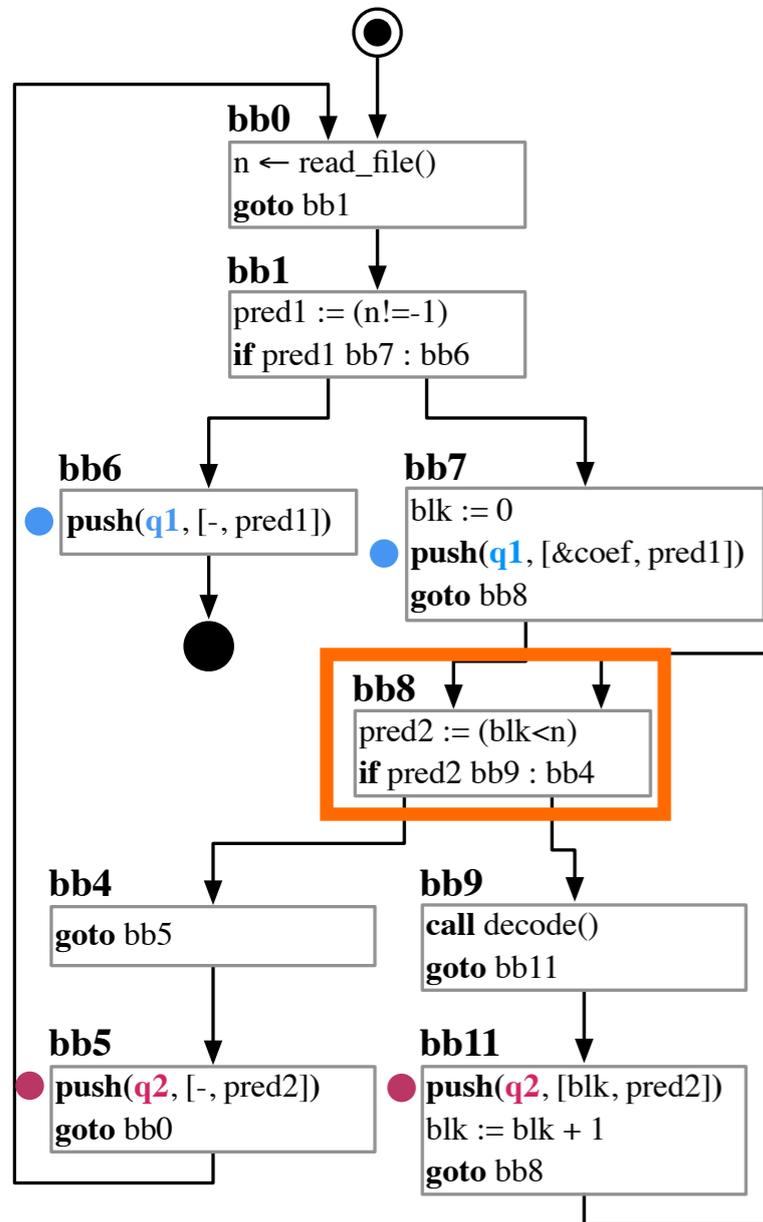
Stage 3



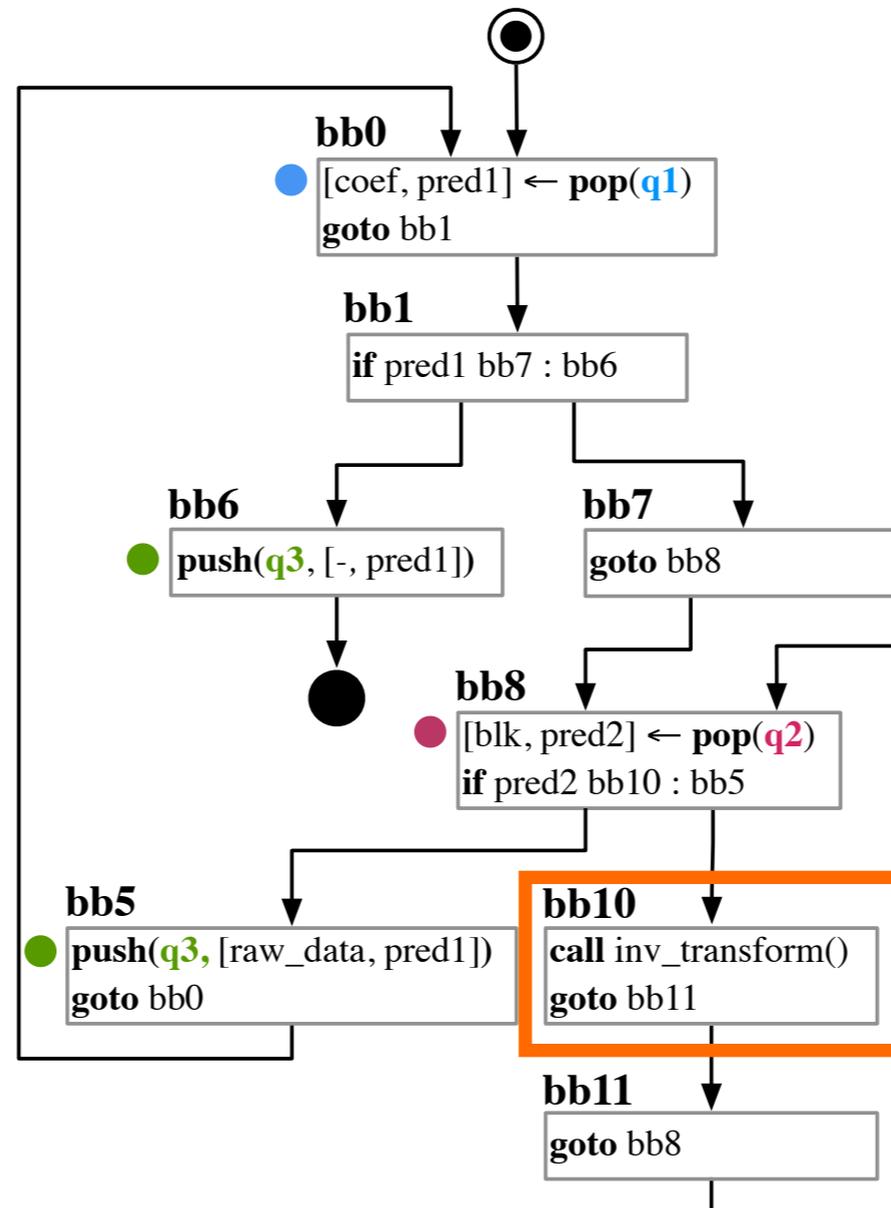
Communication



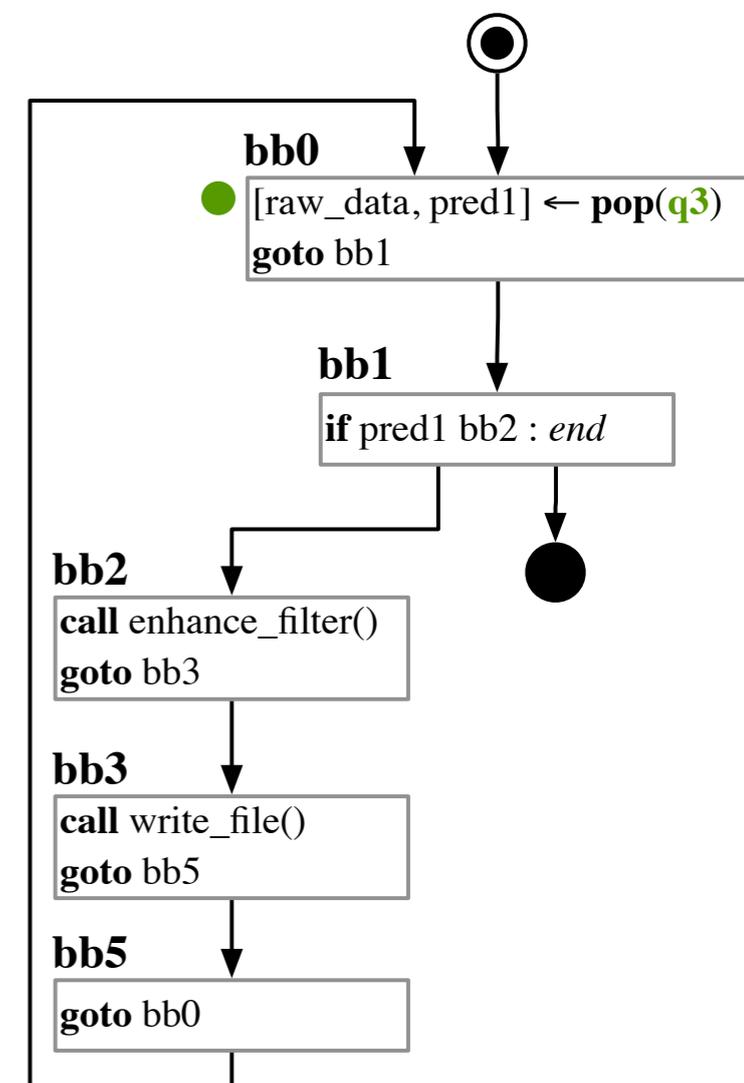
Stage 1



Stage 2



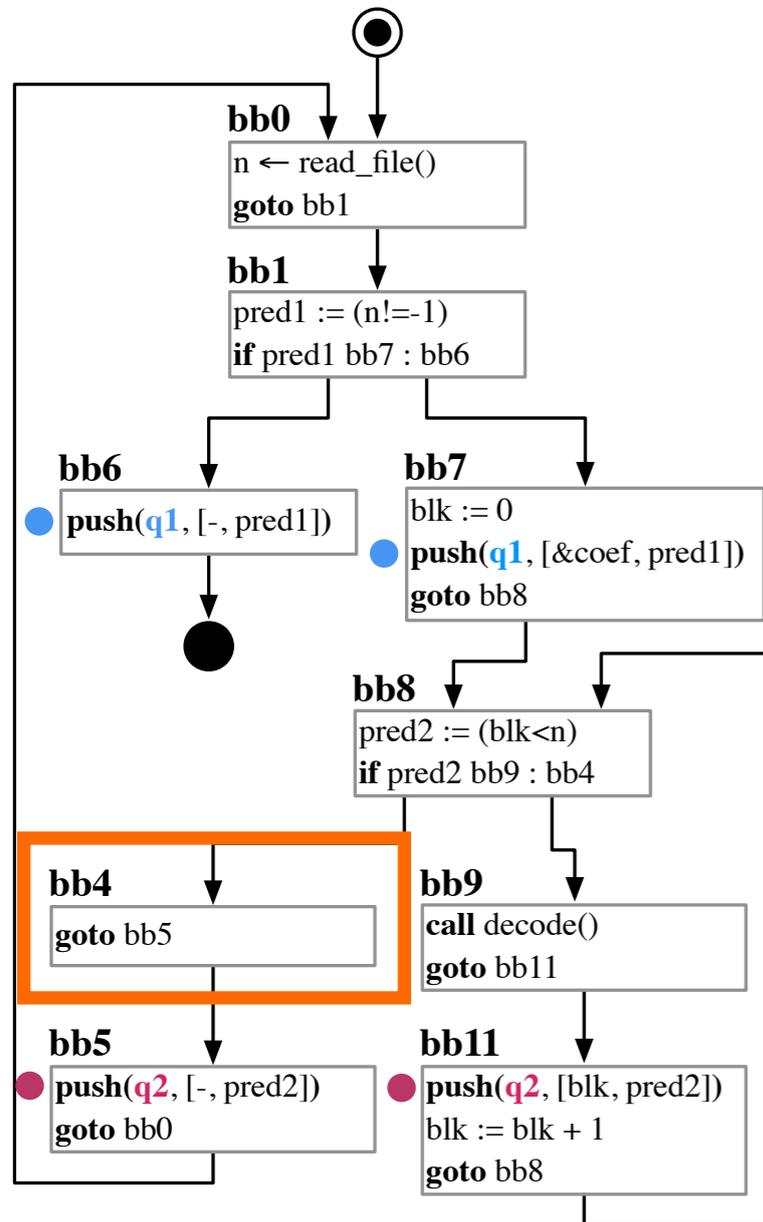
Stage 3



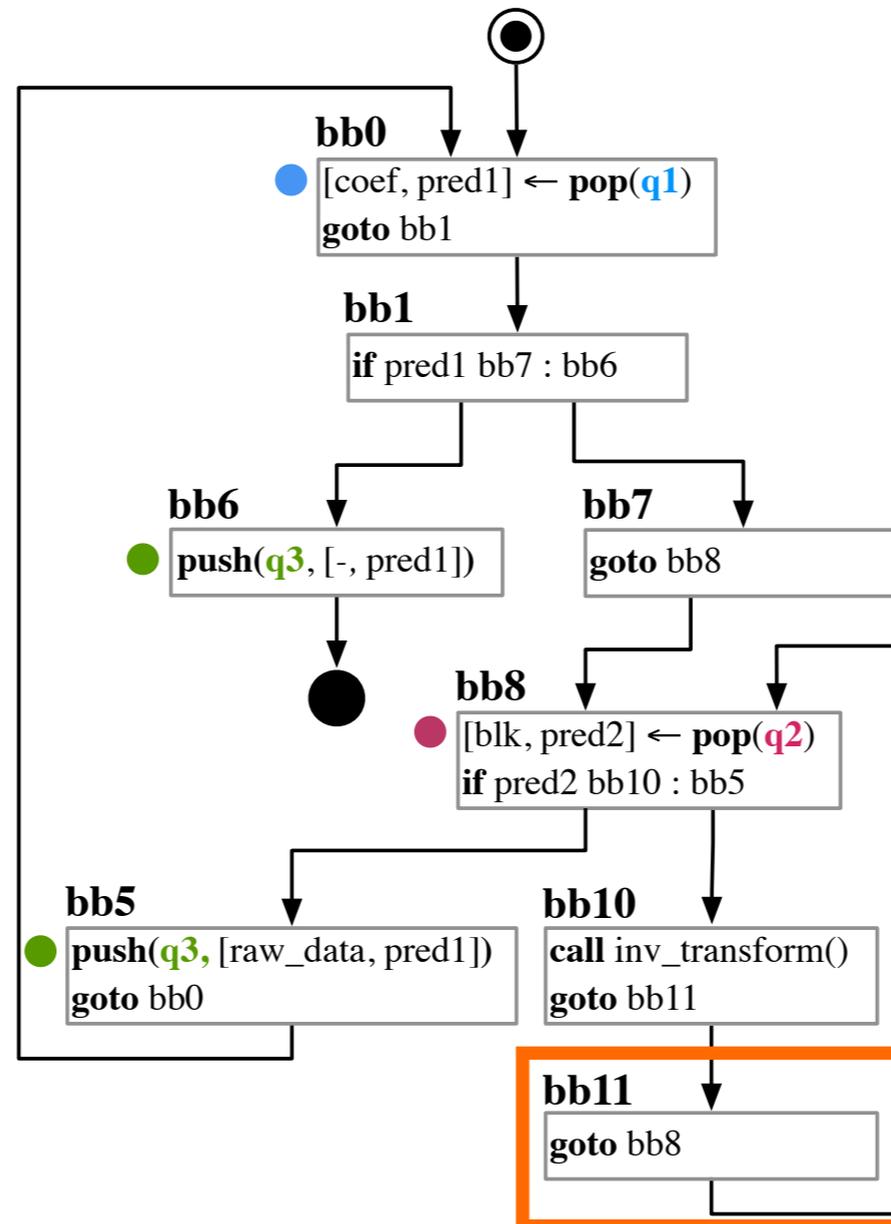
Communication



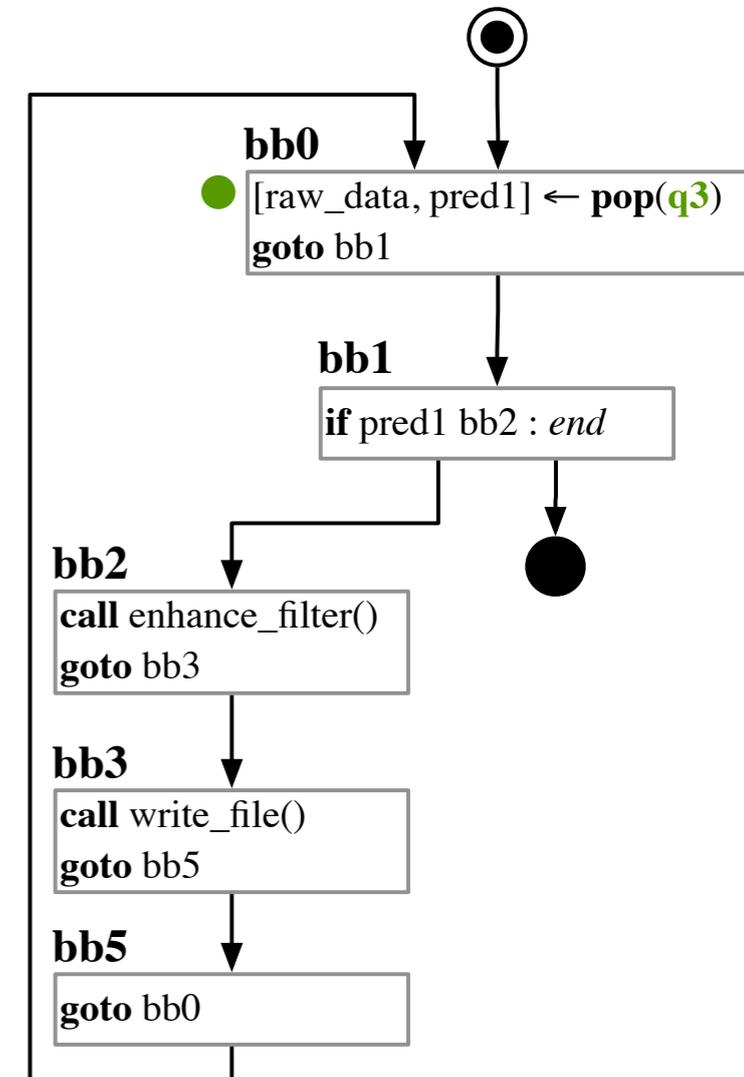
Stage 1



Stage 2



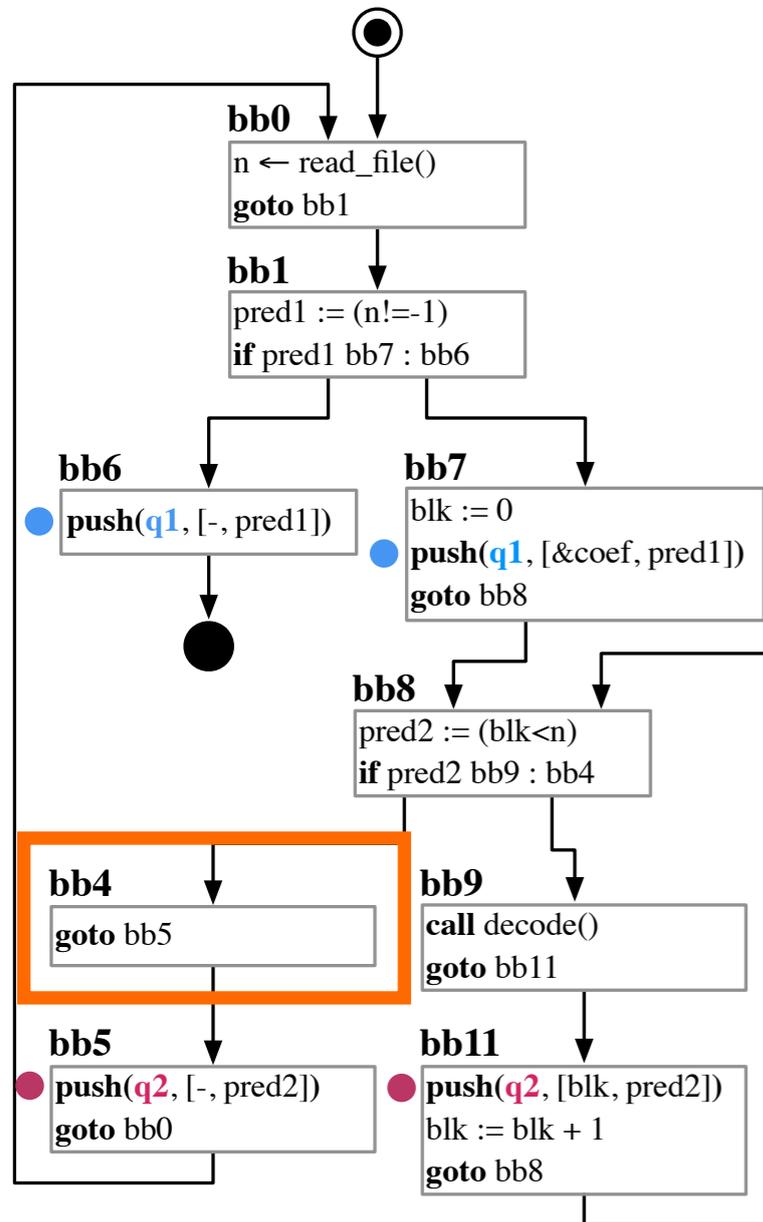
Stage 3



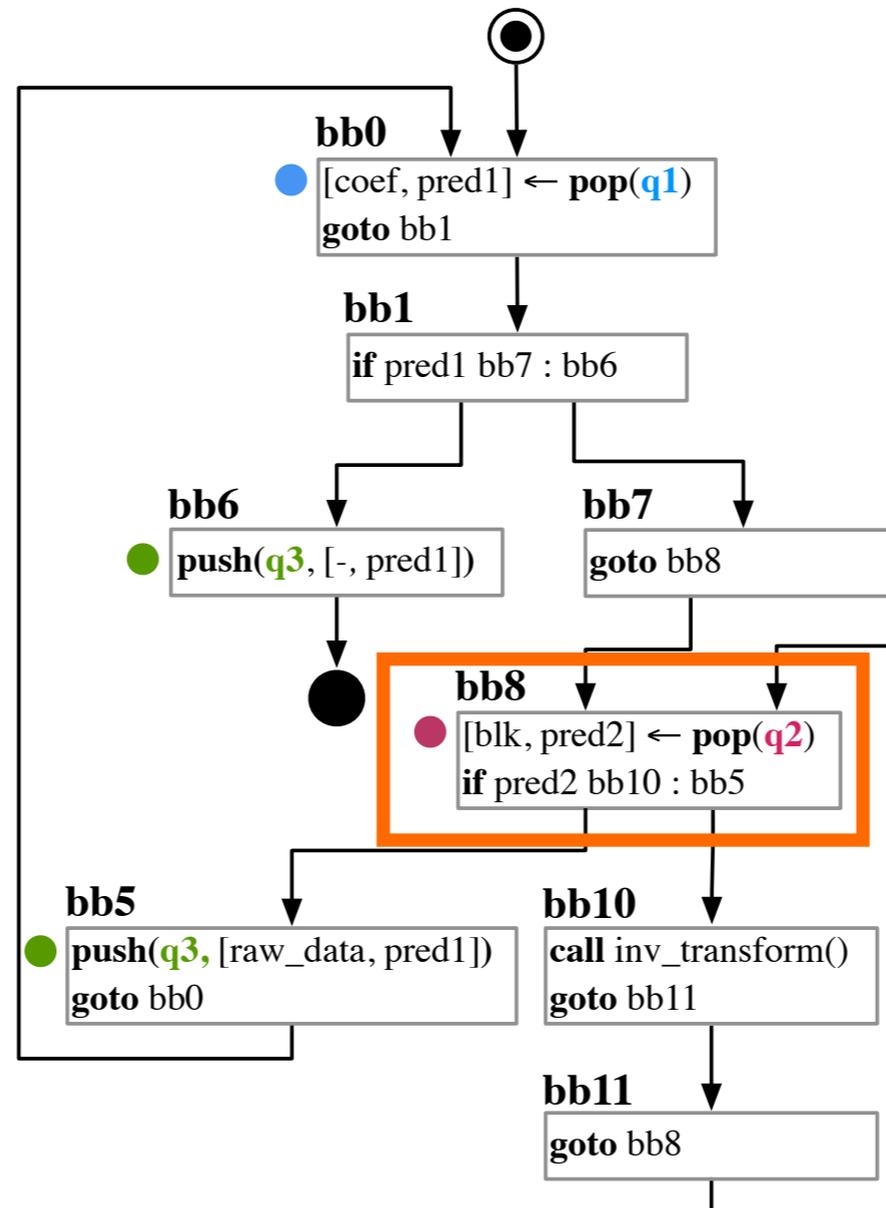
Communication



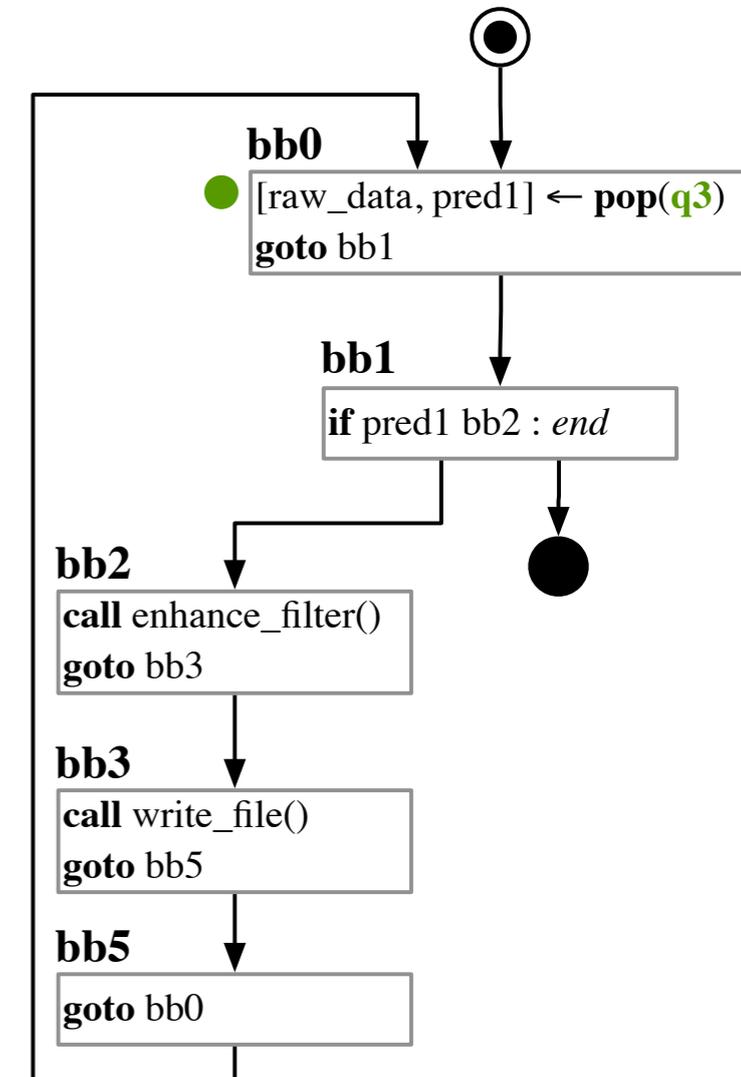
Stage 1



Stage 2



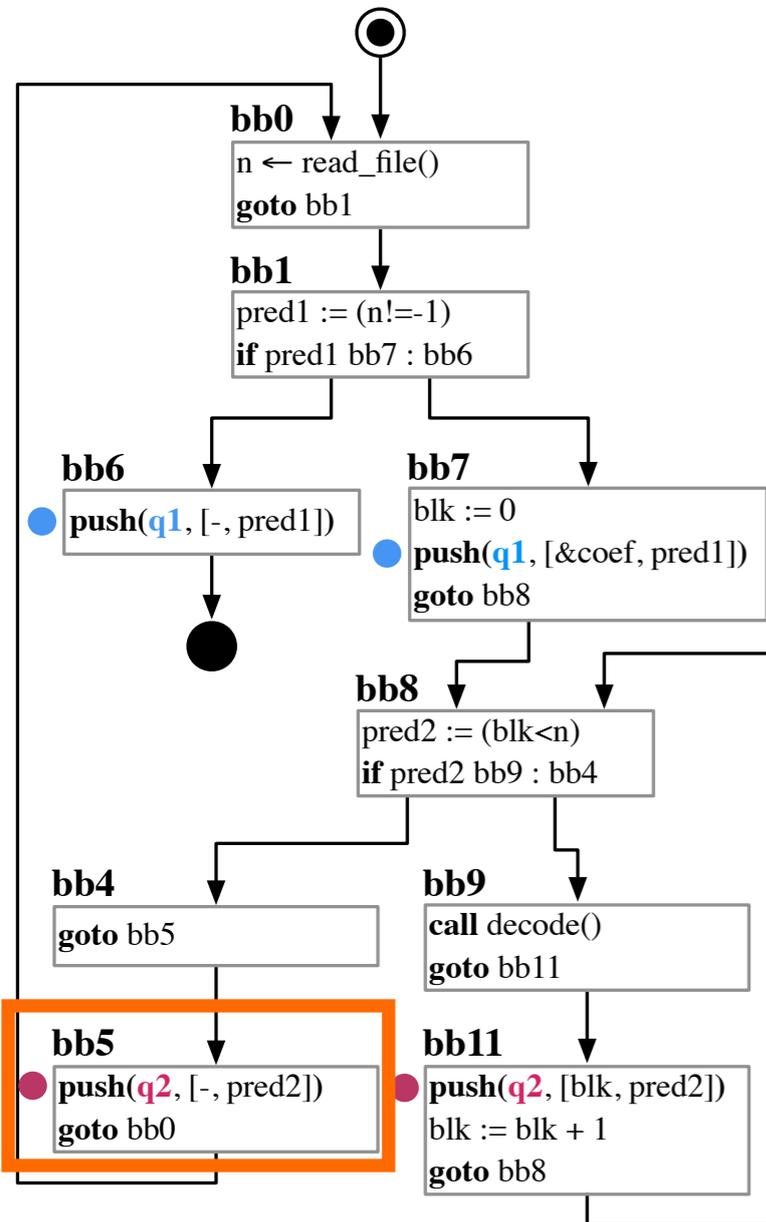
Stage 3



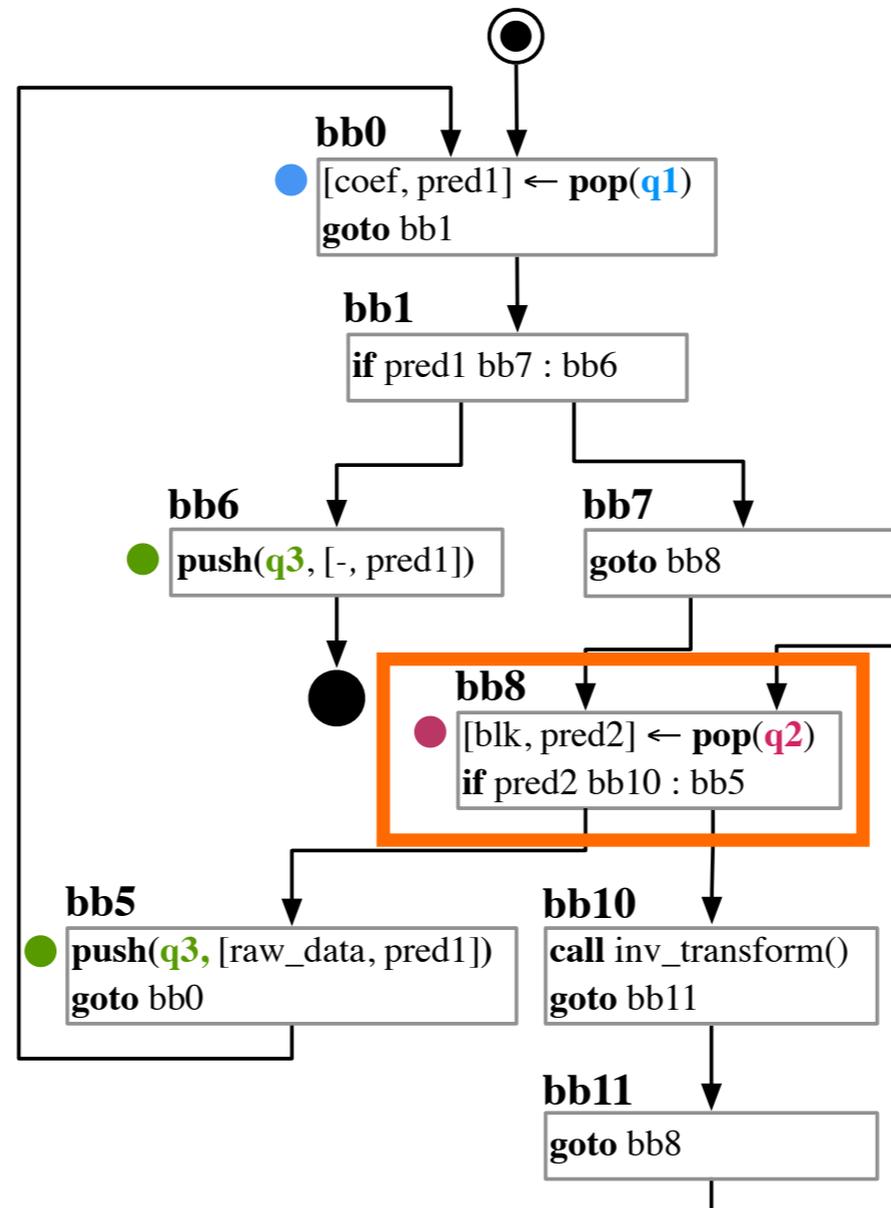
Communication



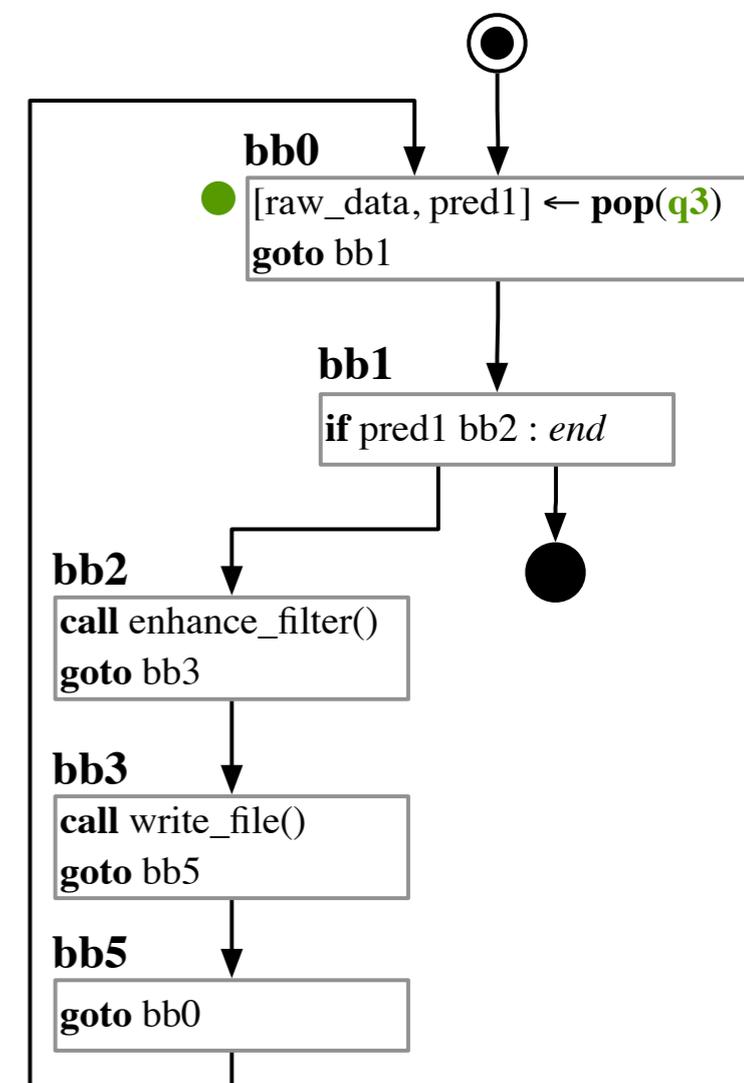
Stage 1



Stage 2



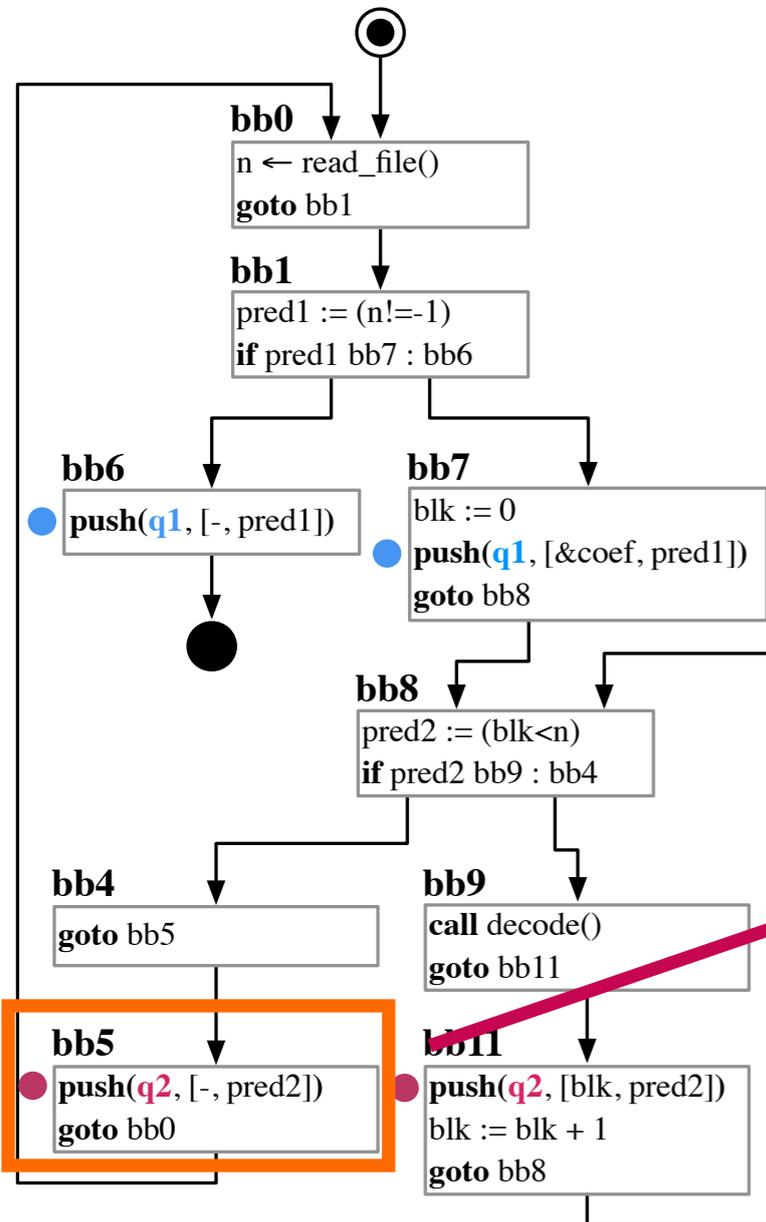
Stage 3



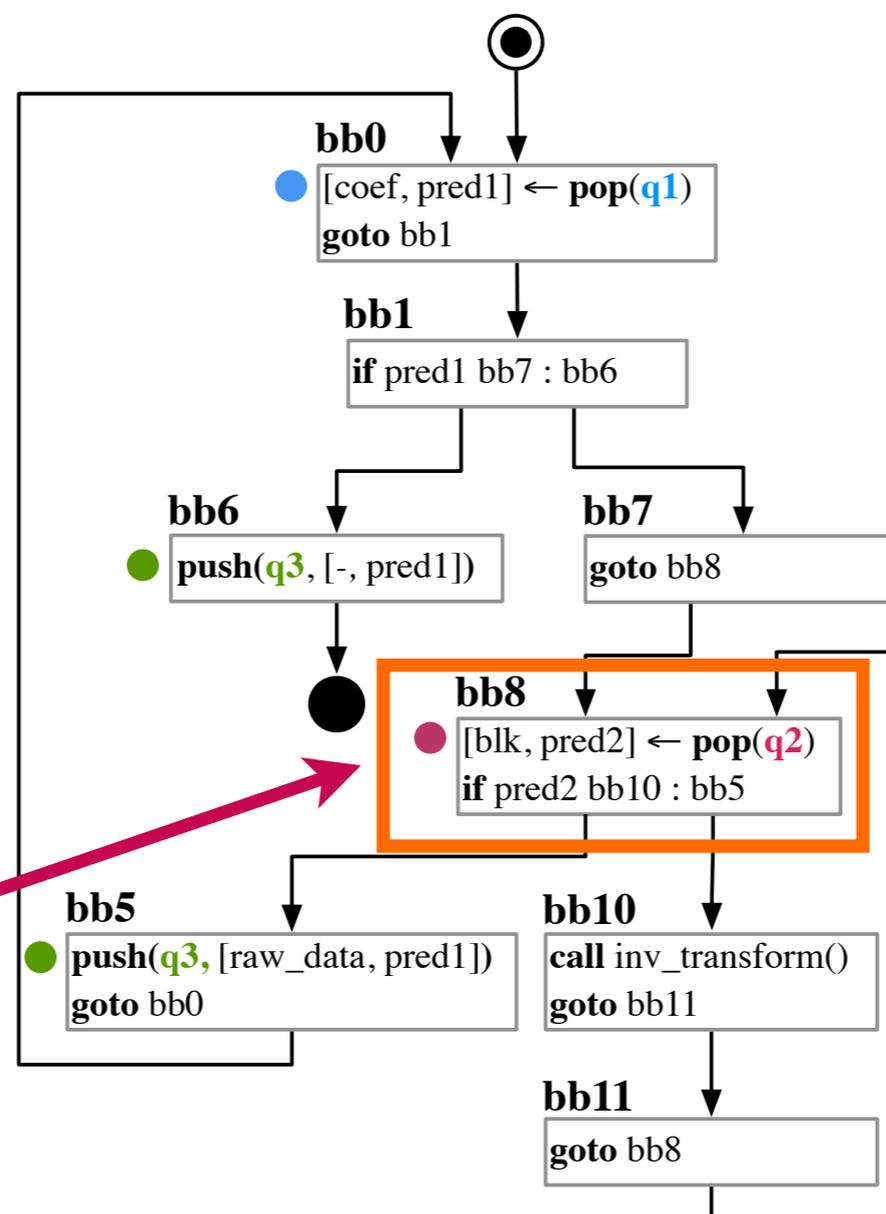
Communication



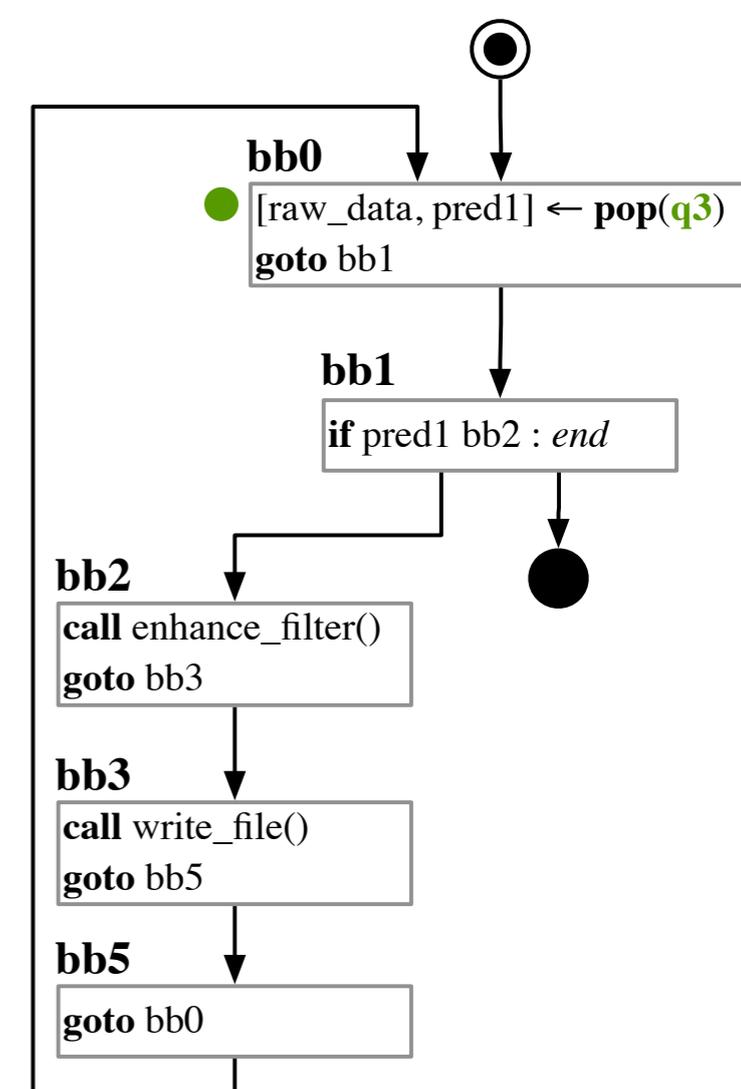
Stage 1



Stage 2



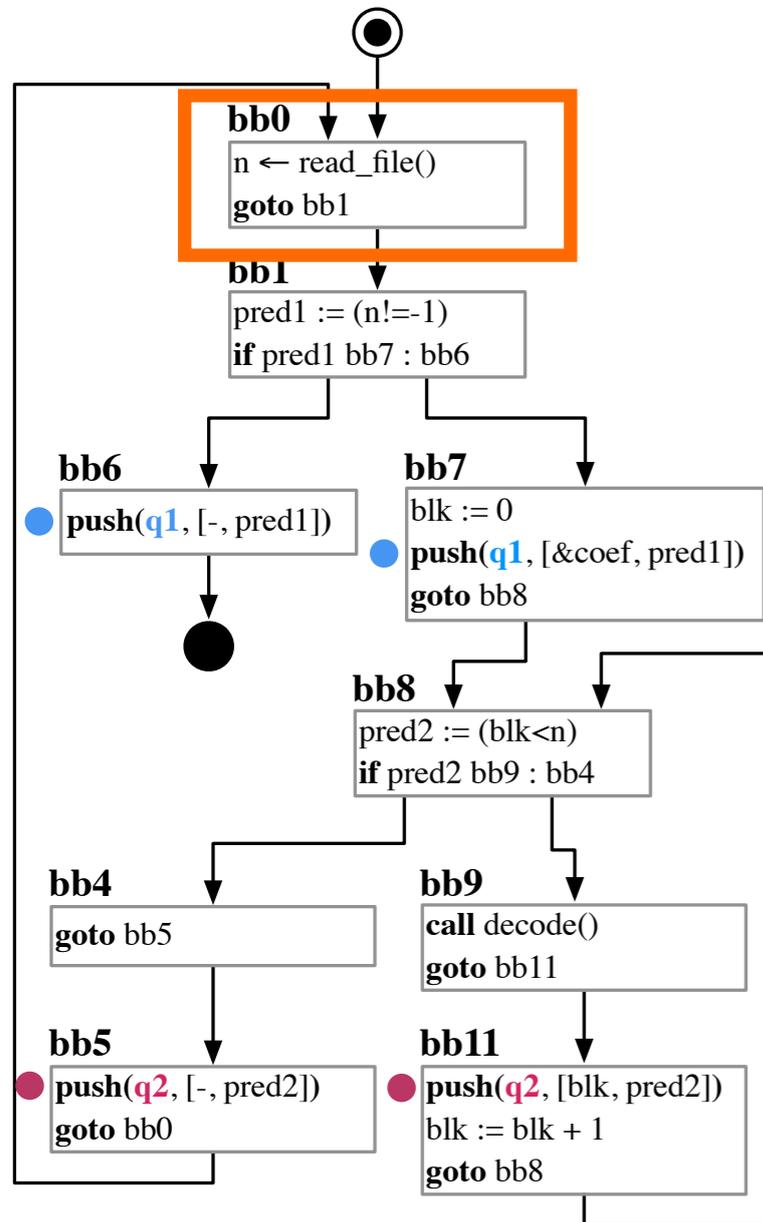
Stage 3



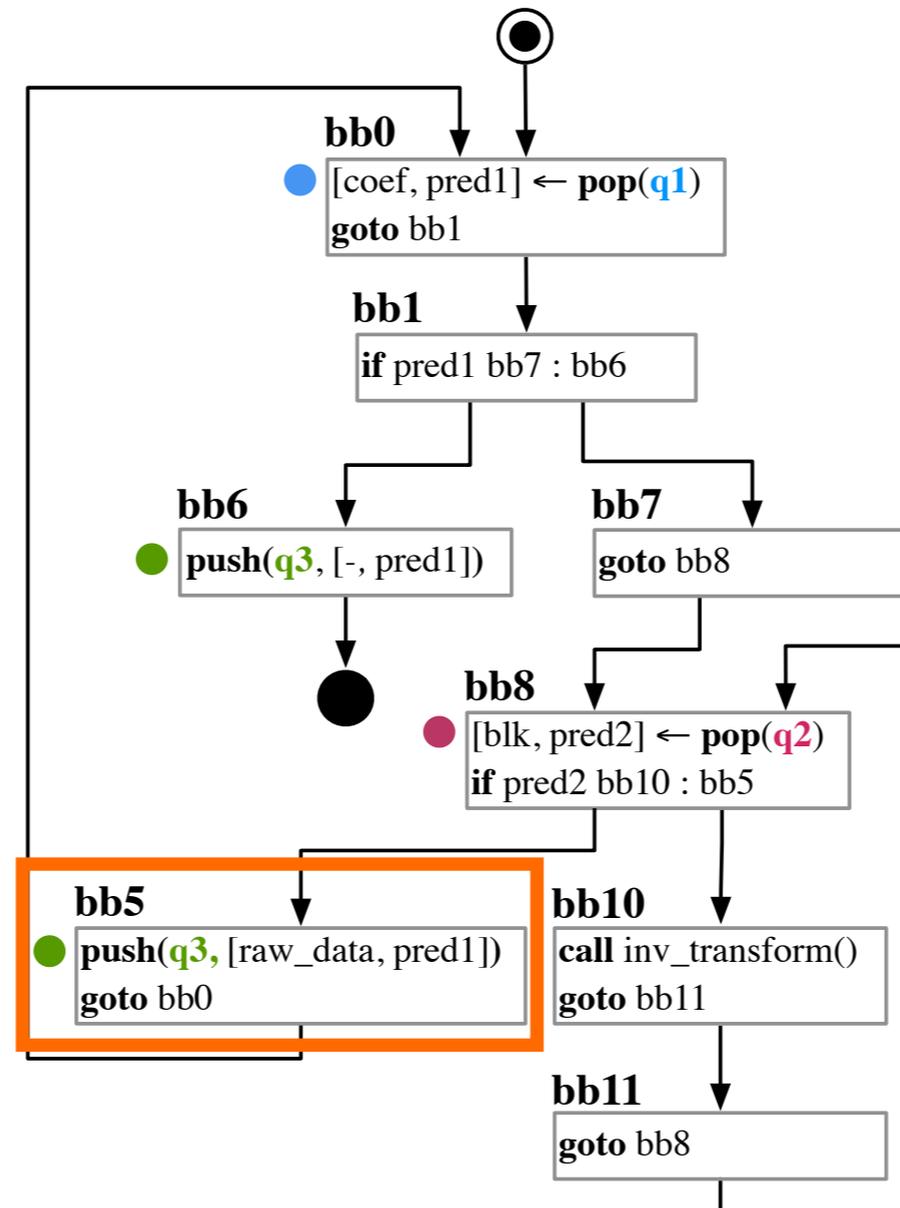
Communication



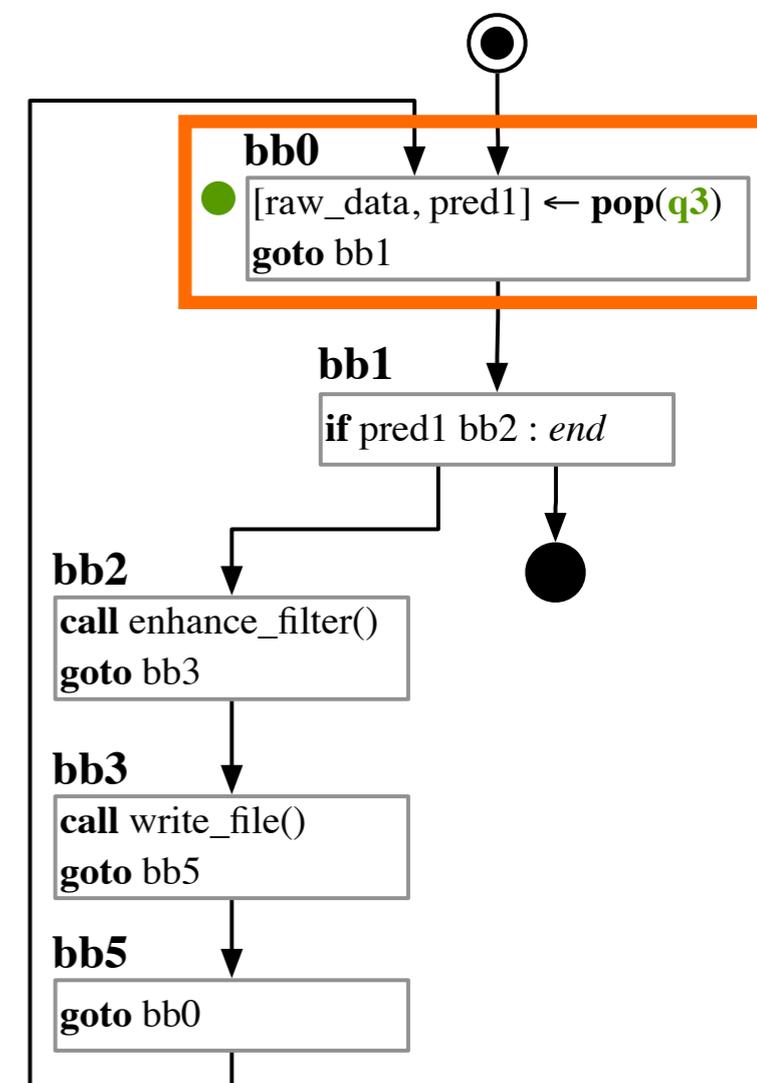
Stage 1



Stage 2



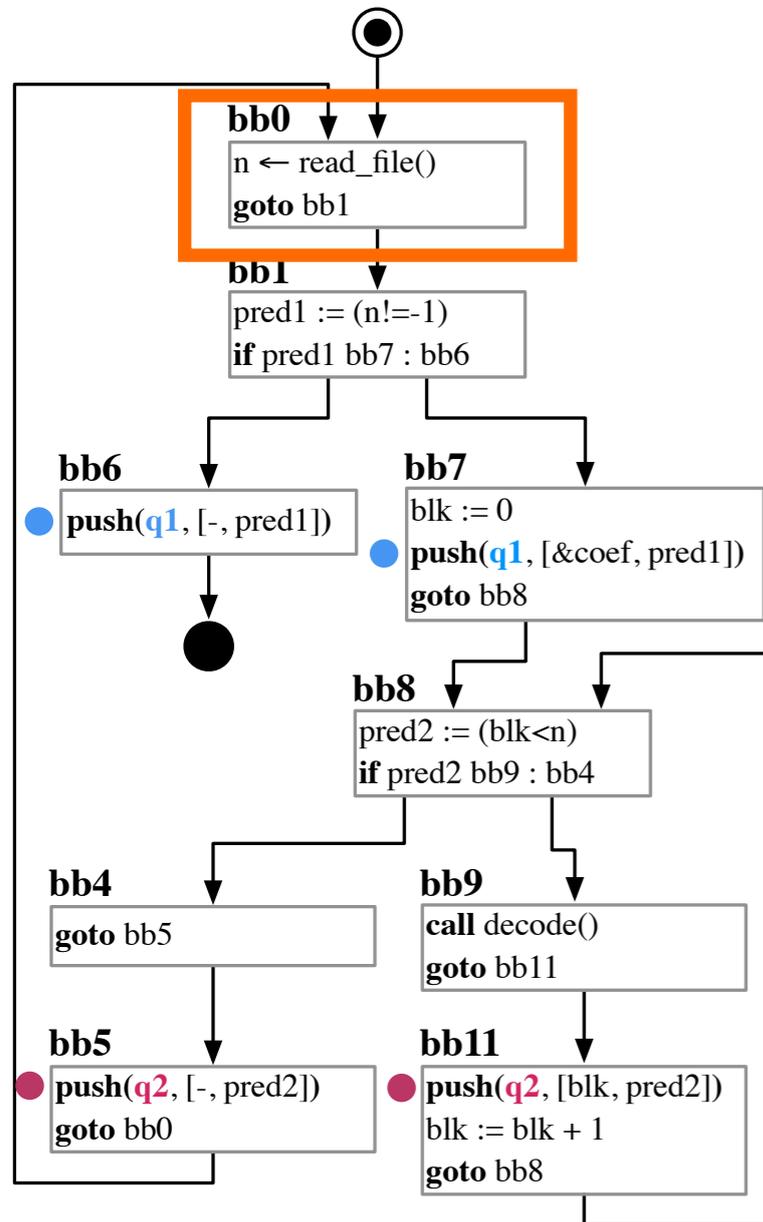
Stage 3



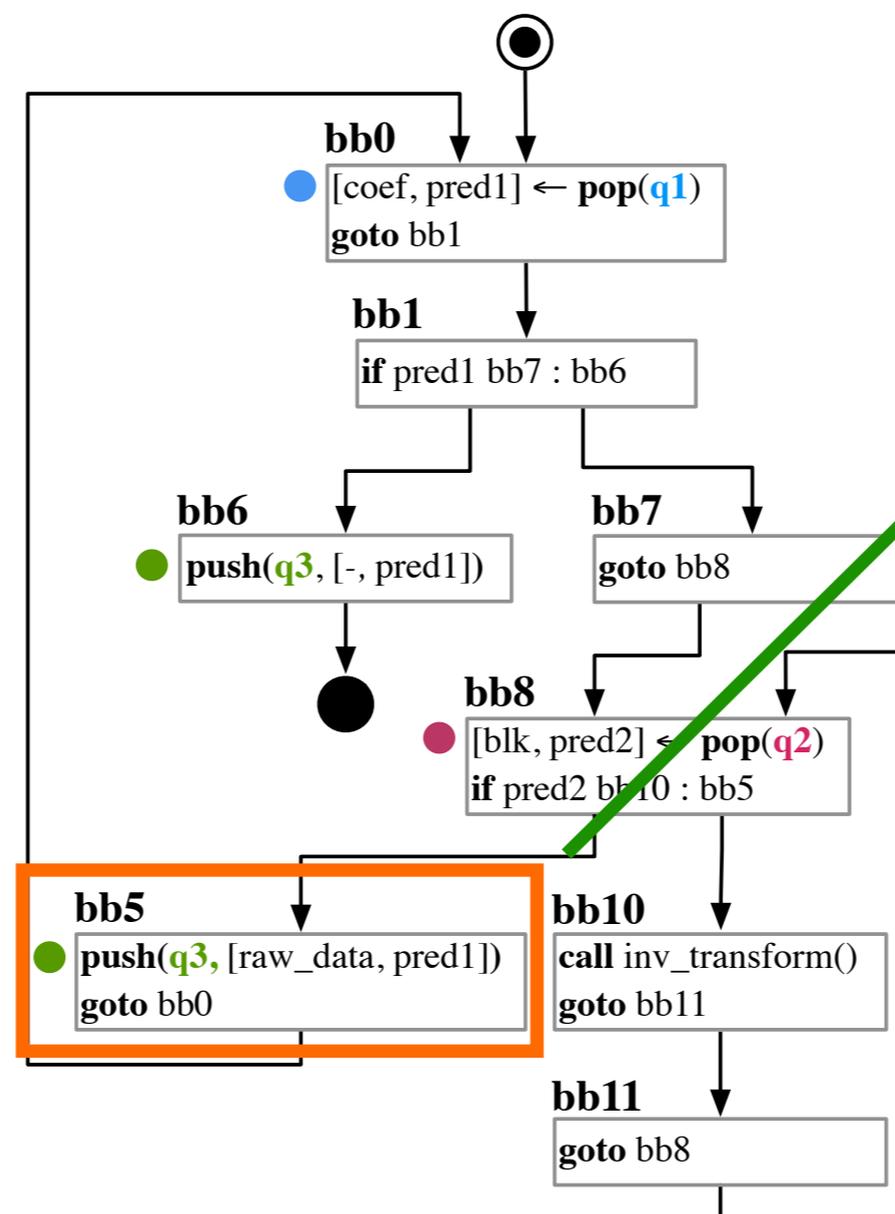
Communication



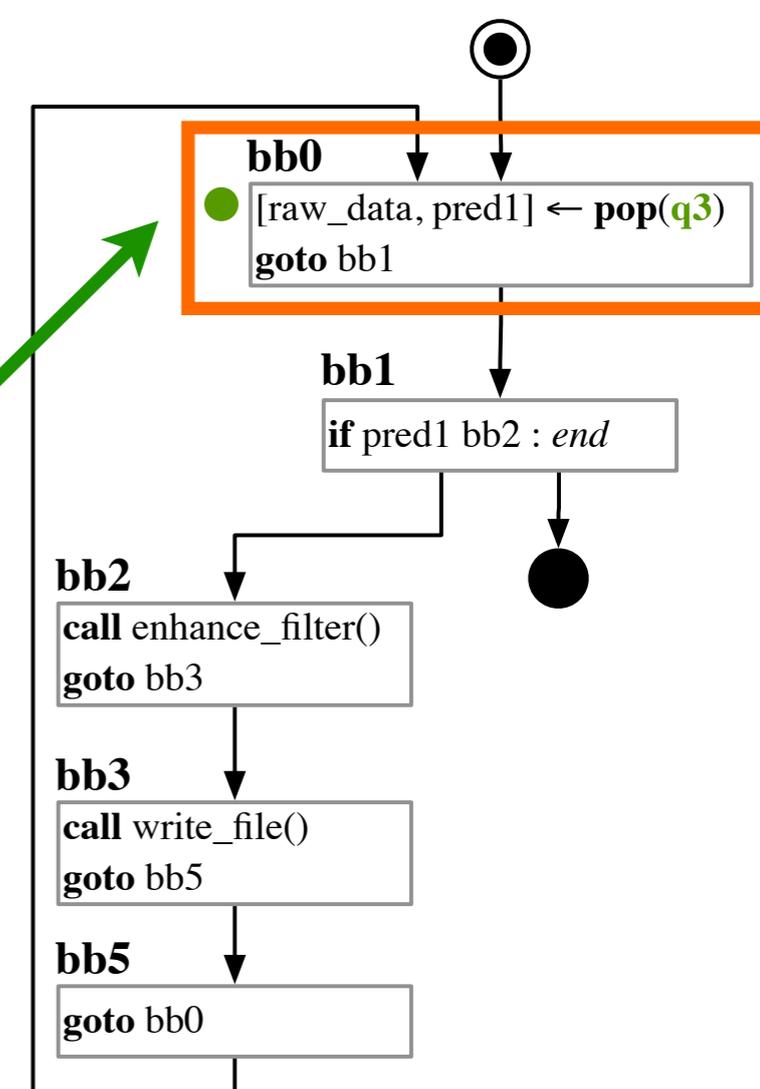
Stage 1



Stage 2



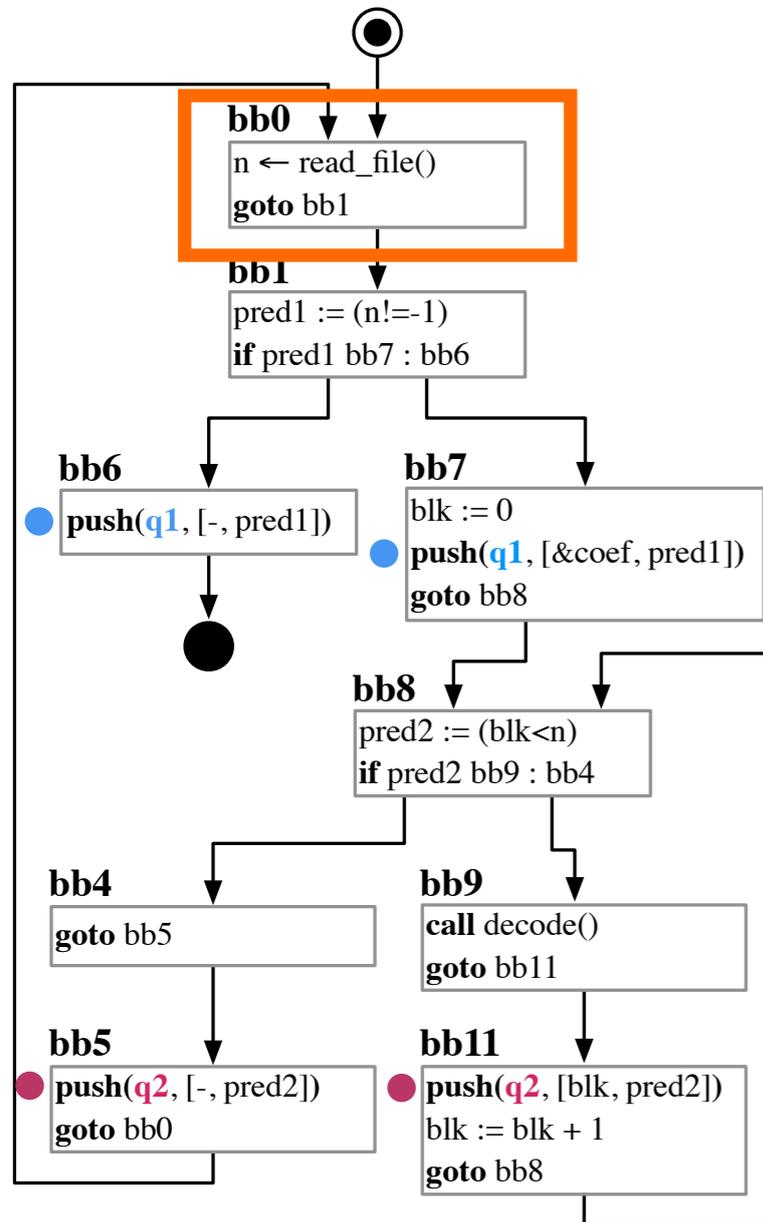
Stage 3



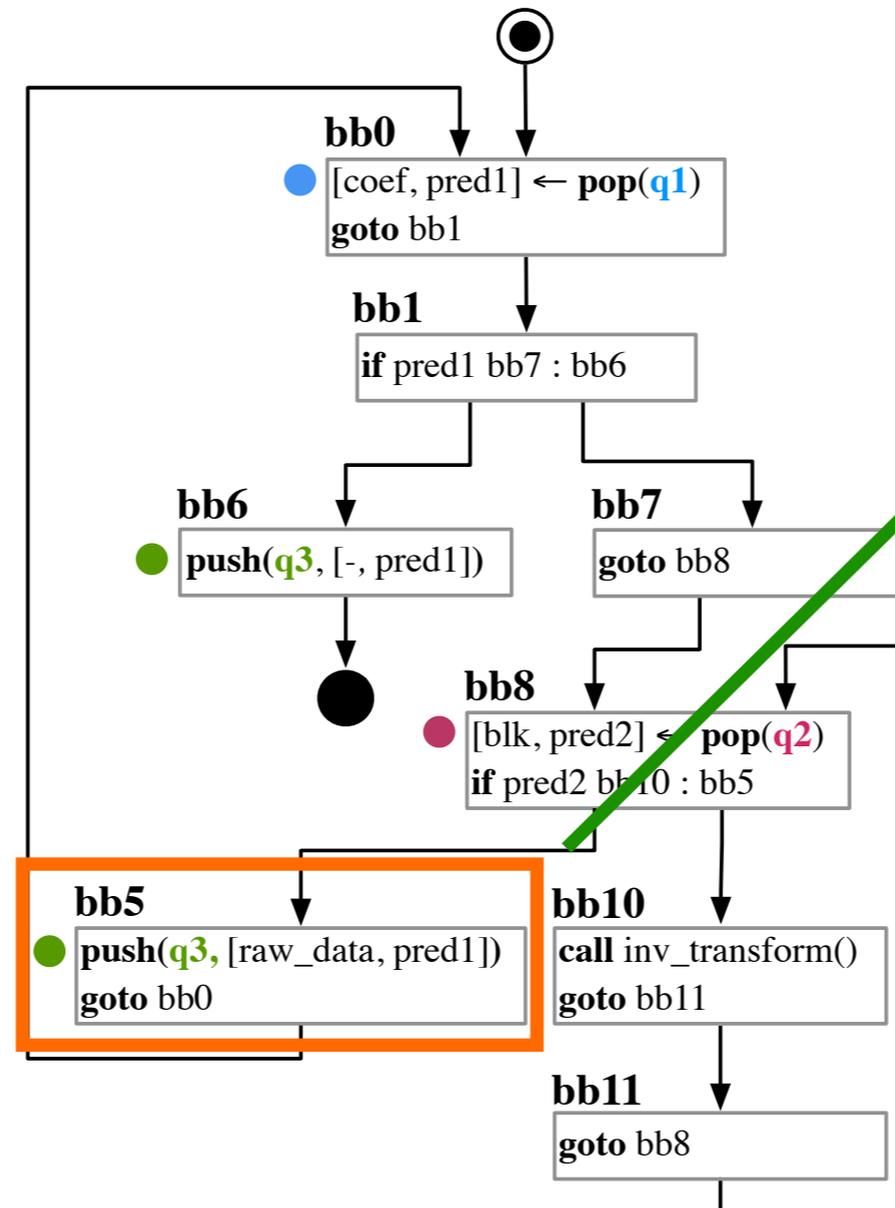
Communication



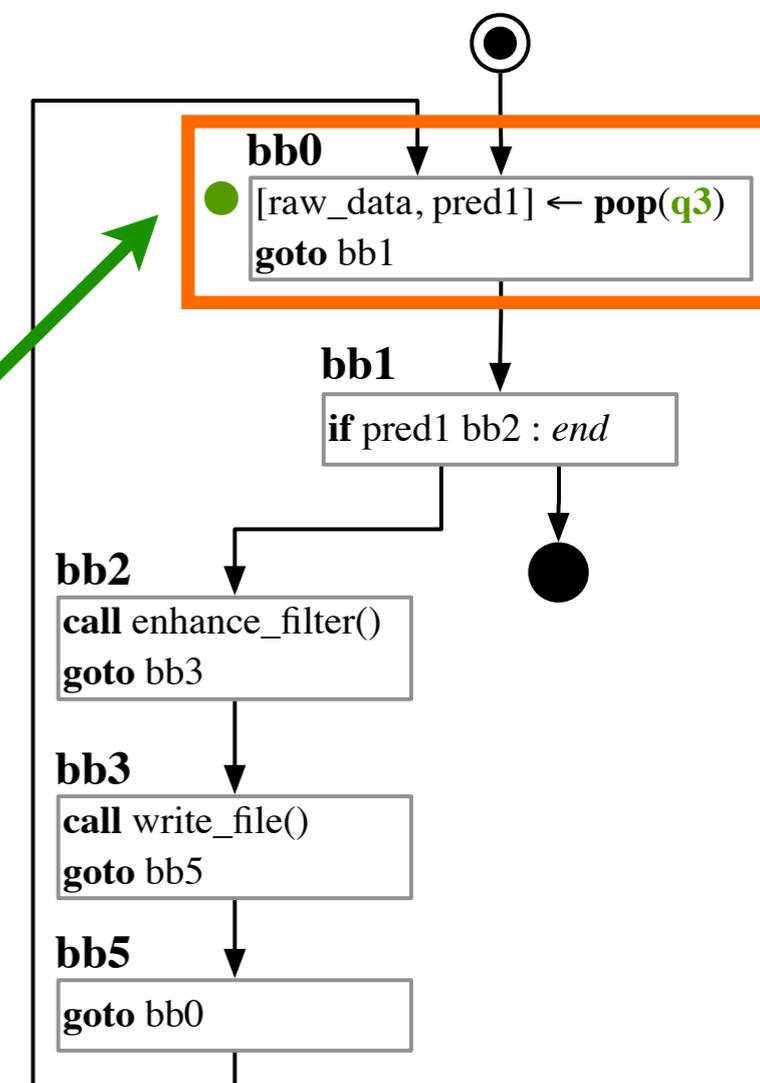
Stage 1



Stage 2



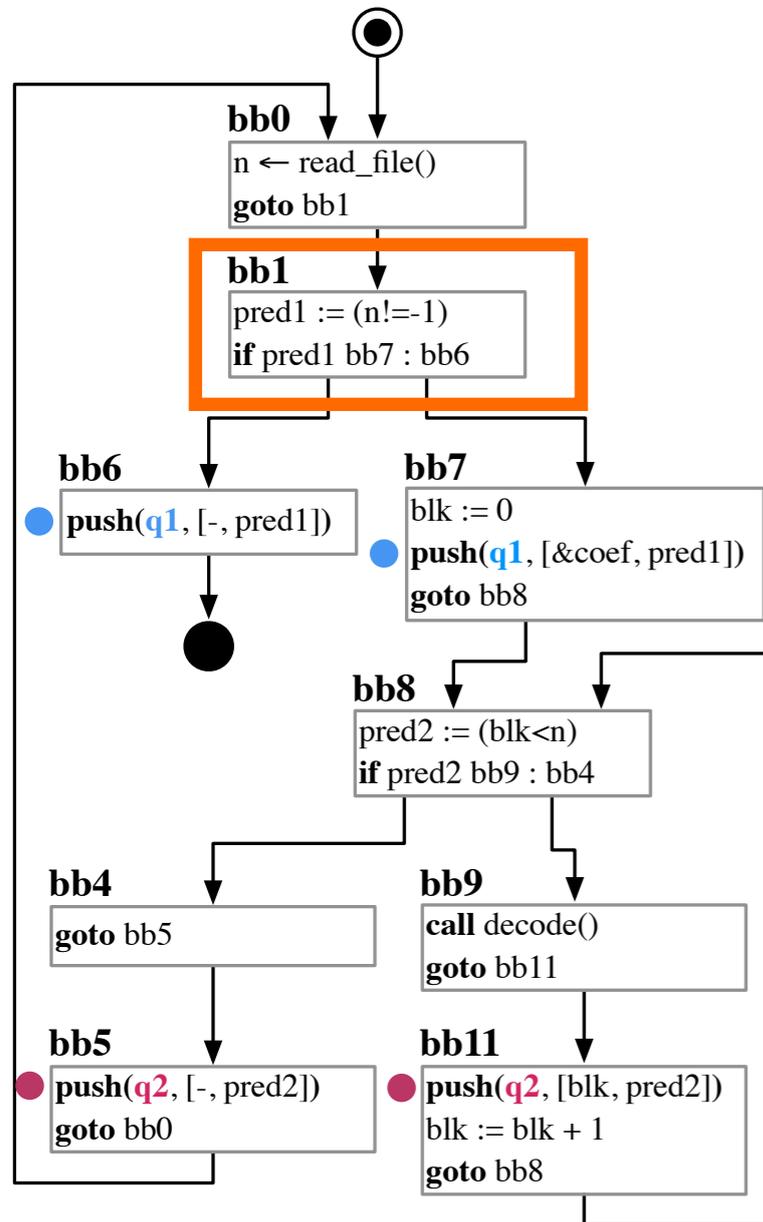
Stage 3



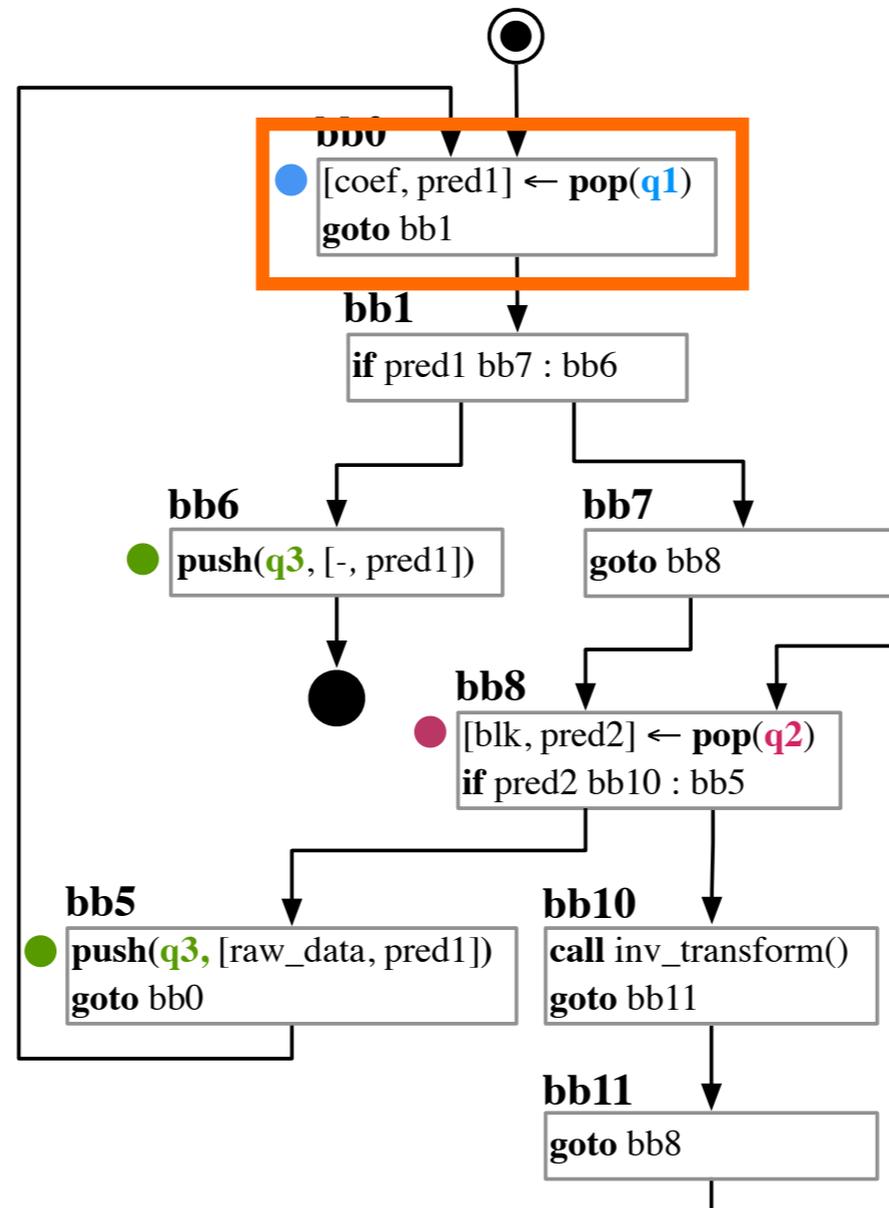
Communication



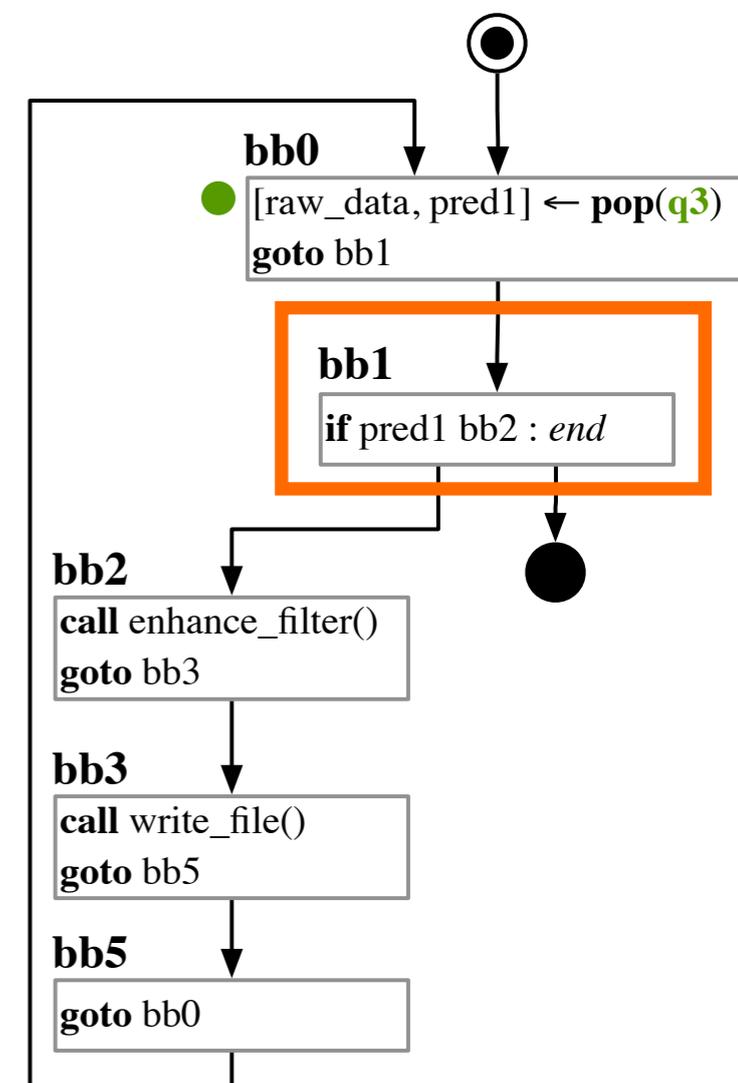
Stage 1



Stage 2



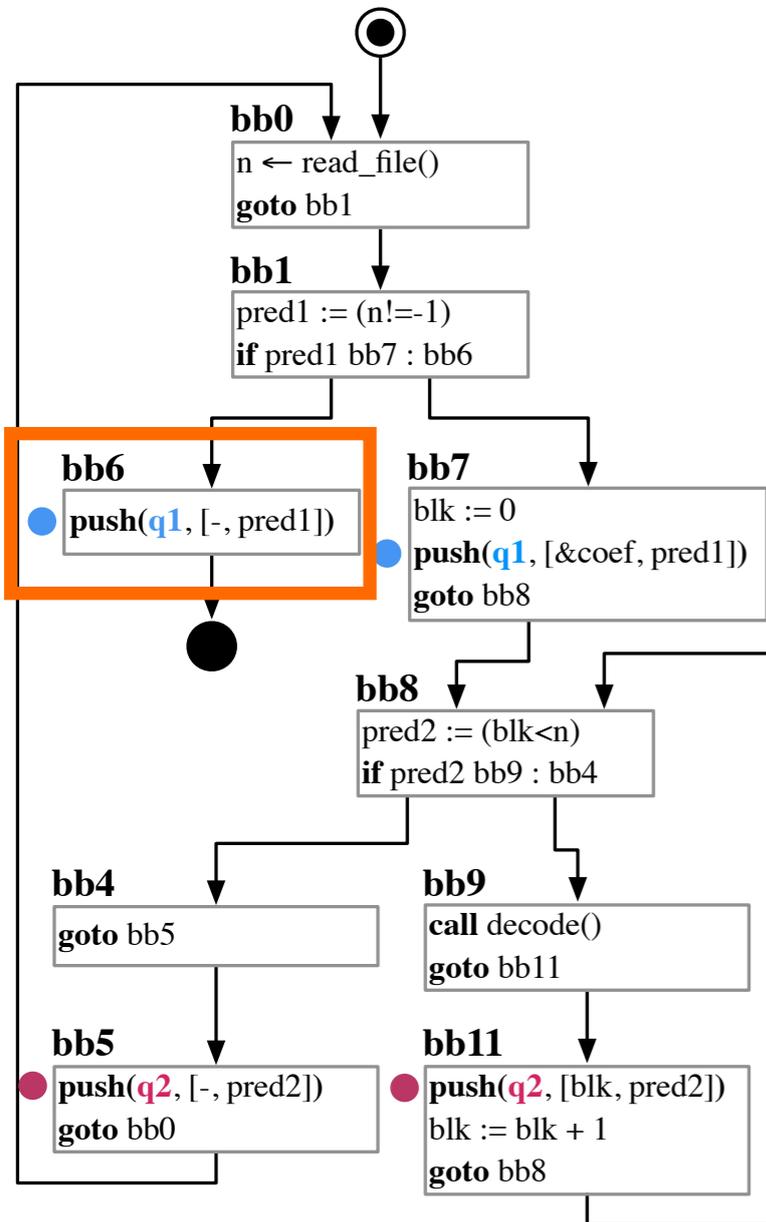
Stage 3



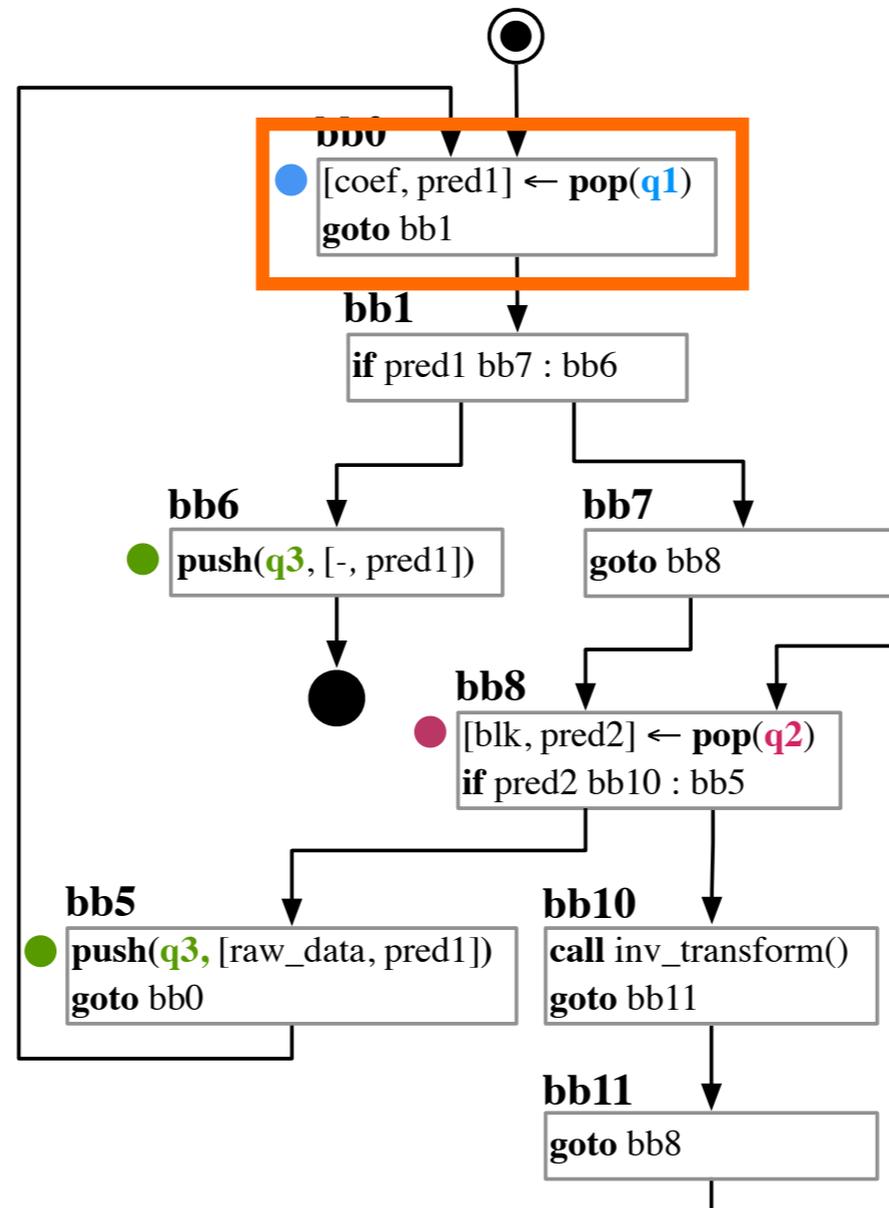
Communication



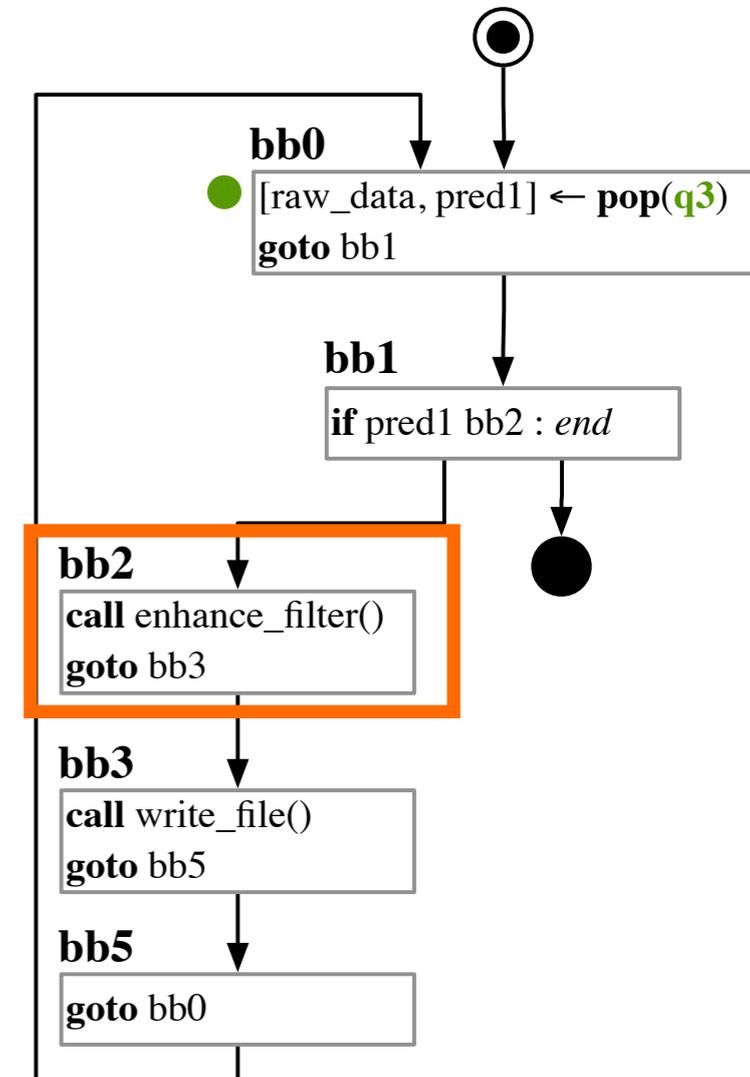
Stage 1



Stage 2



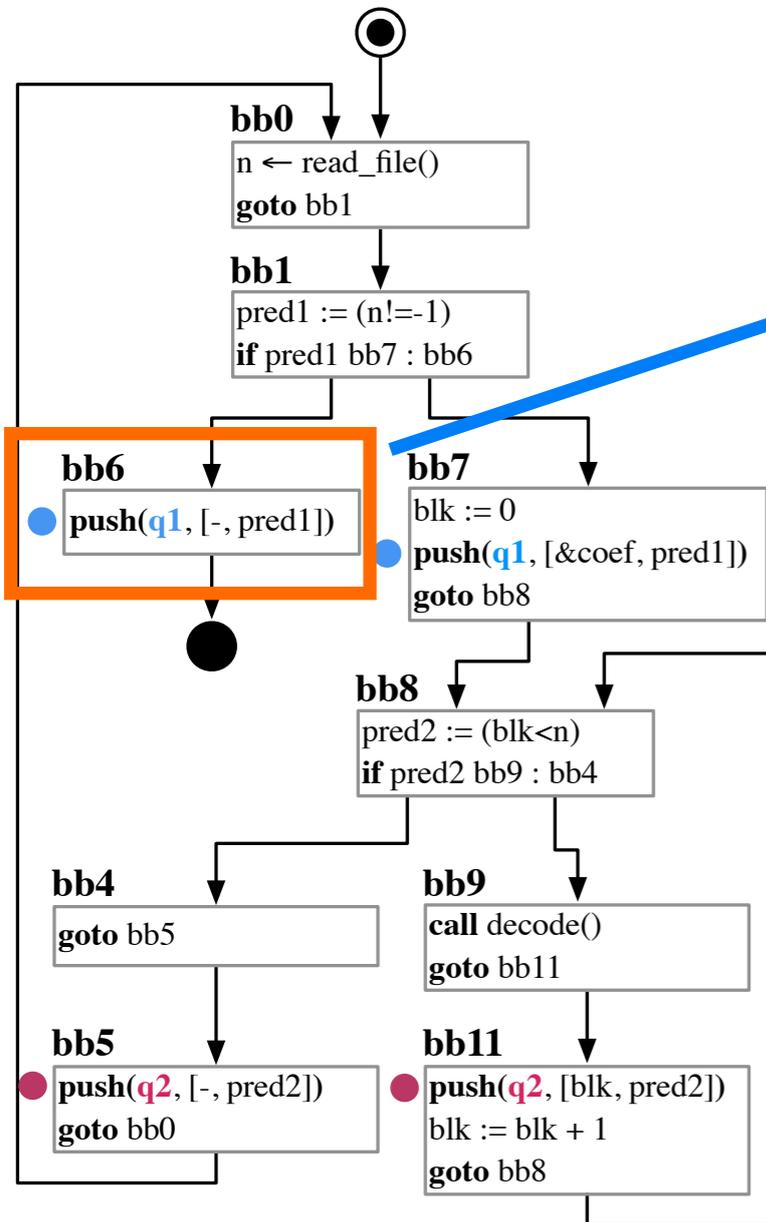
Stage 3



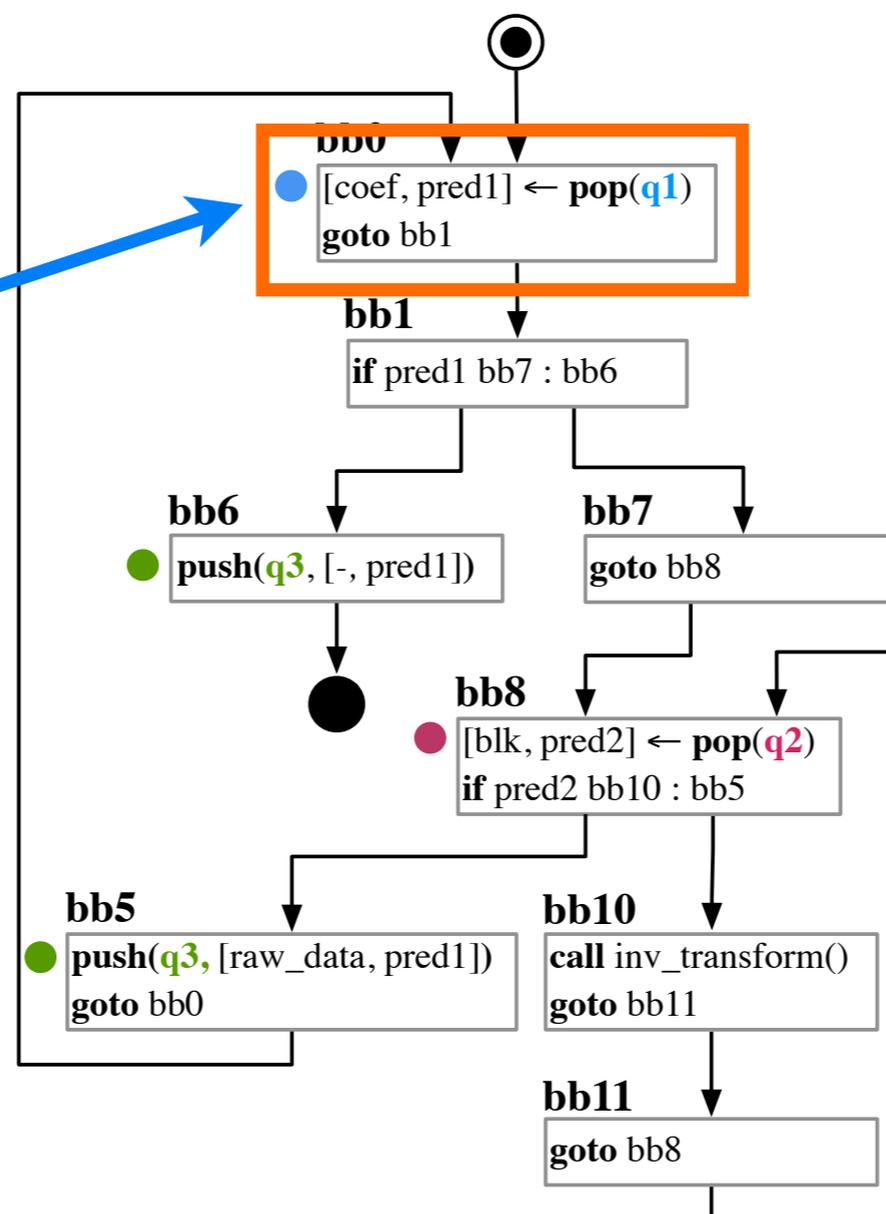
Communication



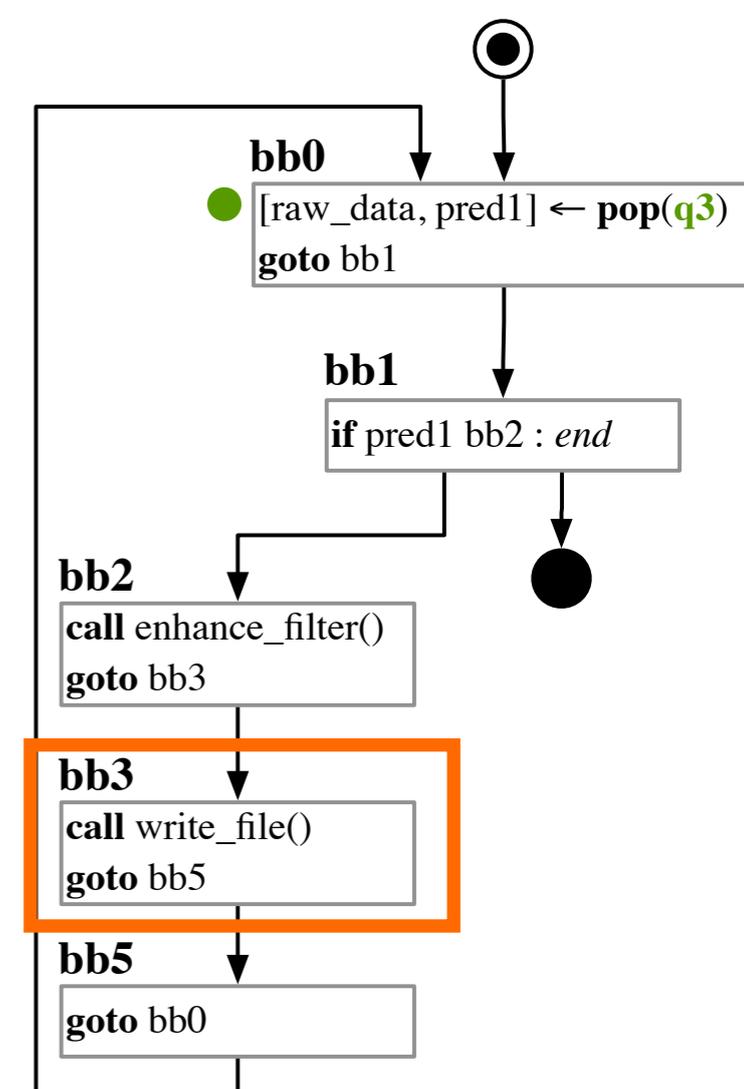
Stage 1



Stage 2



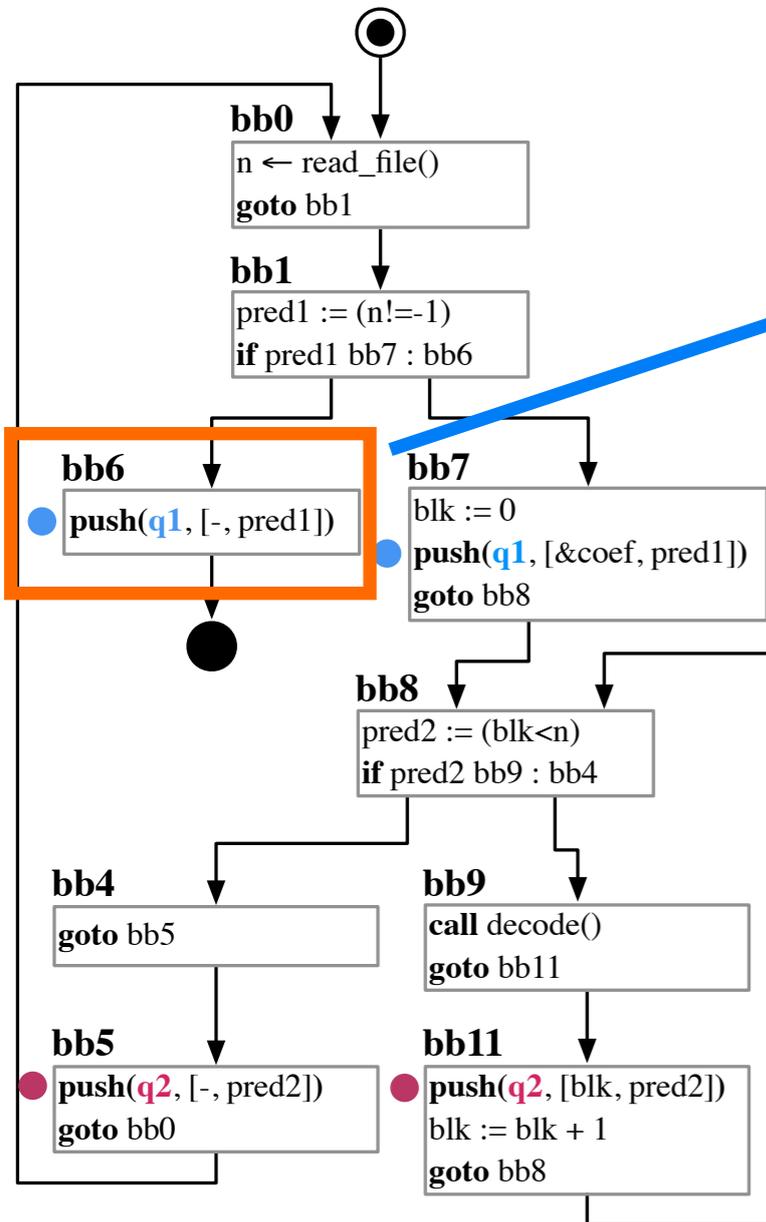
Stage 3



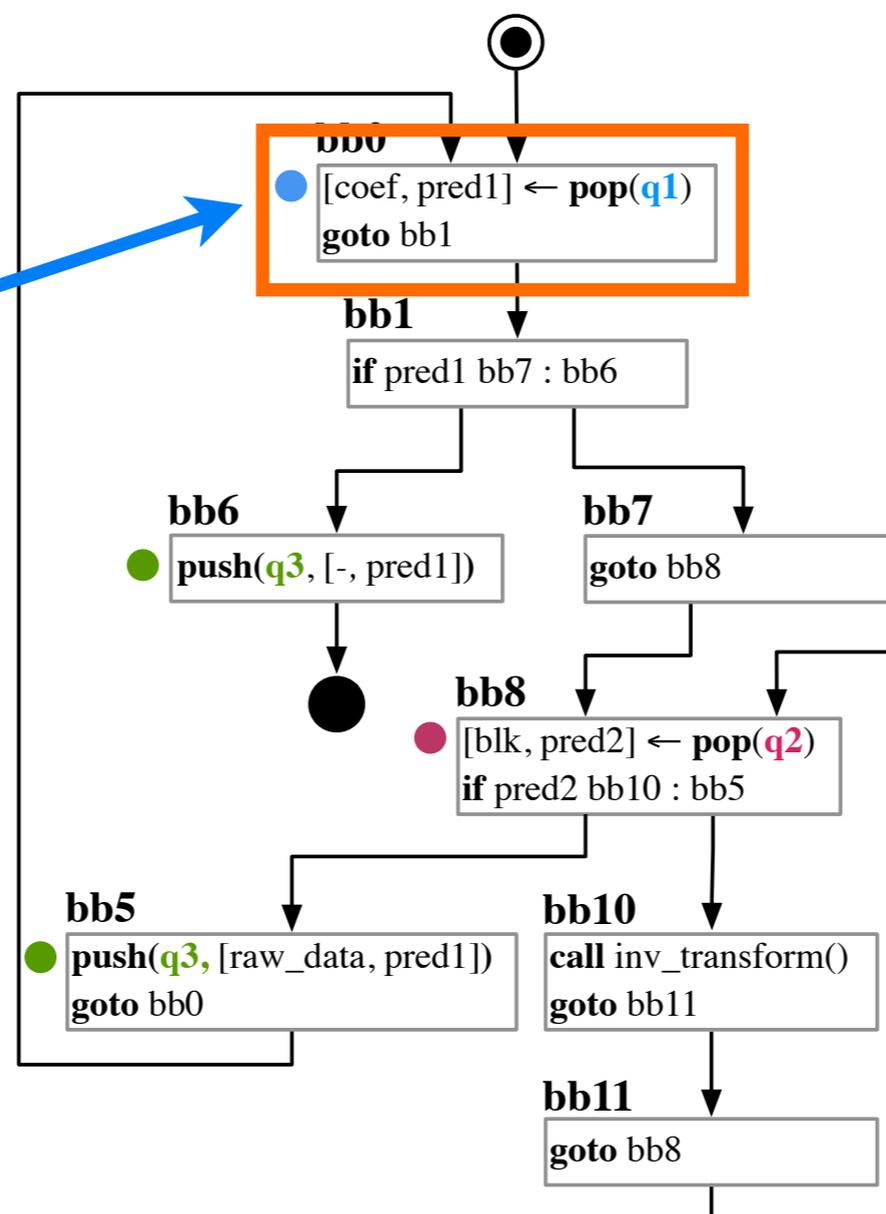
Communication



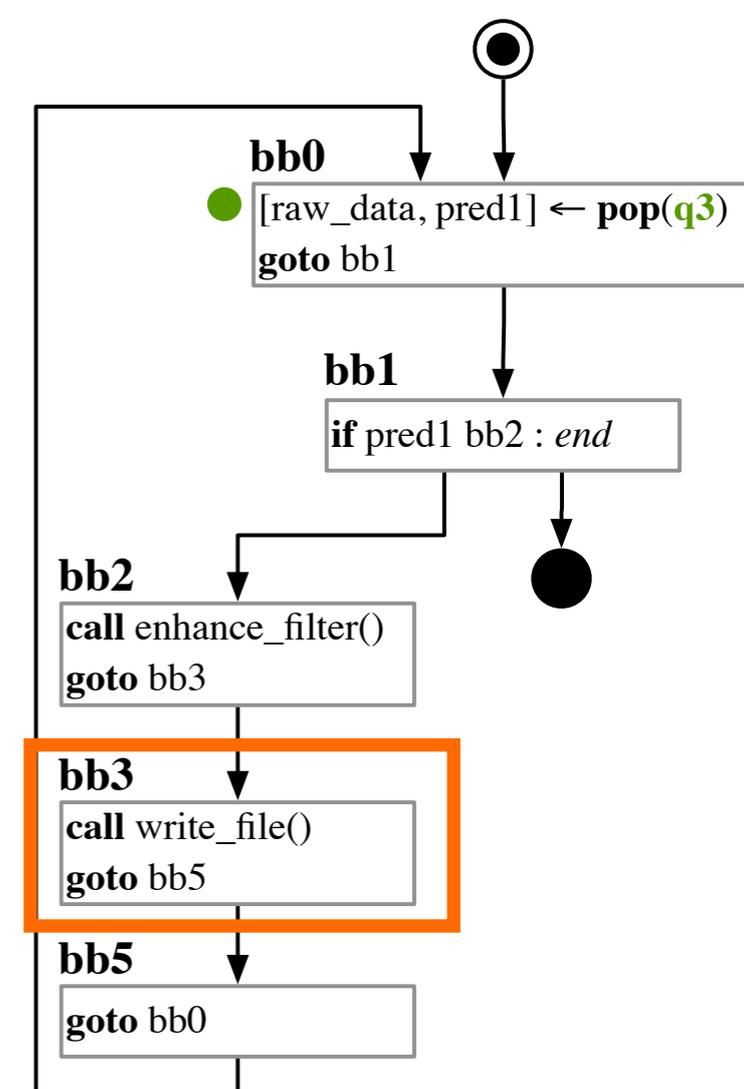
Stage 1



Stage 2



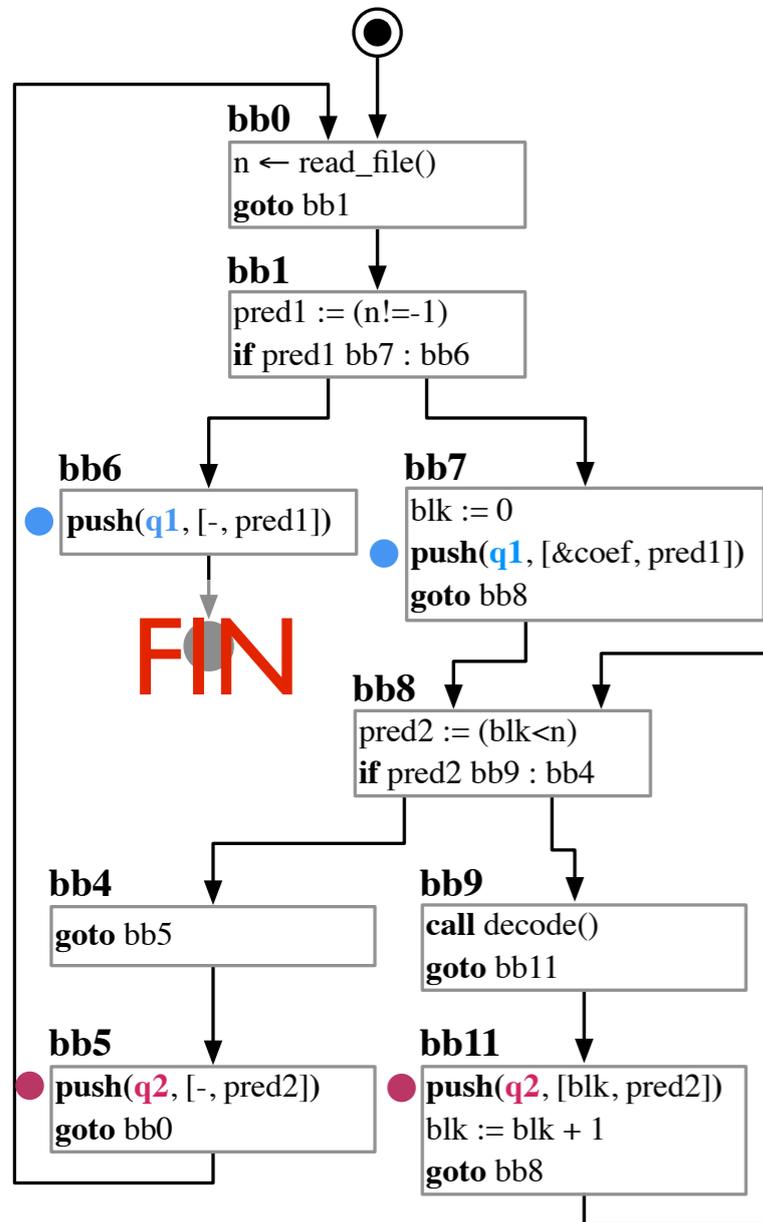
Stage 3



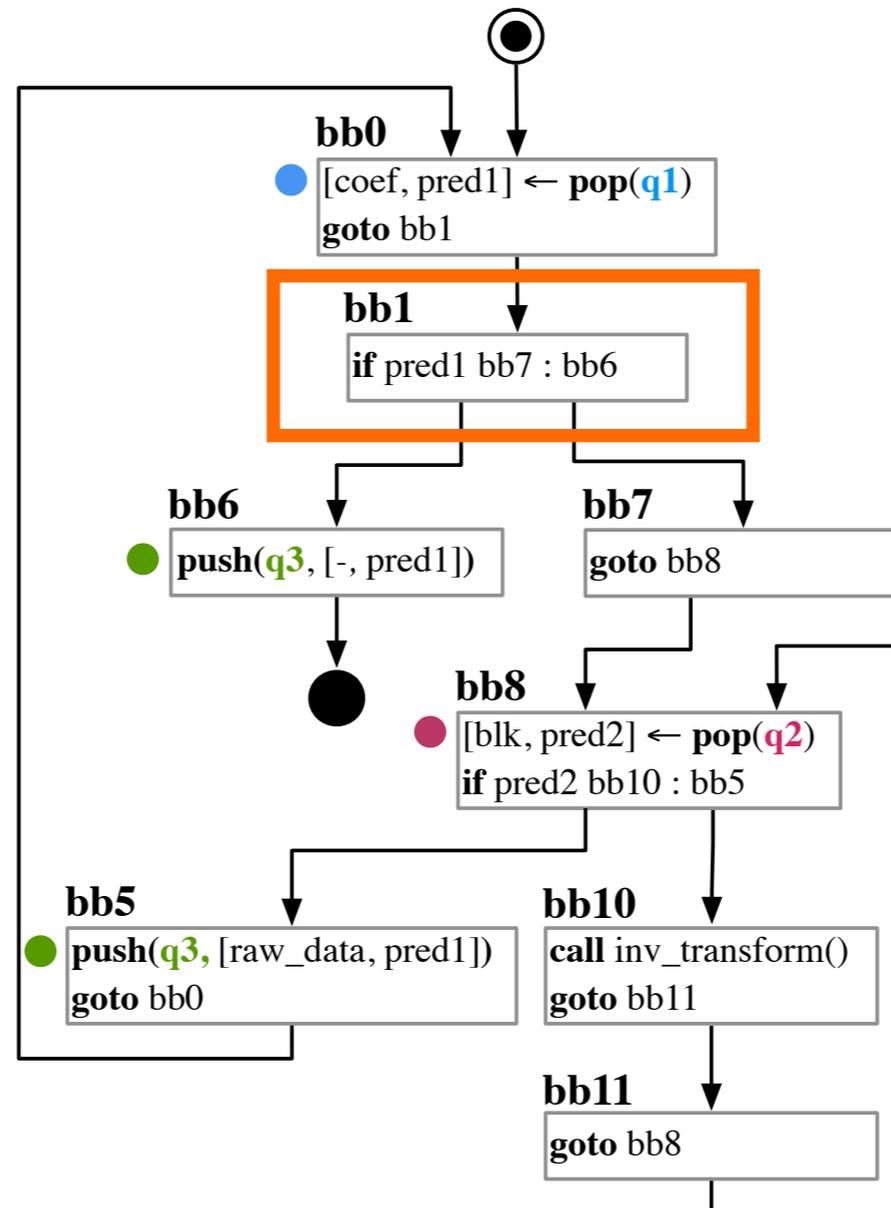
Communication



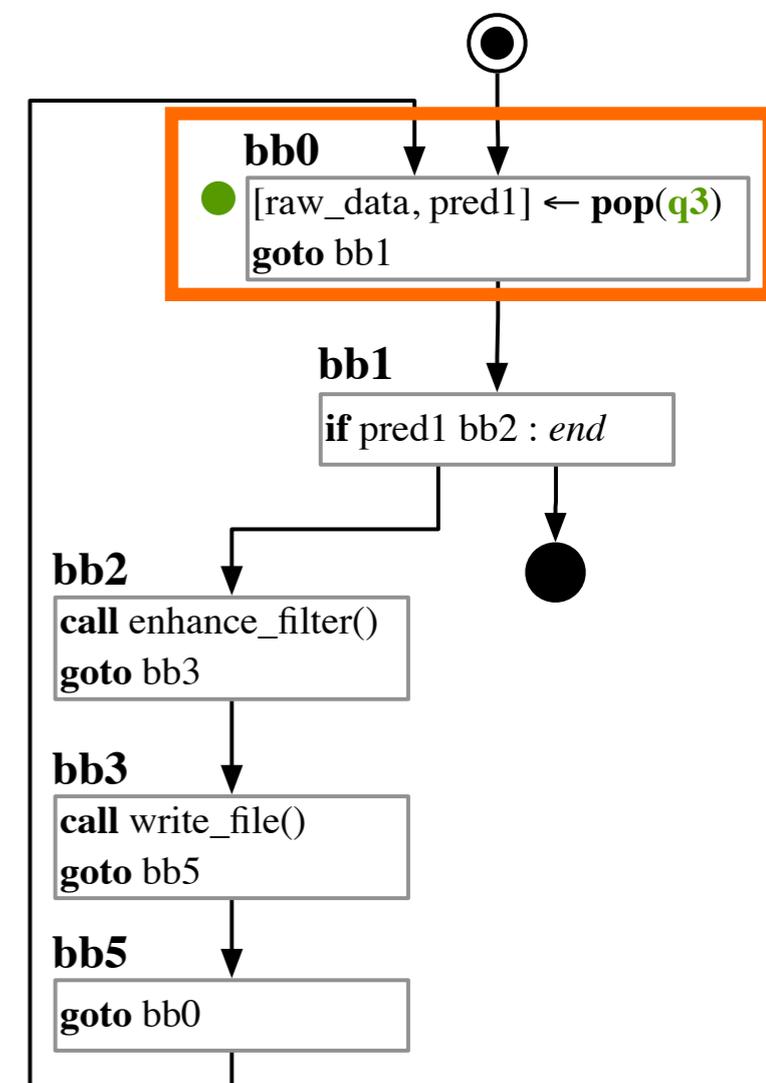
Stage 1



Stage 2



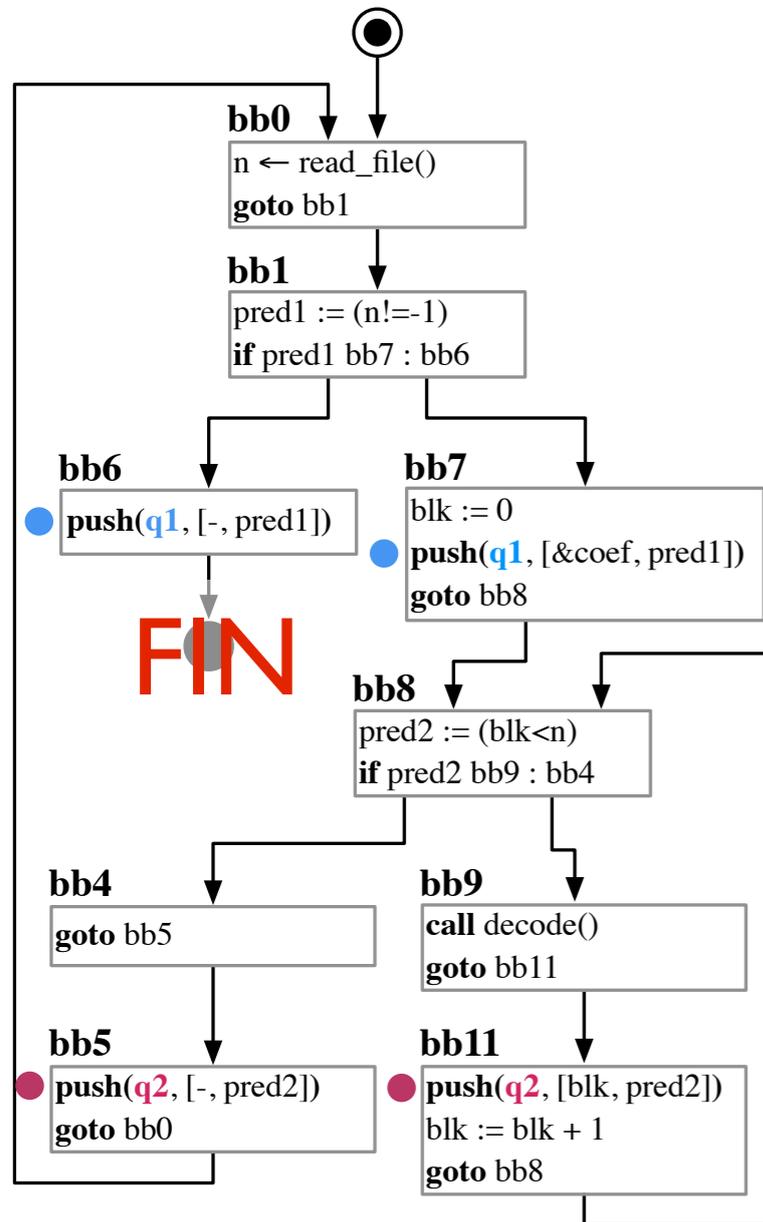
Stage 3



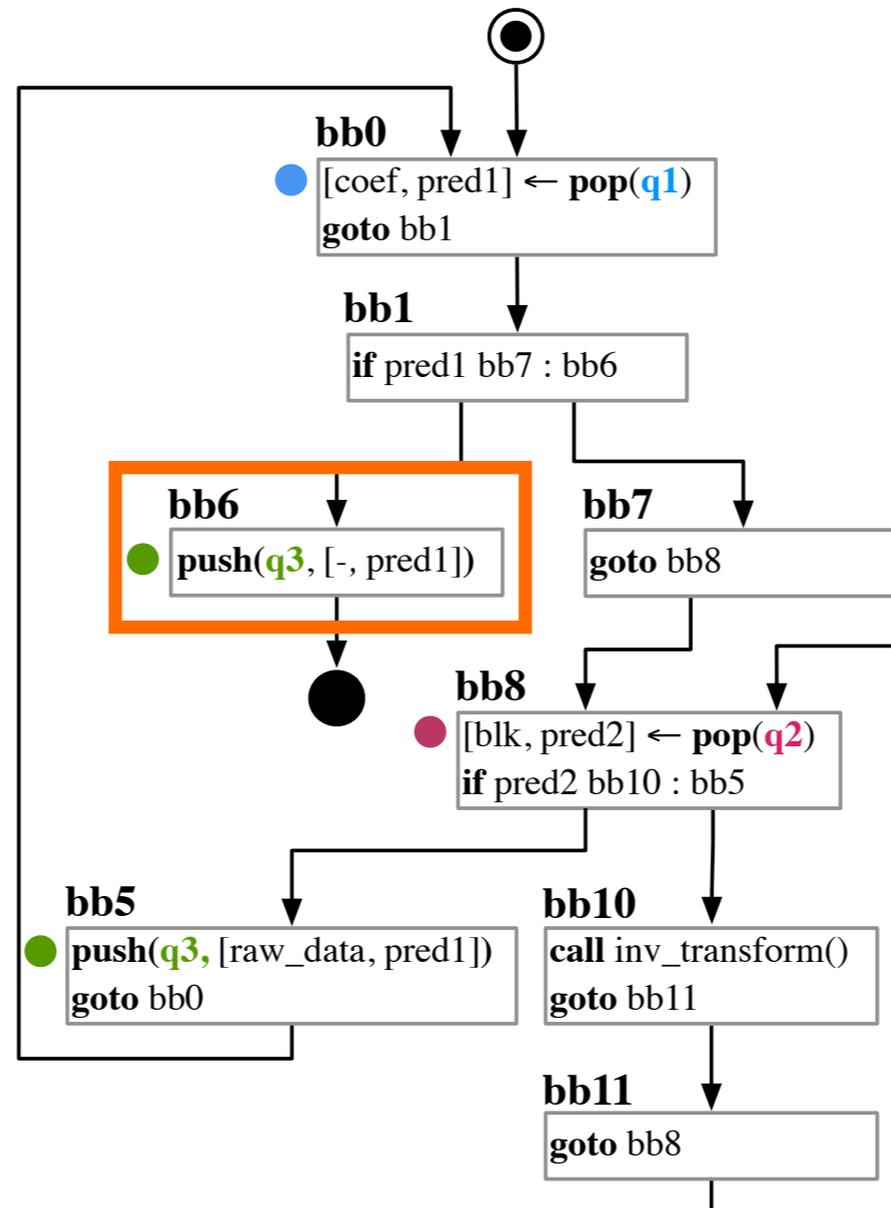
Communication



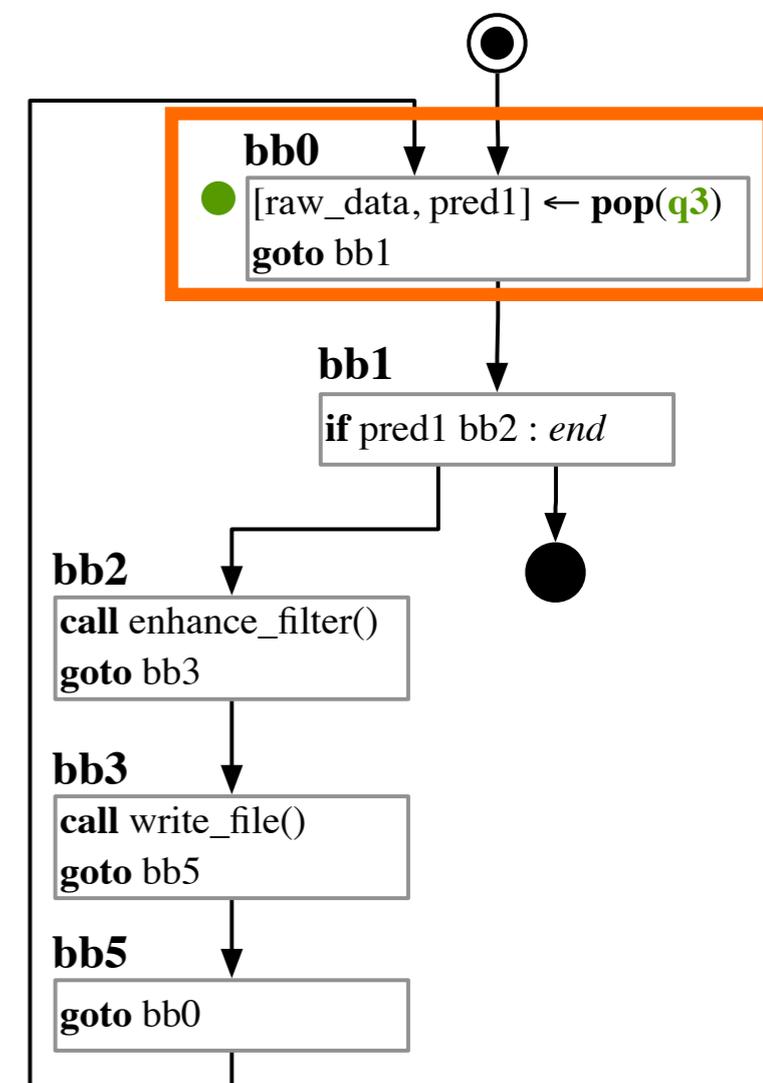
Stage 1



Stage 2



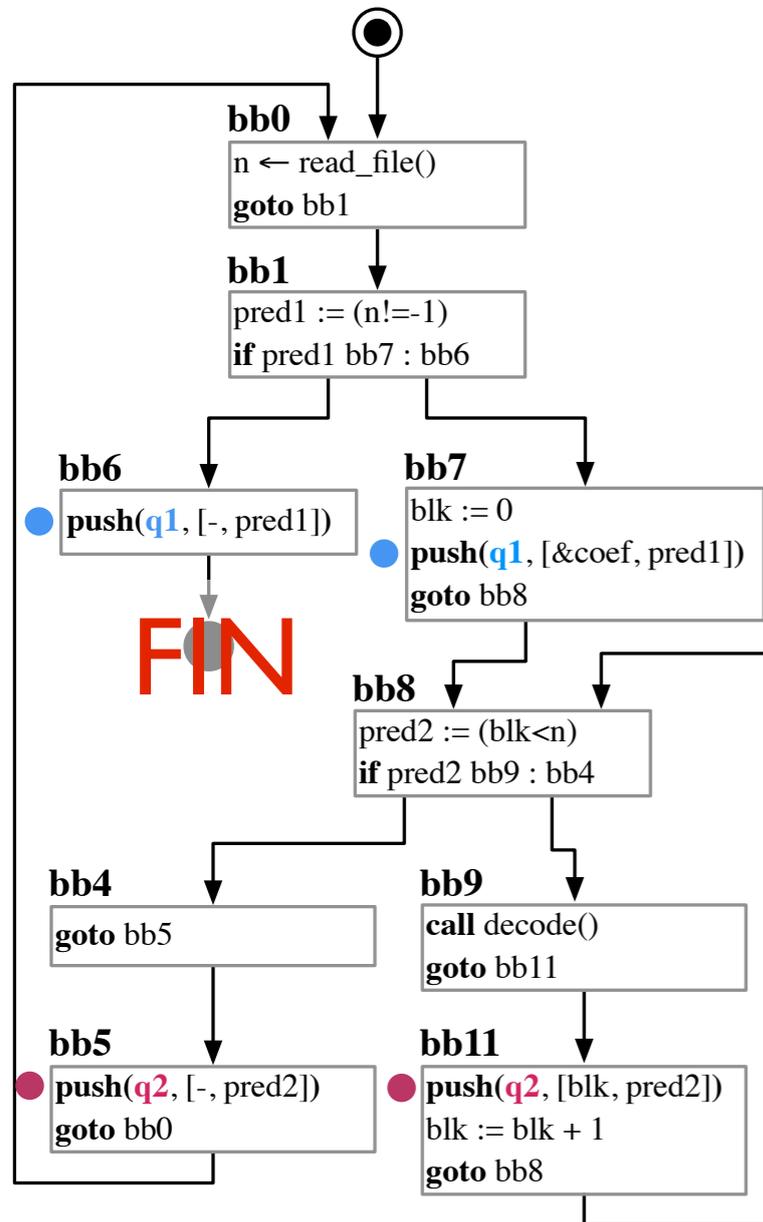
Stage 3



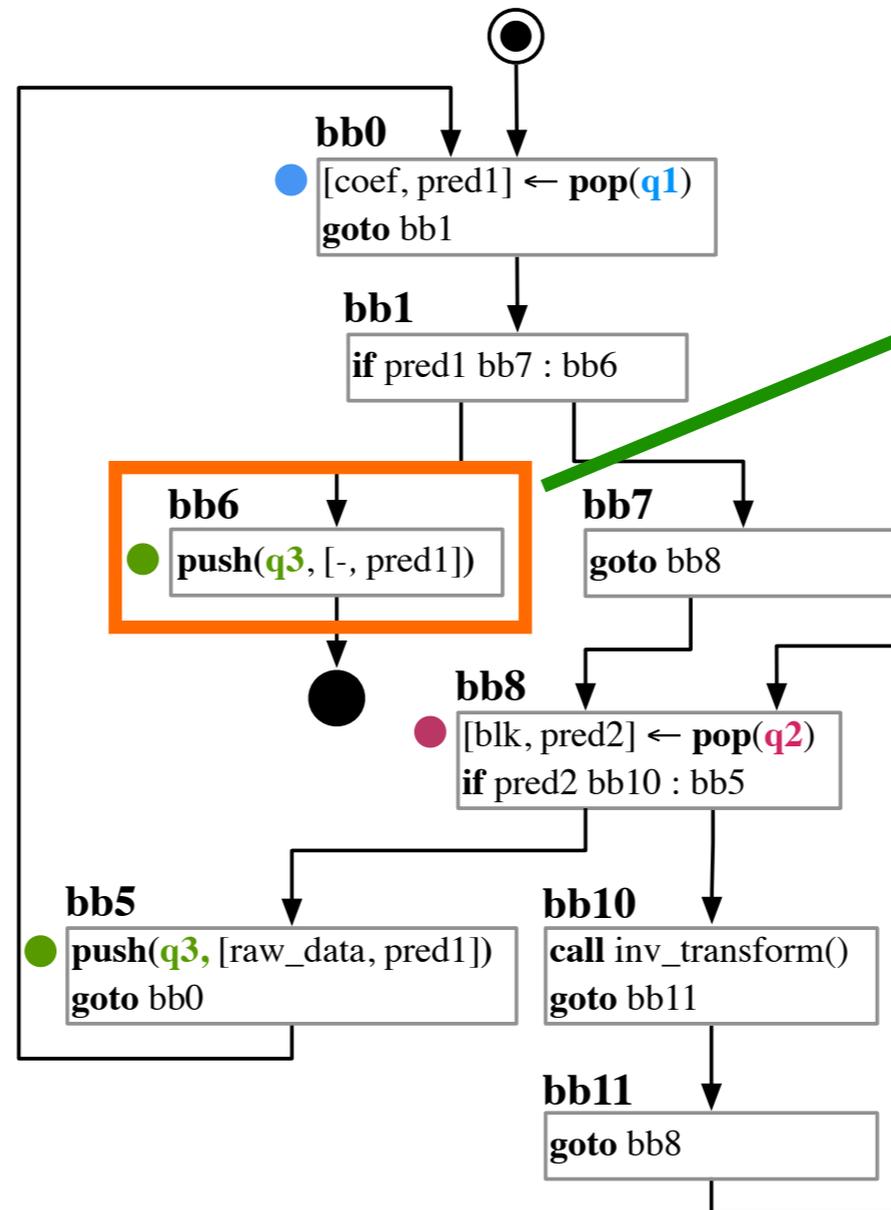
Communication



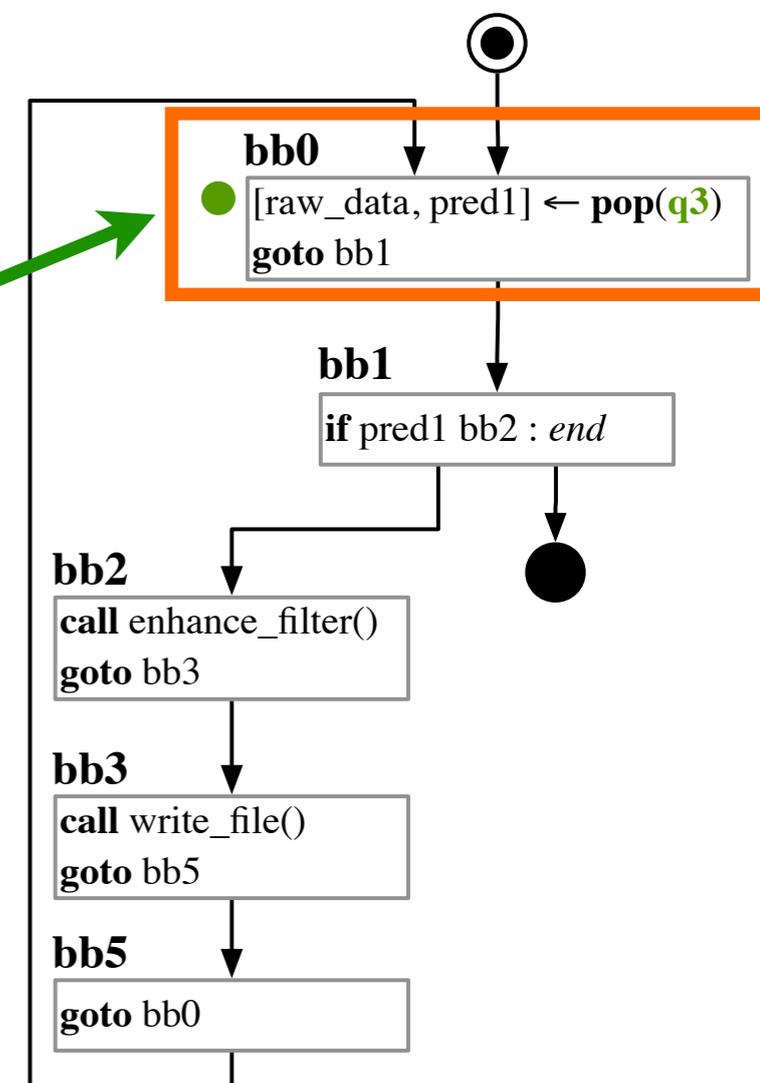
Stage 1



Stage 2



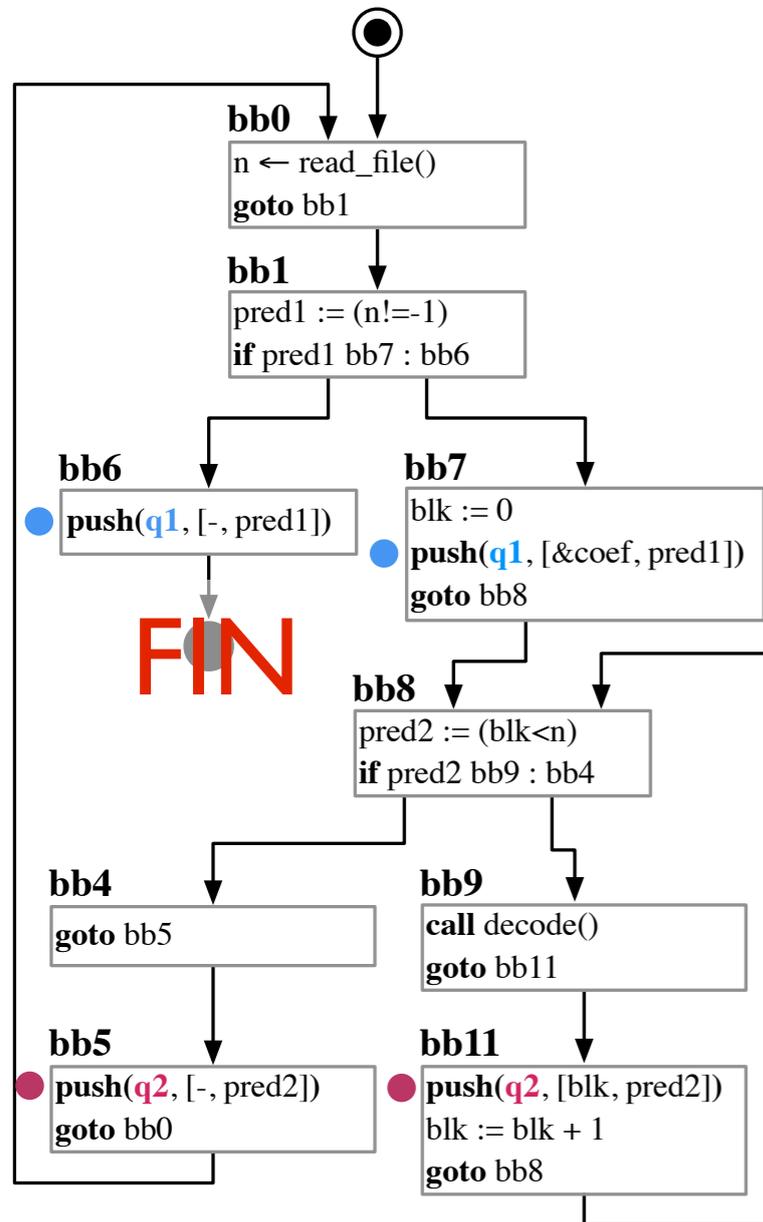
Stage 3



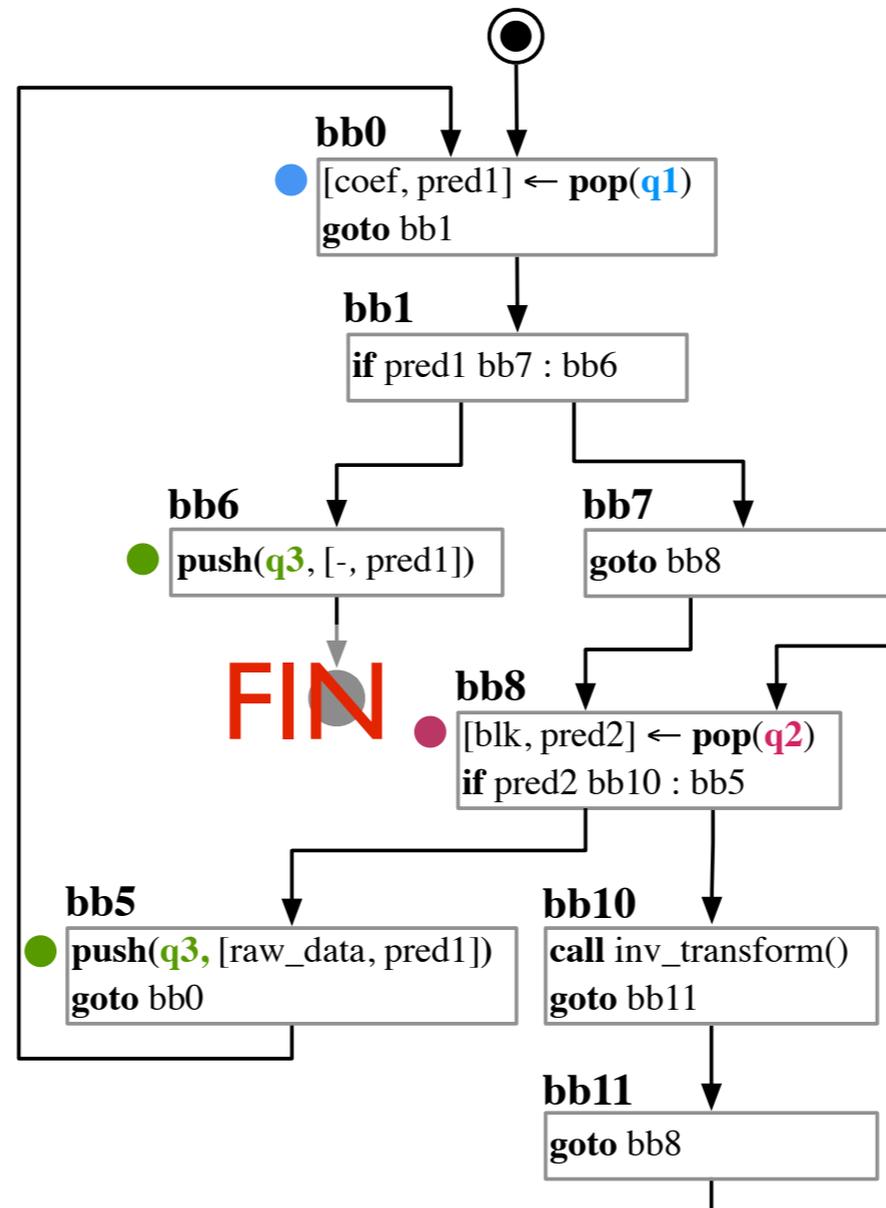
Communication



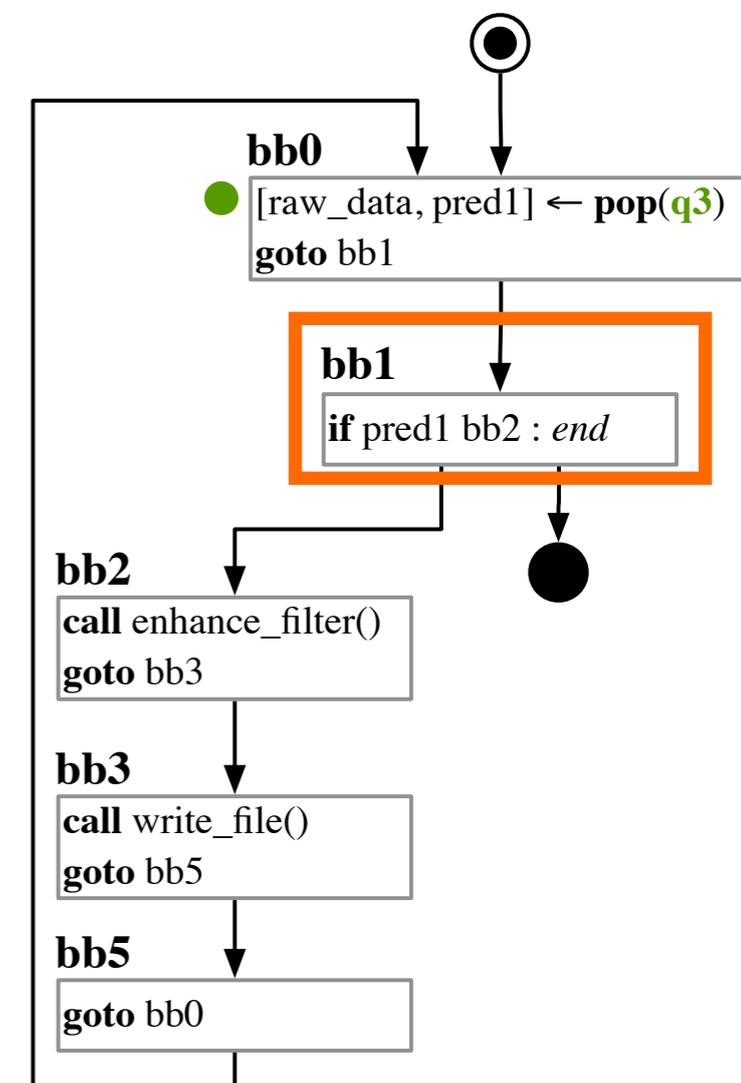
Stage 1



Stage 2



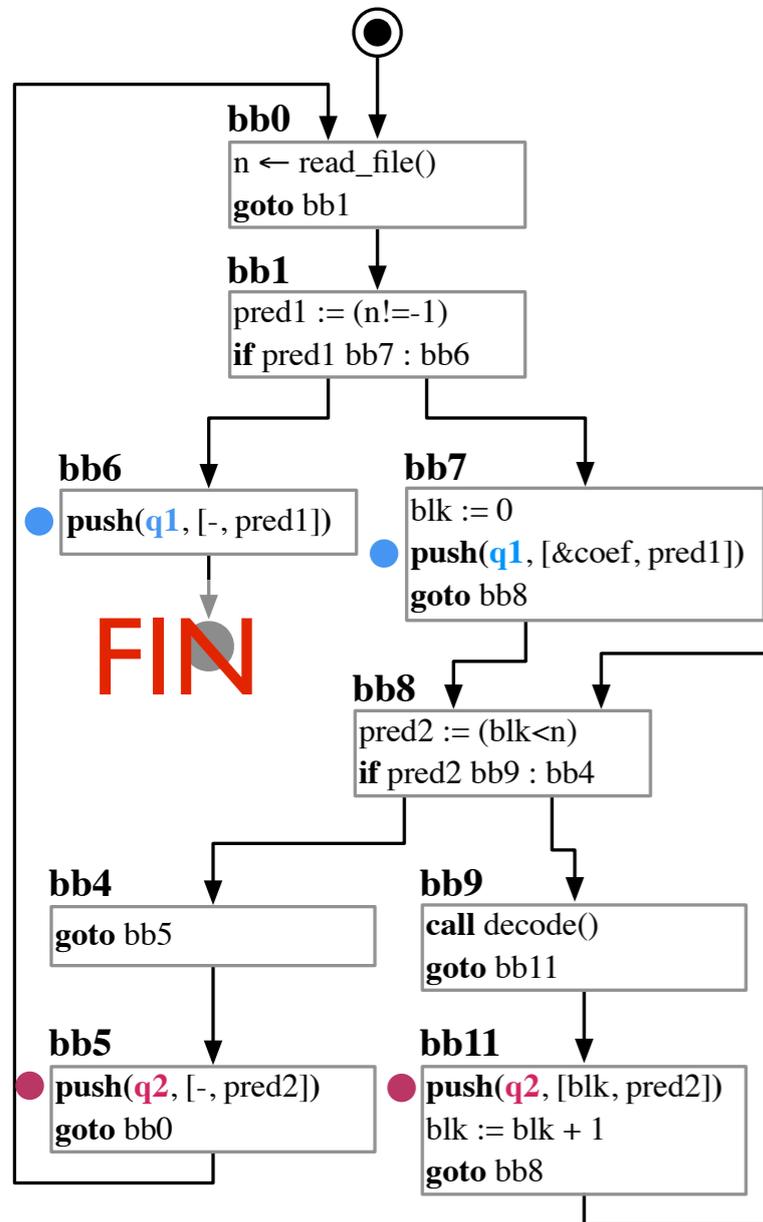
Stage 3



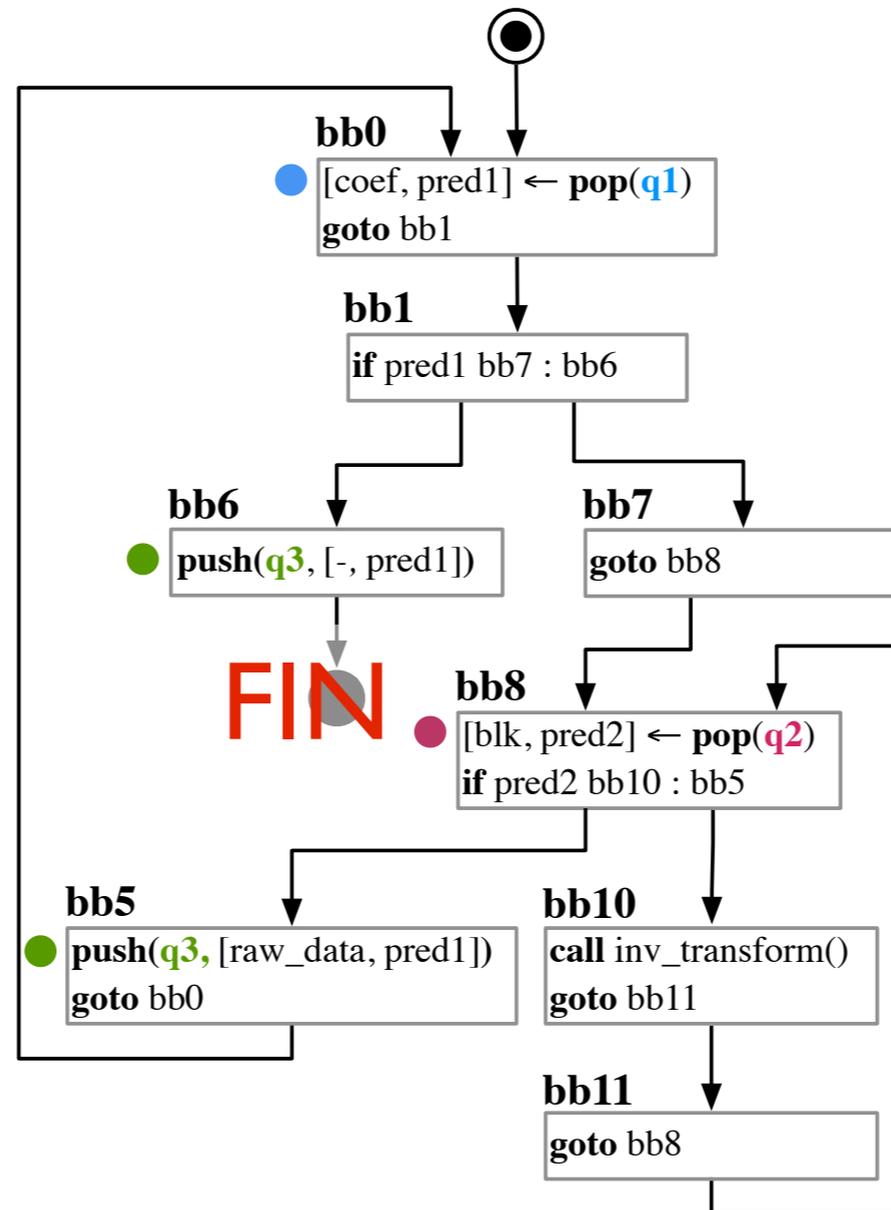
Communication



Stage 1



Stage 2



Stage 3

