Reducing the Memory Footprint of IFDS-based Data-Flow Analyses Using Fine-Grained Garbage Collection

Dongjie He[1], Yujiang Gui[1], Yaoqing Gao[2], Jingling Xue[1]

[1]School of Computer Science and Engineering
University of New South Wales

[2]Huawei Toronto Research Center

Presenter: Yujiang Gui

32nd ISSTA, July 2023

The IFDS Algorithm

- Transforms an analysis into a graph-reachability problem
- Solves a wide range of data-flow analyses
  - Compiler optimization
  - Bug detection (ASE’18)
  - Taint analysis (PLDI’14, FSE’14)
  - Pointer analysis (ECOOP’16, ASE’21, TSE’23)
  - Typestate-like analysis (OOPSLA’08, PLDI’14)
  - …
- Has been implemented in many tools

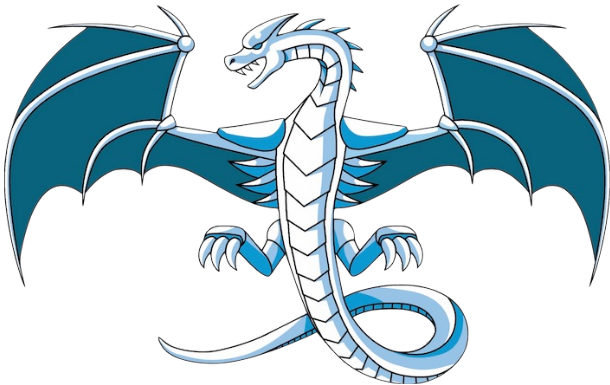

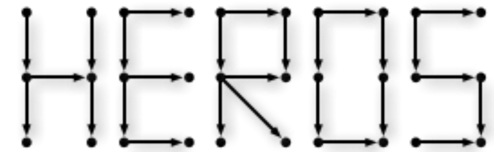

Reps, Thomas, Susan Horwitz, and Mooly Sagiv. “Precise Interprocedural Dataflow Analysis via Graph Reachability.”, POPL’95.

A Running Example: Taint Analysis

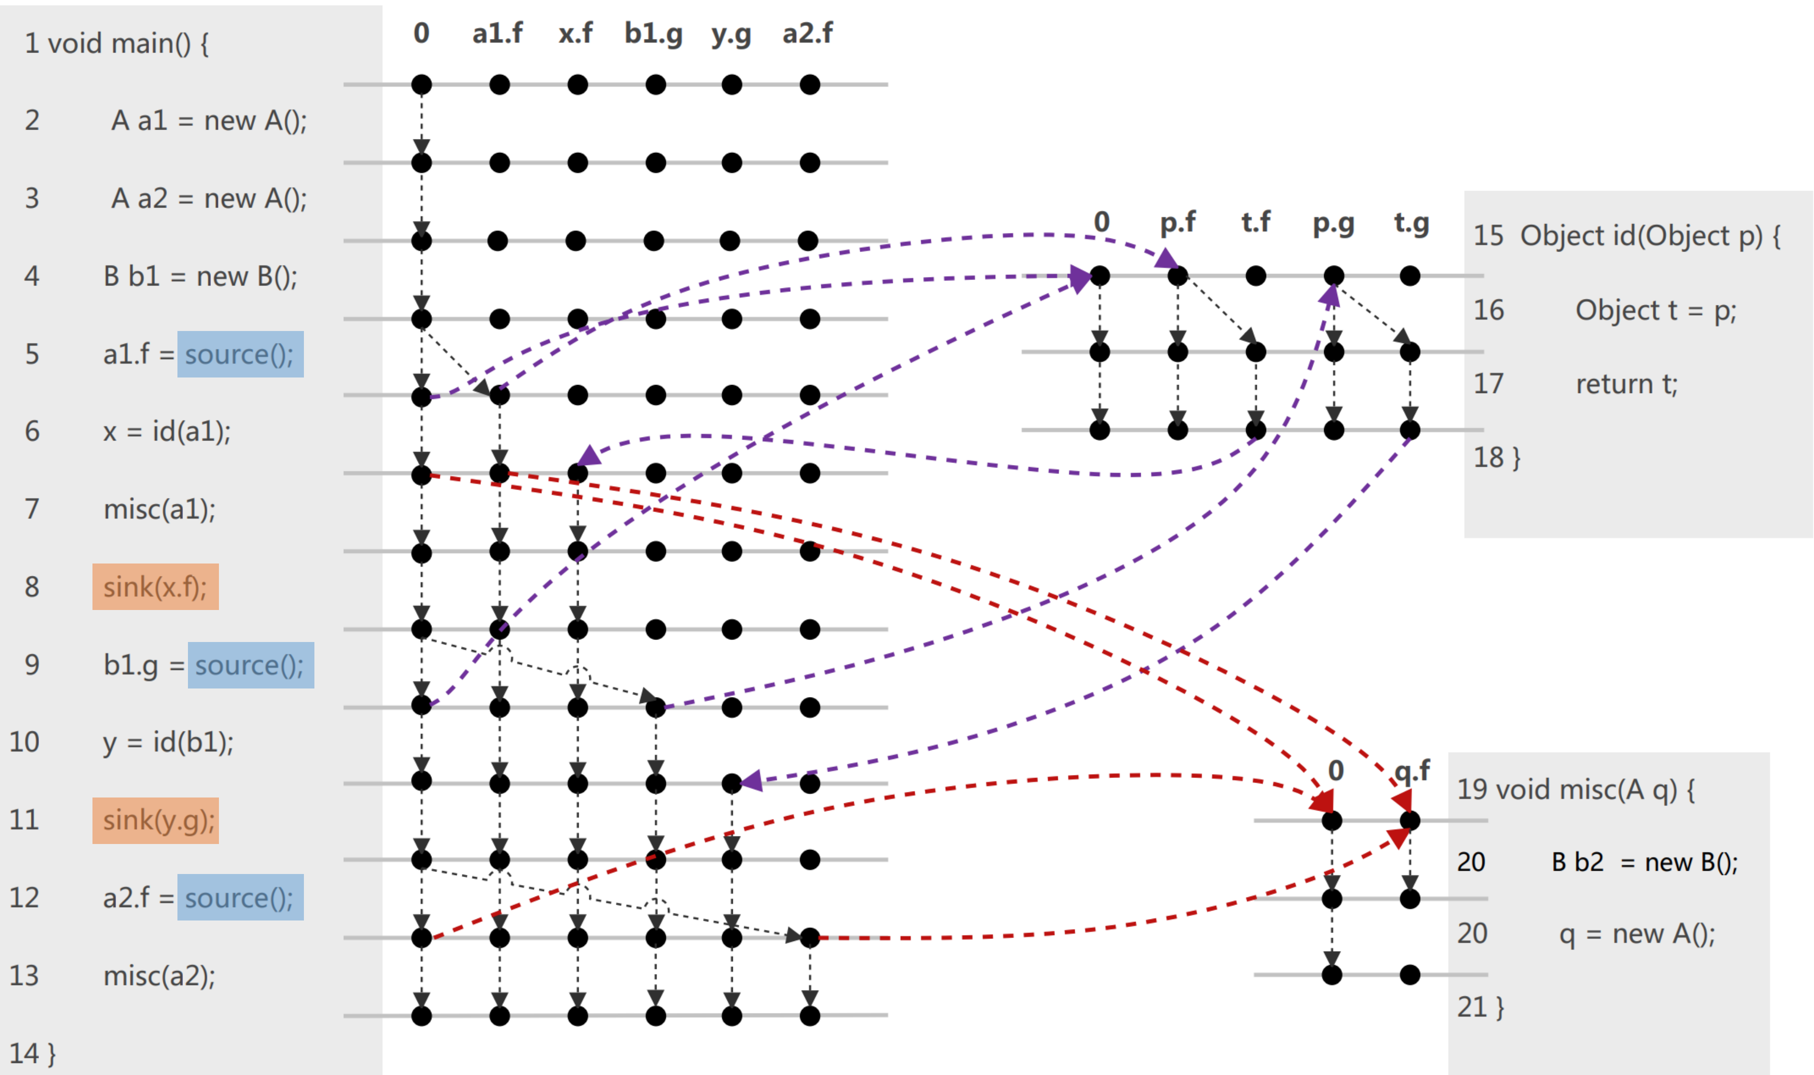

A Running Example: Taint Analysis

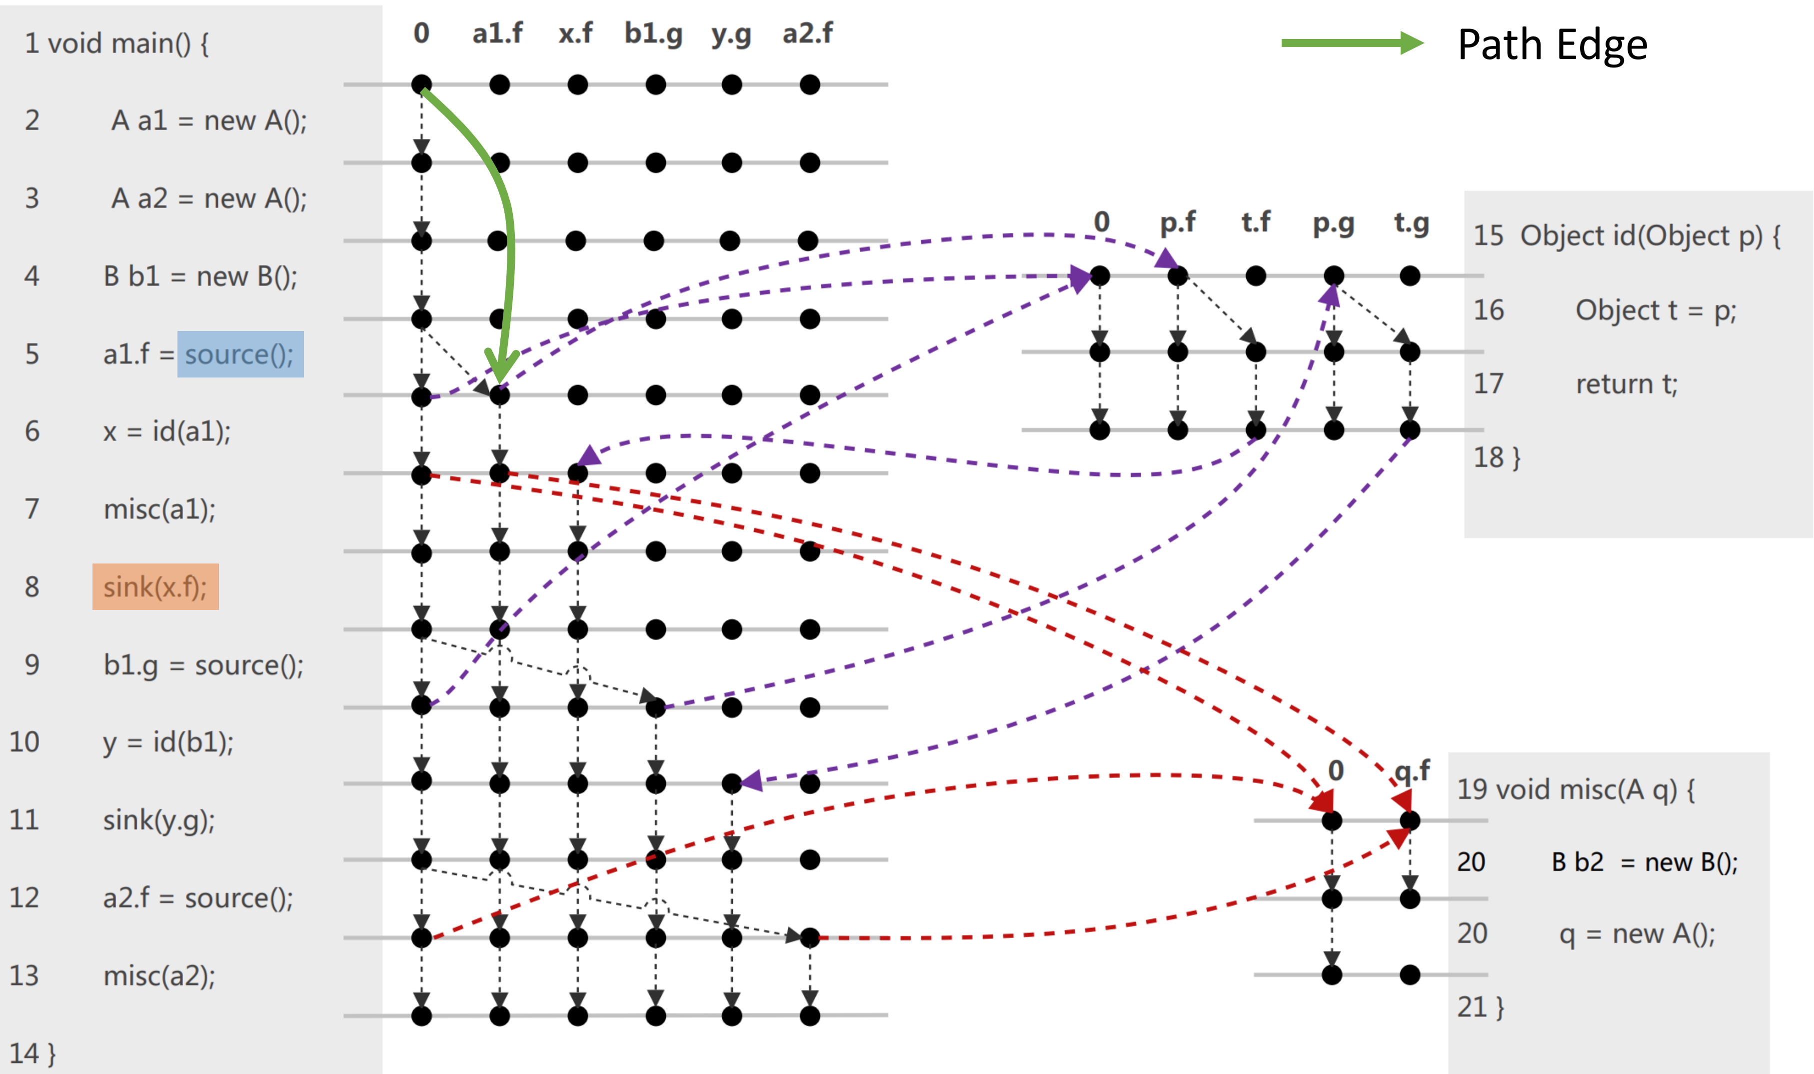

A Running Example: Taint Analysis

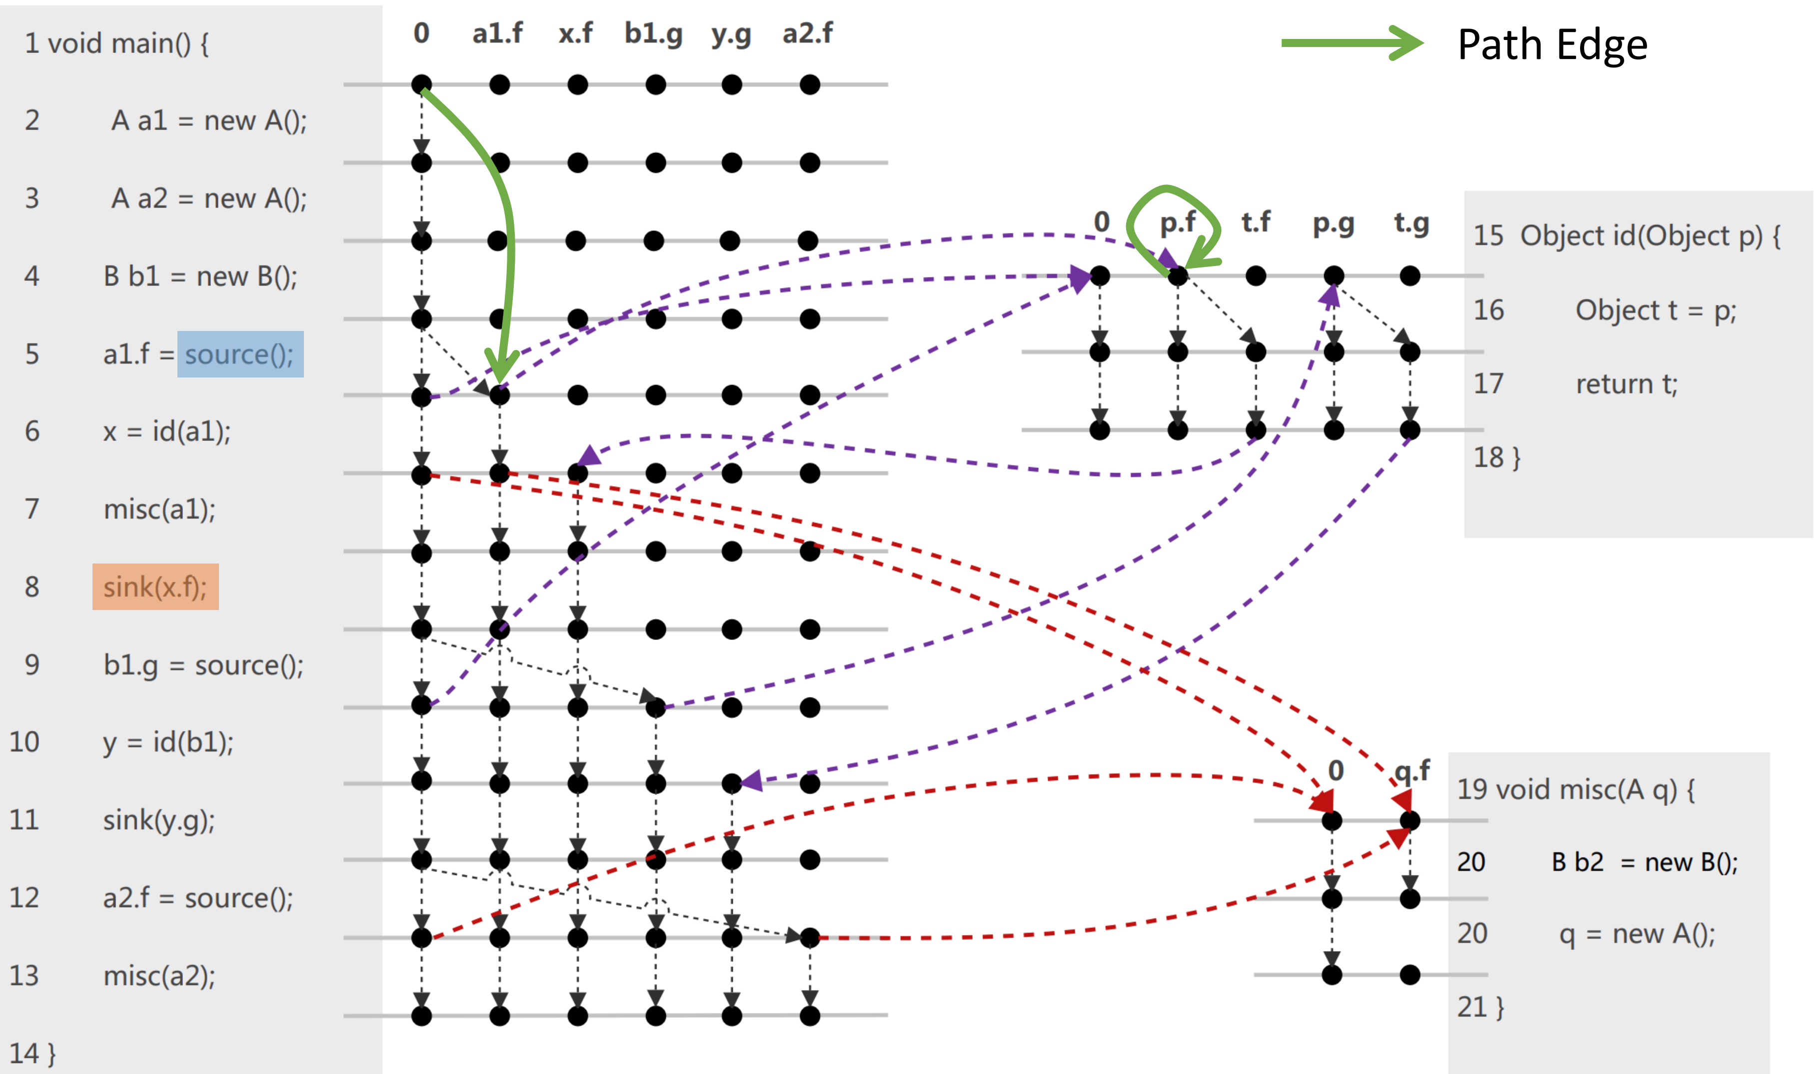

A Running Example: Taint Analysis

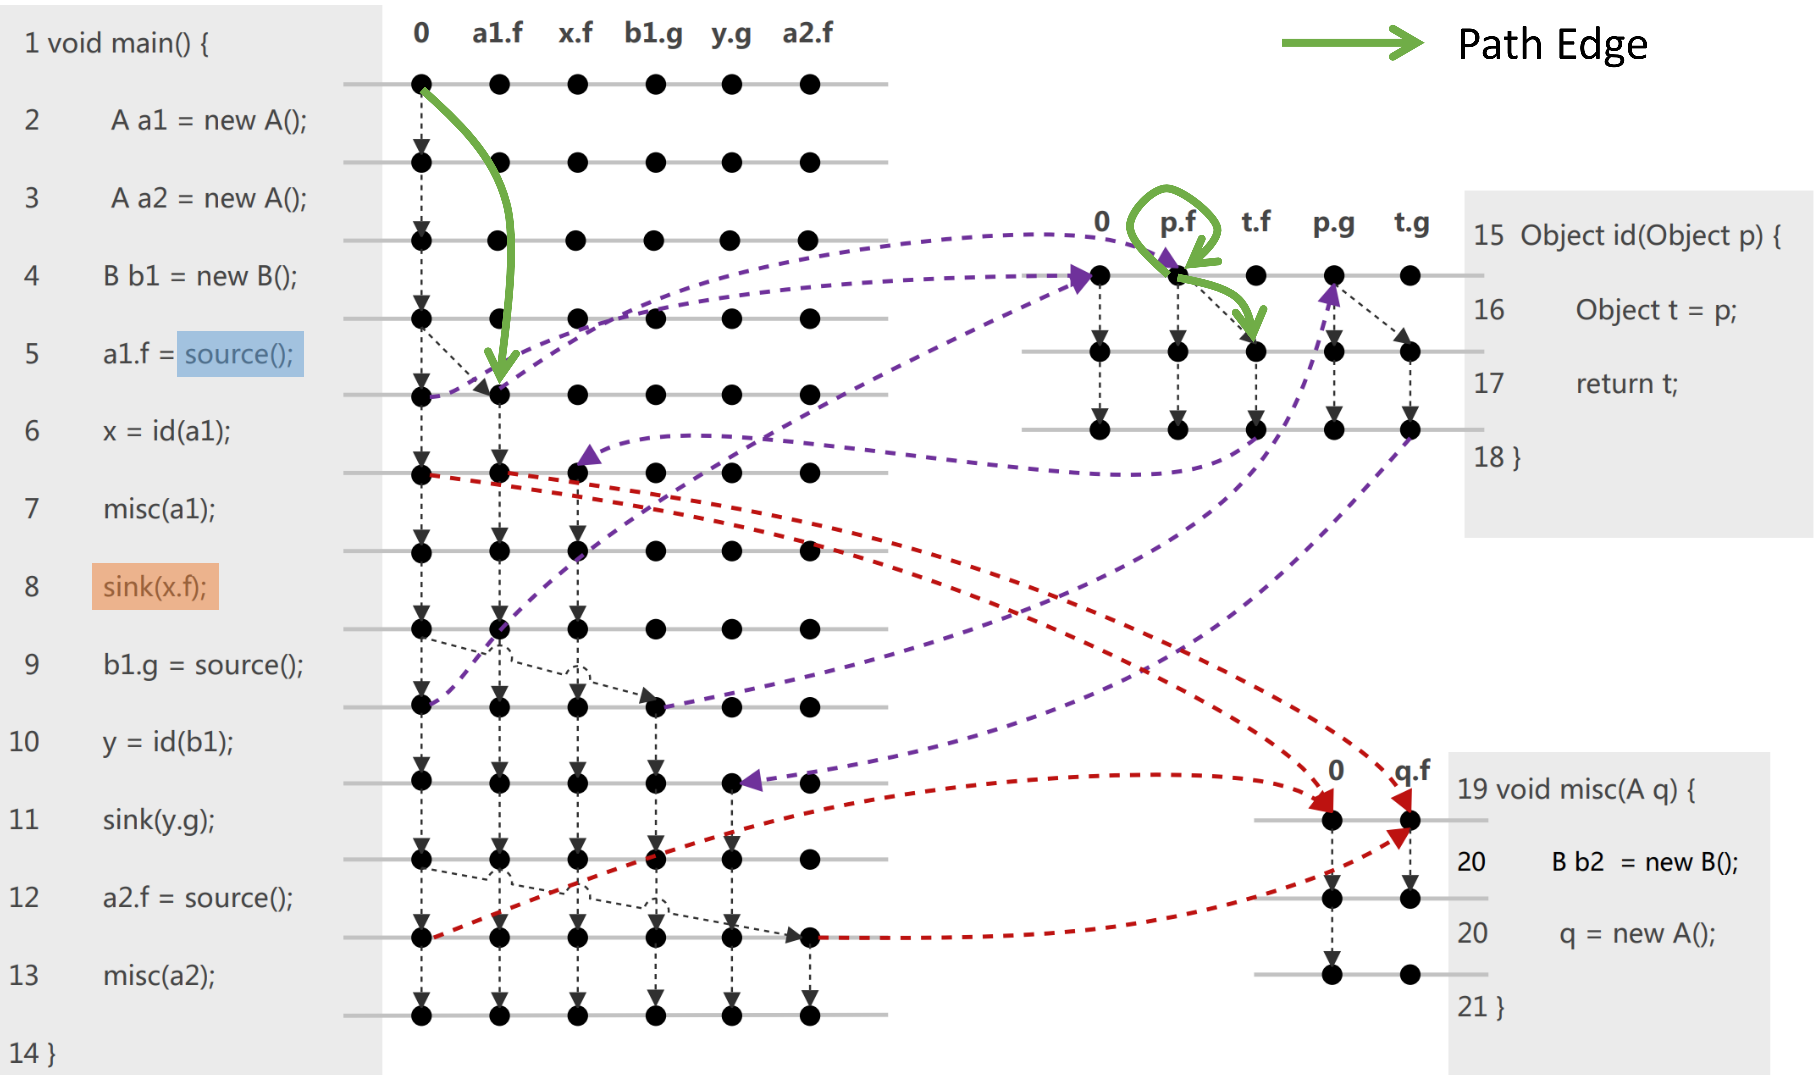

A Running Example: Taint Analysis

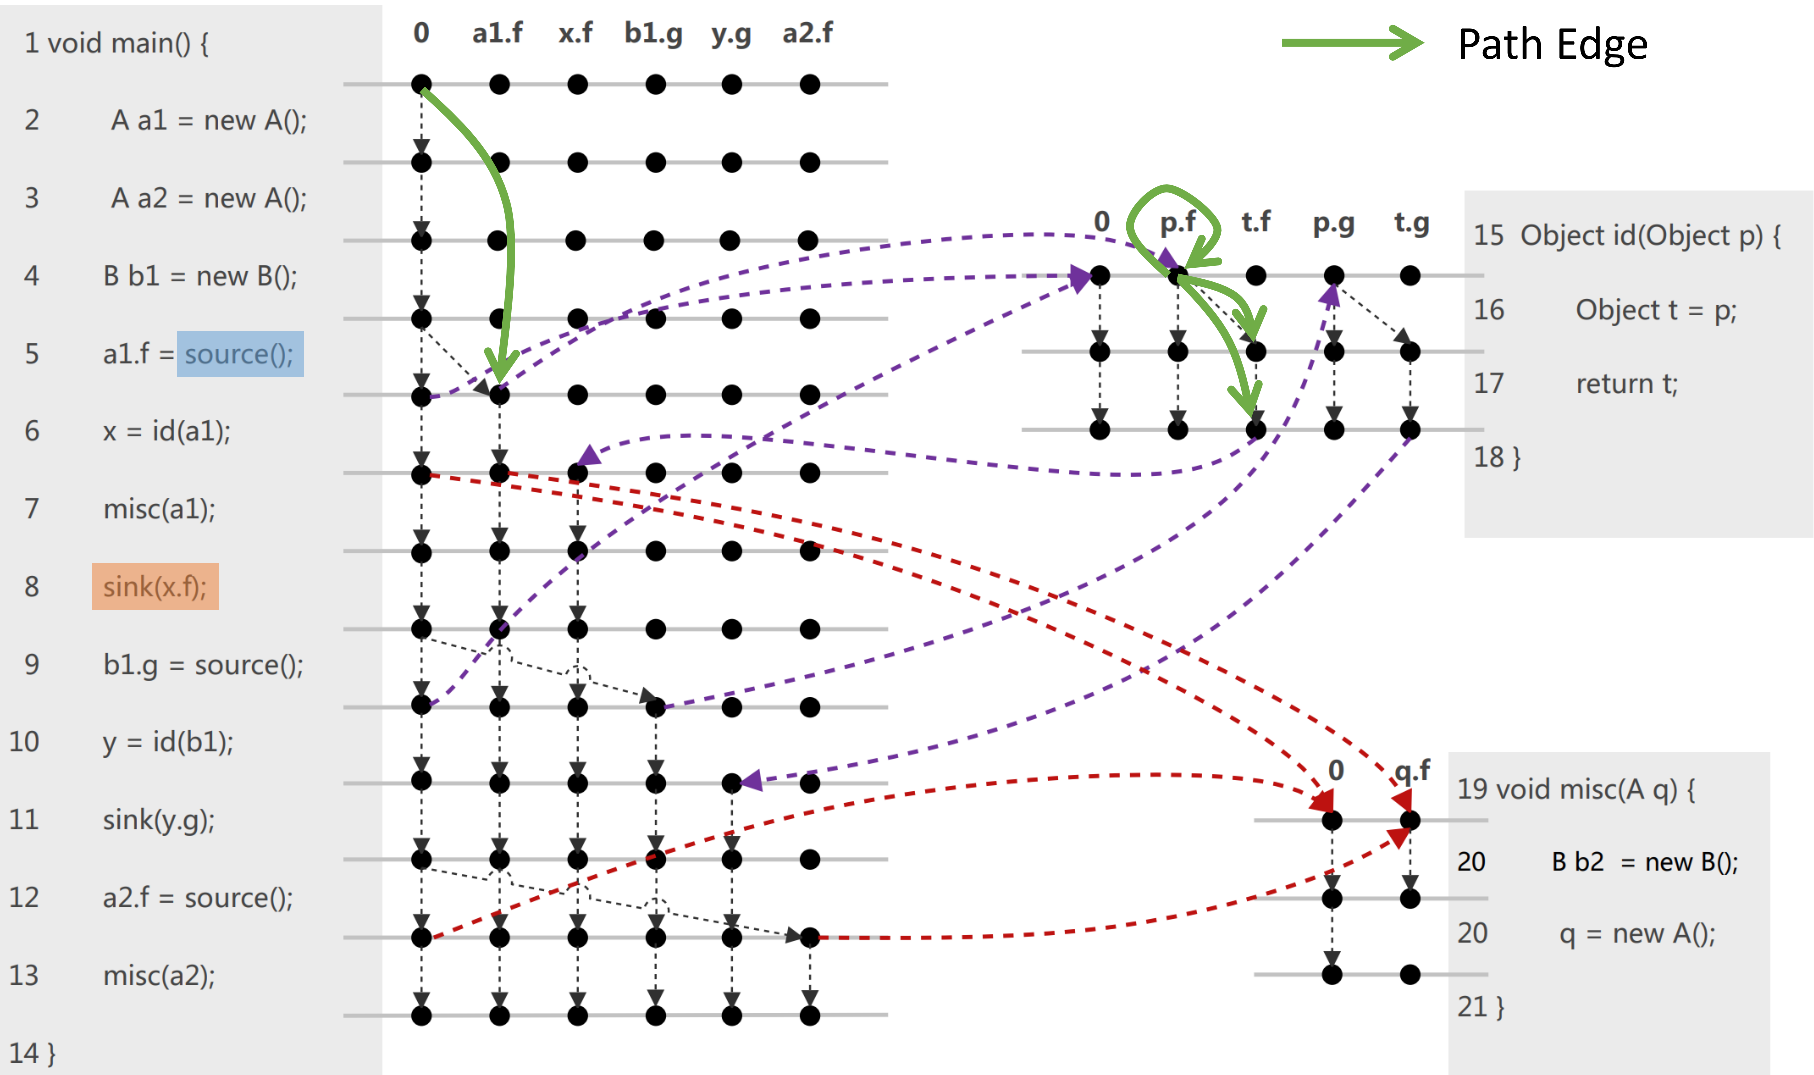

A Running Example: Taint Analysis

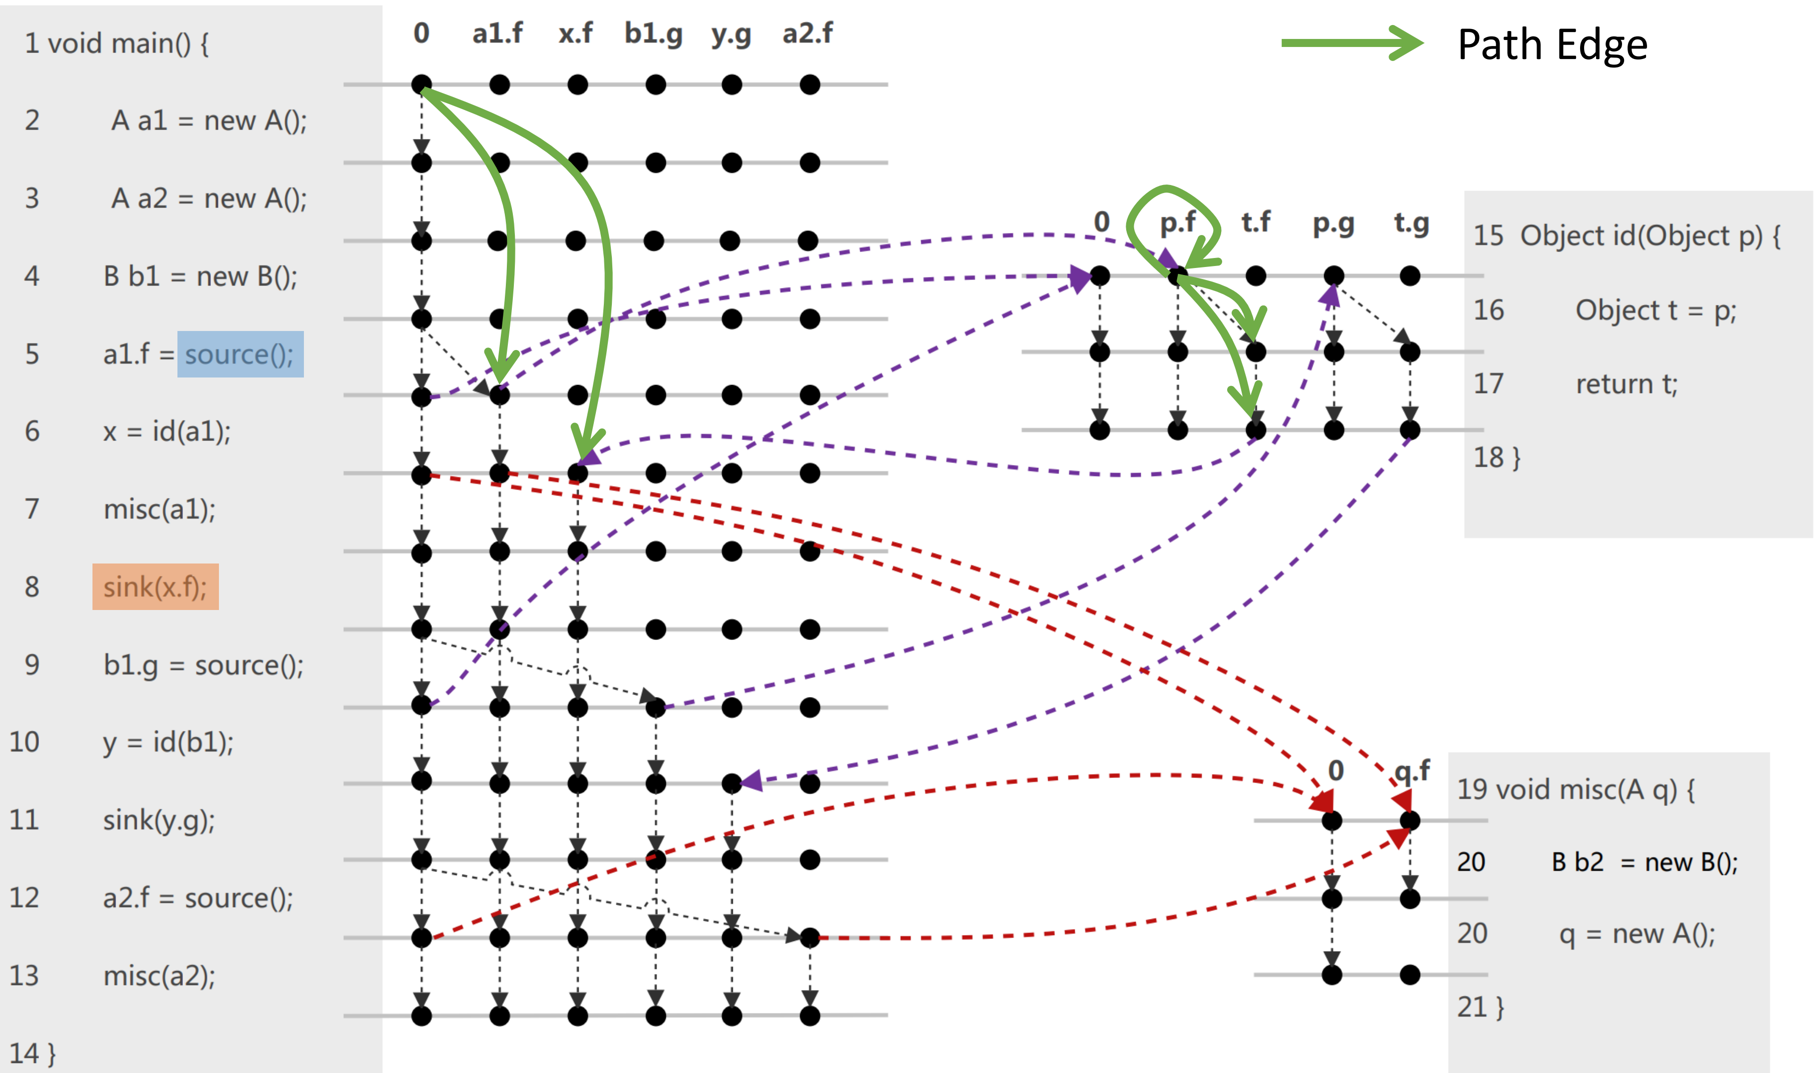

A Running Example: Taint Analysis

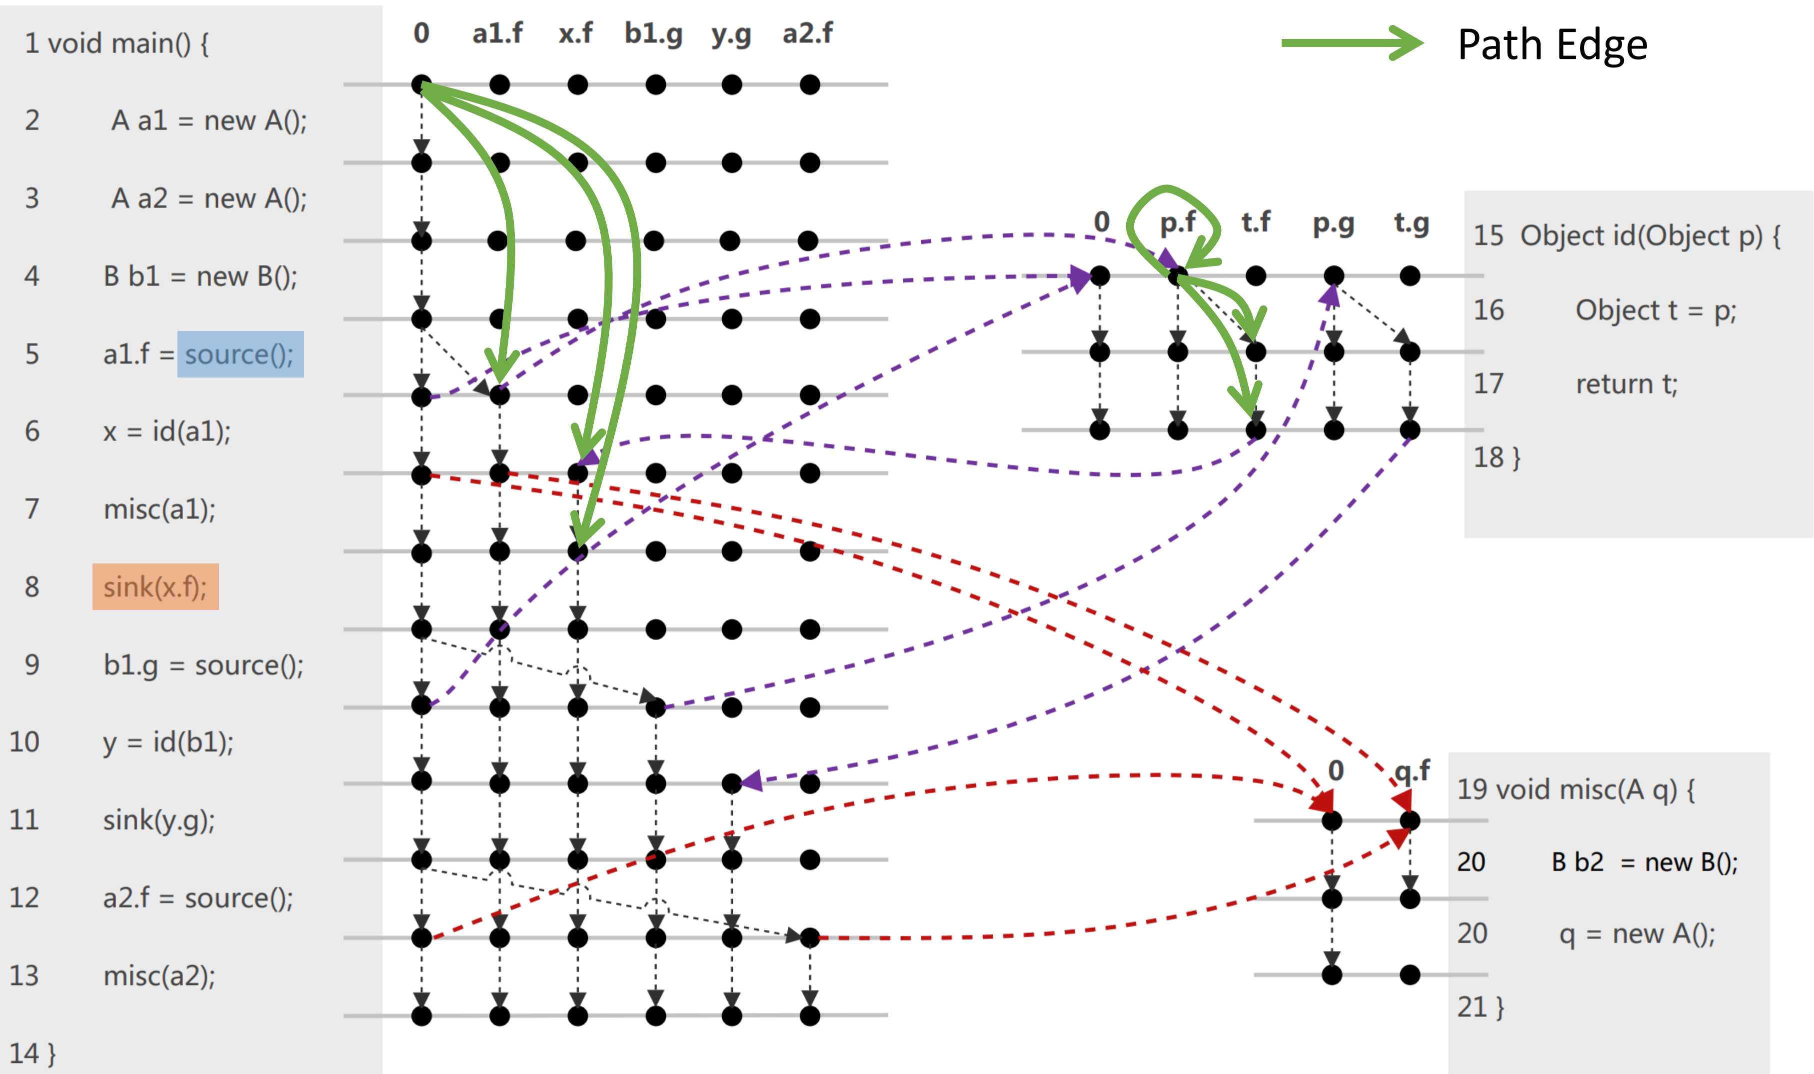

A Running Example: Taint Analysis

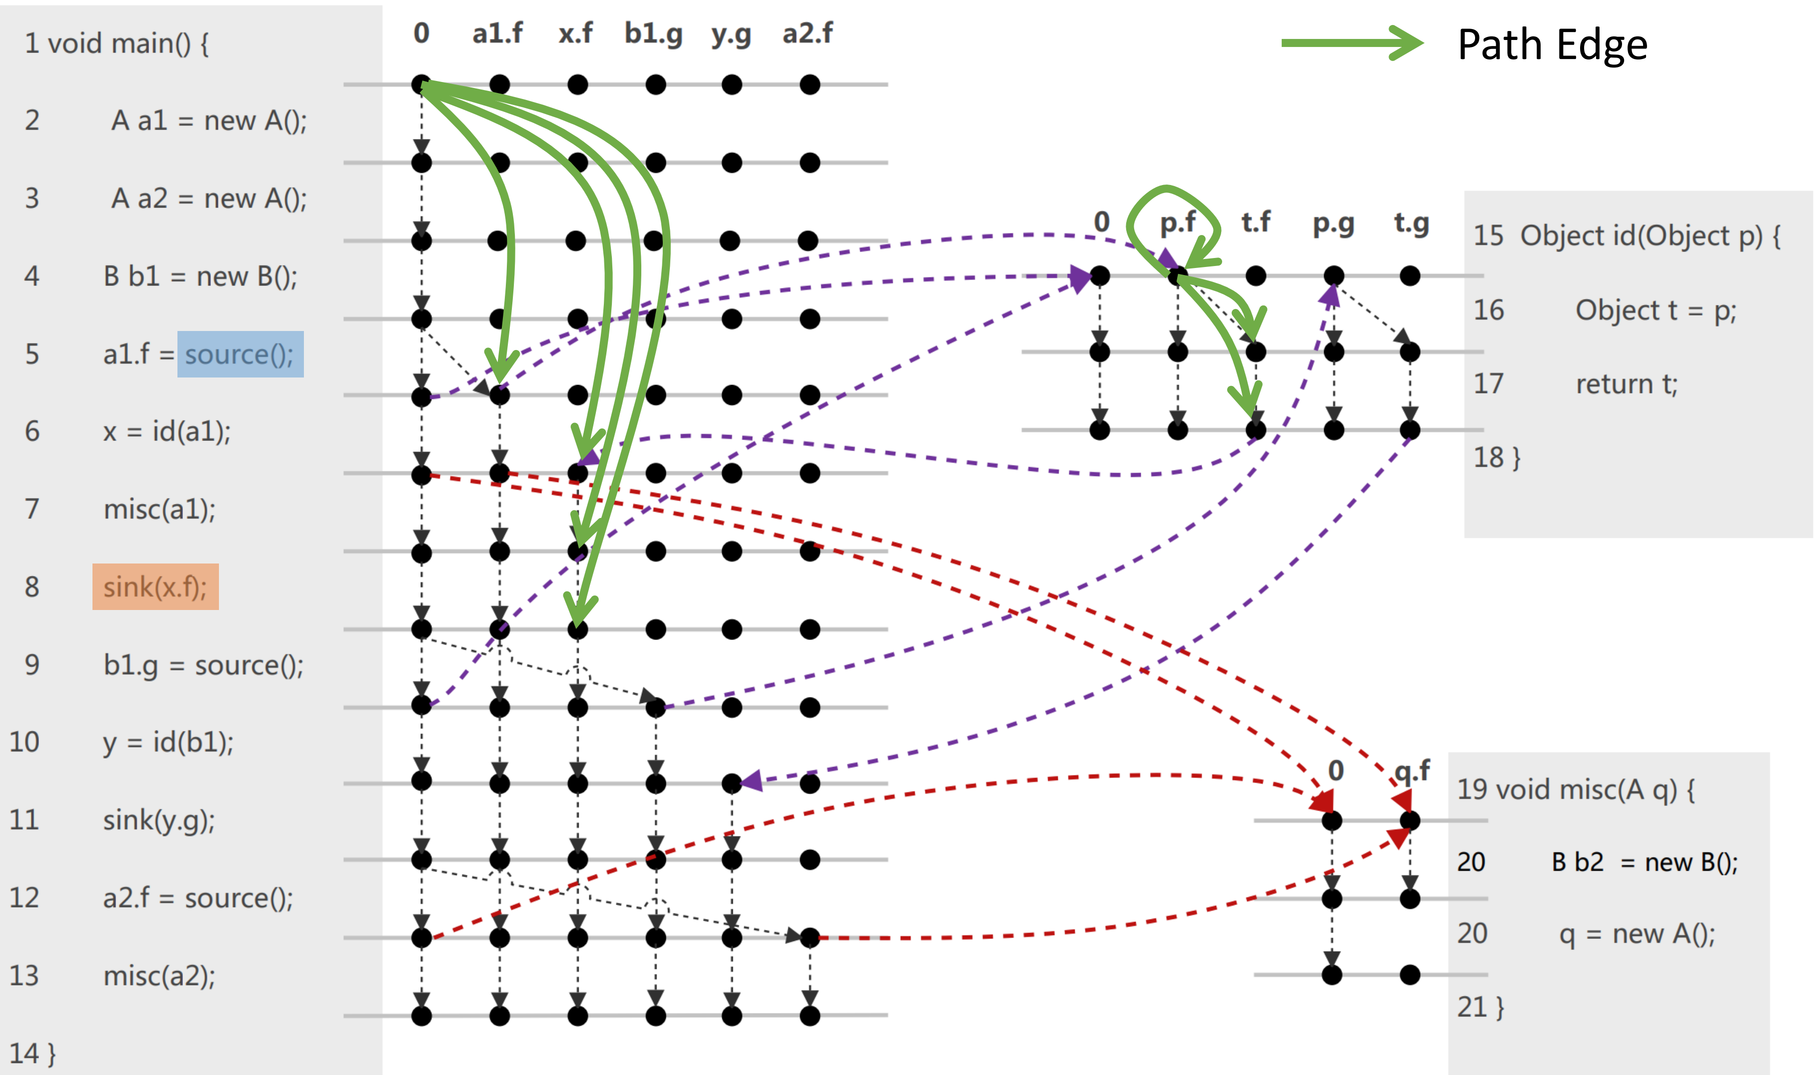

The Drawback: Memory-Intensive

- Maintains a huge number of path edges

#PathEdge (M) | Memory Usage (GB)

140 | 120 | 100 | 80 | 60 | 40 | 20 | 0

600 | 500 | 400 | 300 | 200 | 100 | 0

Weather | Metadata | Lumicall | Book Reader | Overchan

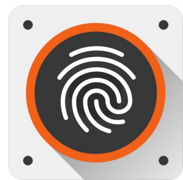

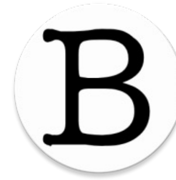

The Drawback: Memory-Intensive

- Maintains a huge number of path edges

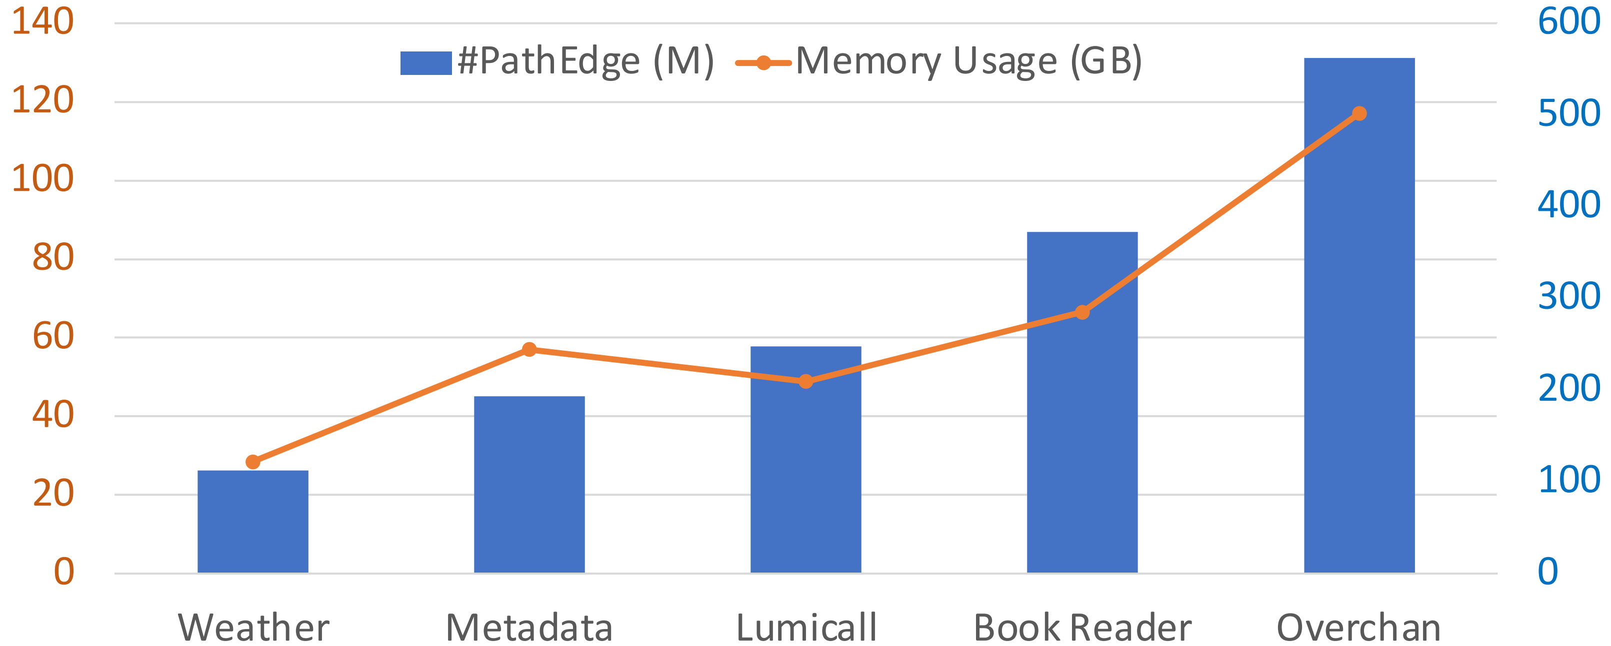

- Fails to solve taint analysis on some Android apps
  - Even on a server with **730GB RAM**

- Slows down the analysis
  - Frequent re-hashing operations

Objective

Improving the Scalability and Efficiency

Objective

Improving the Scalability and Efficiency

- Discovers facts only at some program points
    - Taint analysis

Objective

Improving the Scalability and Efficiency

- Discovers facts only at some program points
  - Taint analysis
- Non-live path edges
  - Visited only once (**86.97%**)
  - Waste memory resources

Objective

Improving the Scalability and Efficiency

- Discovers facts only at some program points
  - Taint analysis
- Non-live path edges
  - Visited only once (**86.97%**)
  - Waste memory resources
- **Garbage collection**! But how to…
  - Preserve important properties
    - Correctness
    - Precision
    - Termination
  - Avoid redundant computations

Existing Approach

Sustainable Solving: Reducing The Memory Footprint of IFDS-Based Data Flow Analyses Using Intelligent Garbage Collection

Steven Arzt
Secure Software Engineering Group,
Fraunhofer Institute for Secure Information Technology
Darmstadt, Germany
Email: steven.arzt@sit.fraunhofer.de

CleanDroid (ICSE'21)

- With **2 major limitations**
  - Coarse-grained
  - Allows redundant computations

Limitation 1: Coarse Granularity

Running garbage collection just before **line 10...**

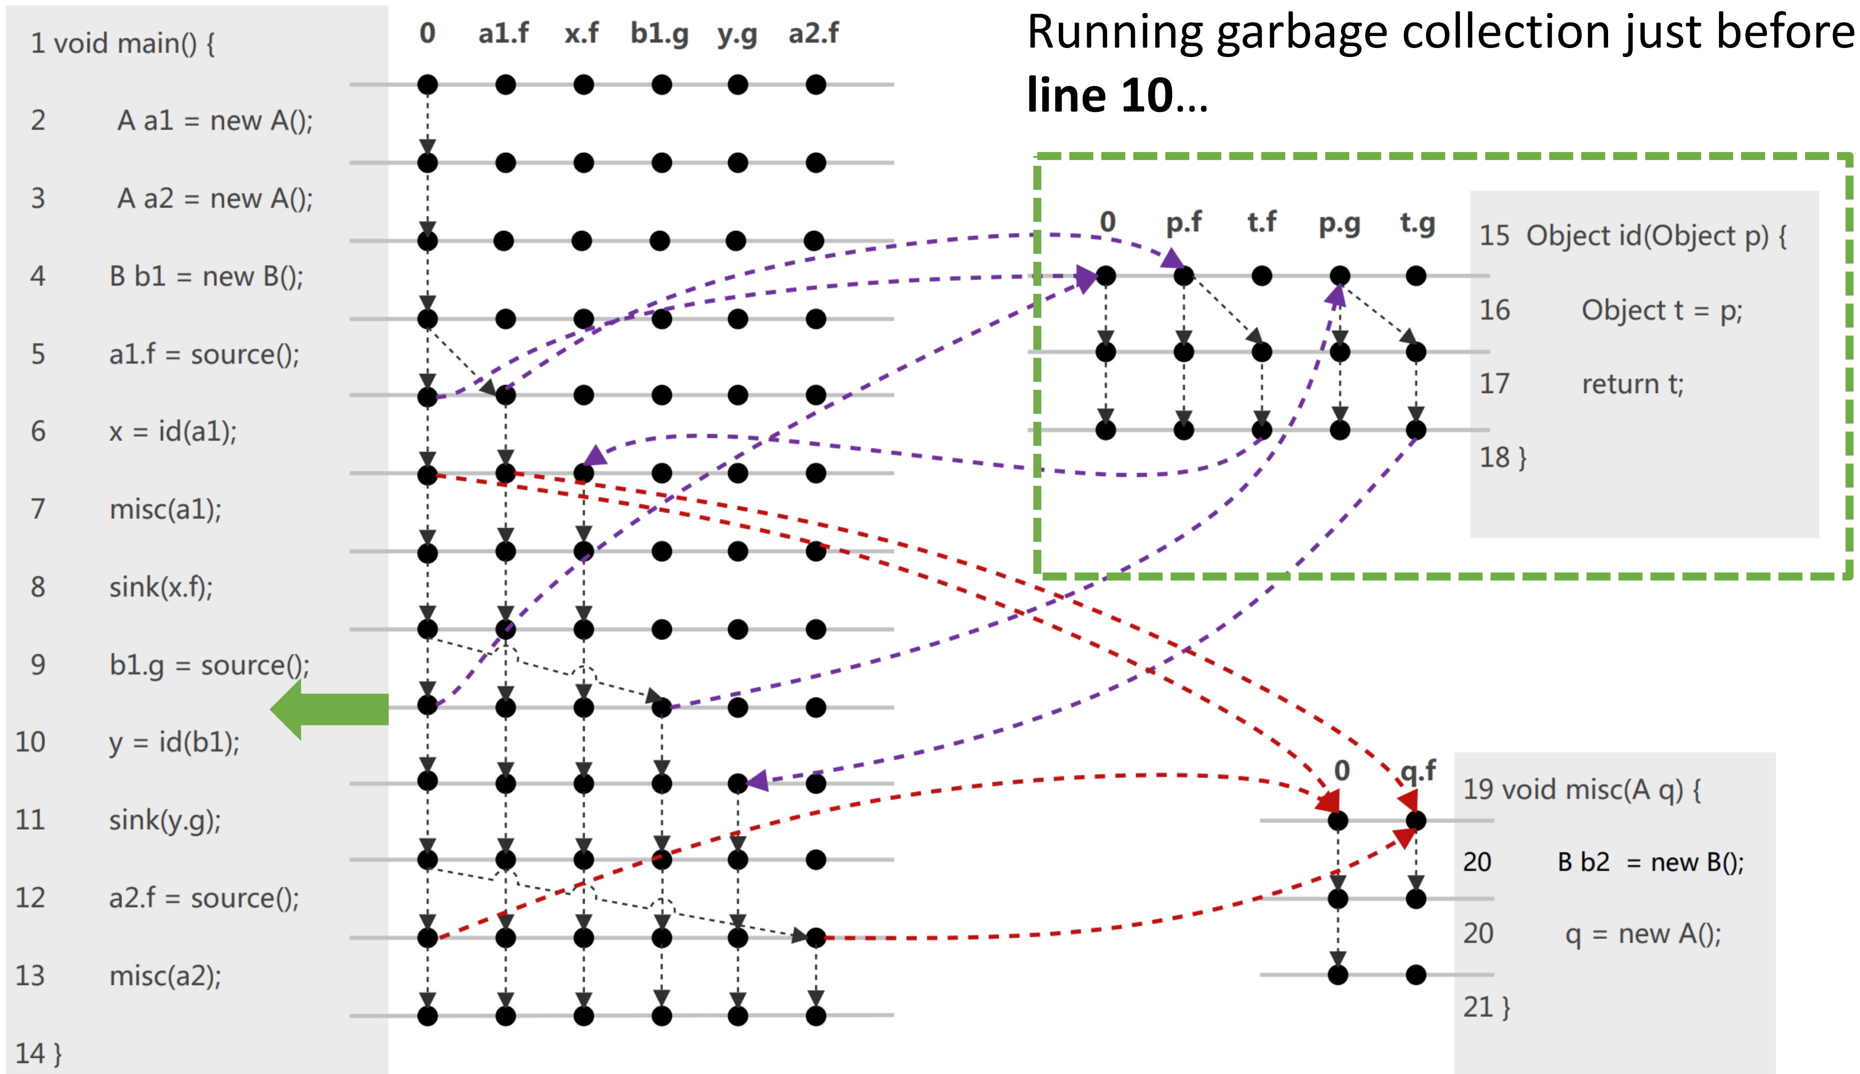

Limitation 1: Coarse Granularity

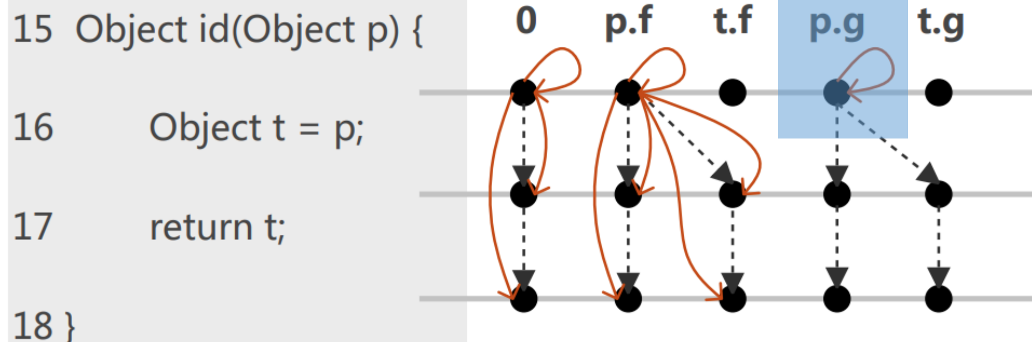

(a) before GC

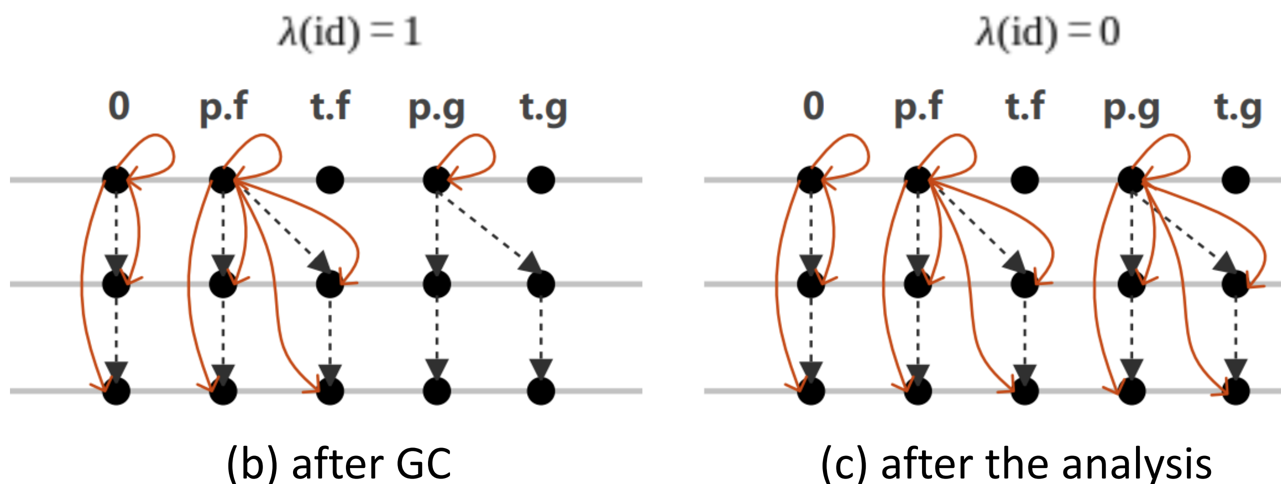

(b) after GC

(c) after the analysis

Maximum #PathEdge maintained: 13.

Limitation 1: Coarse Granularity

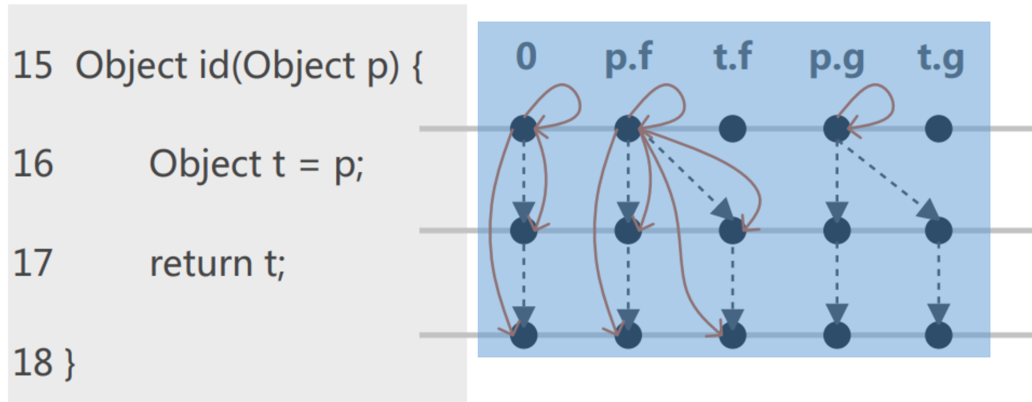

(a) before GC

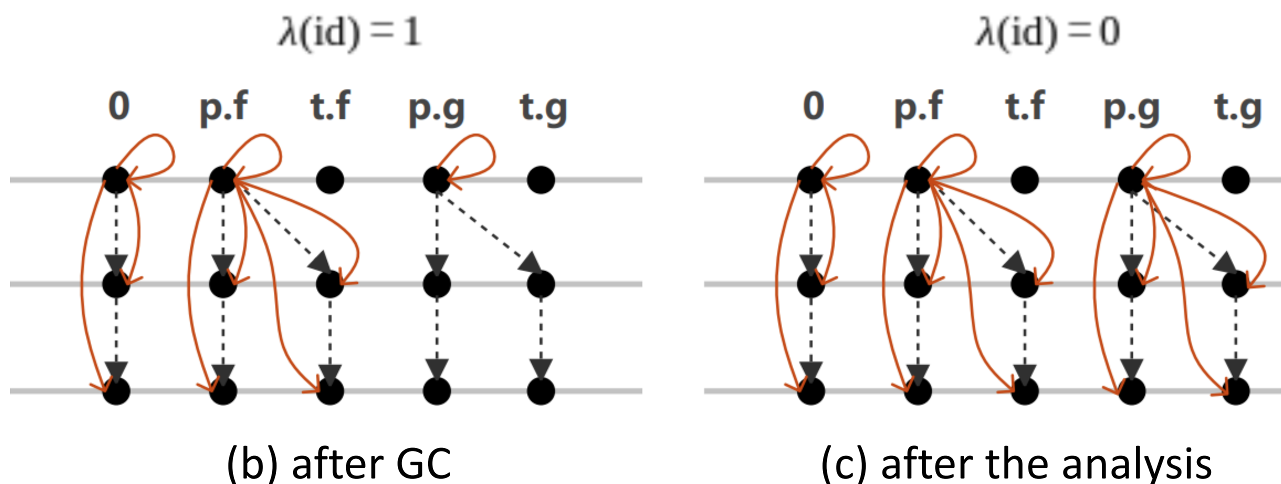

(b) after GC

(c) after the analysis

Method-Level: Coarse-Grained

Fine-Grained Data-Fact-Level GC

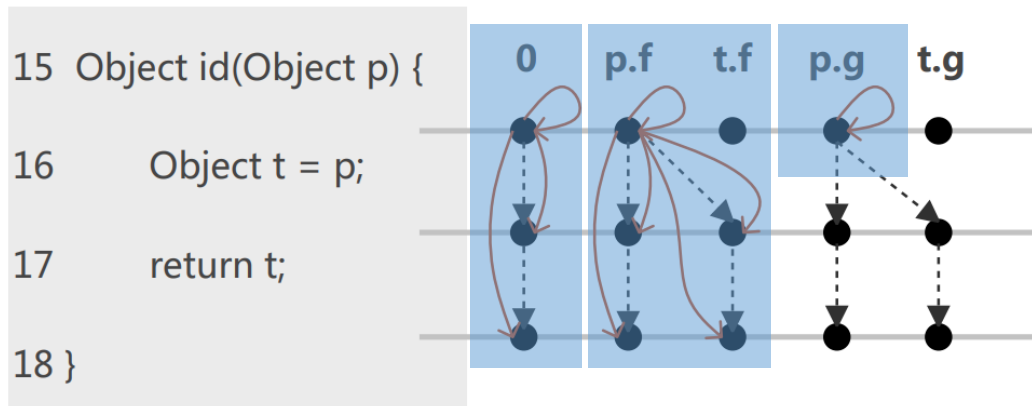

(a) before GC

Observation 1: The path edges with different anchor sites are handled **independently**.

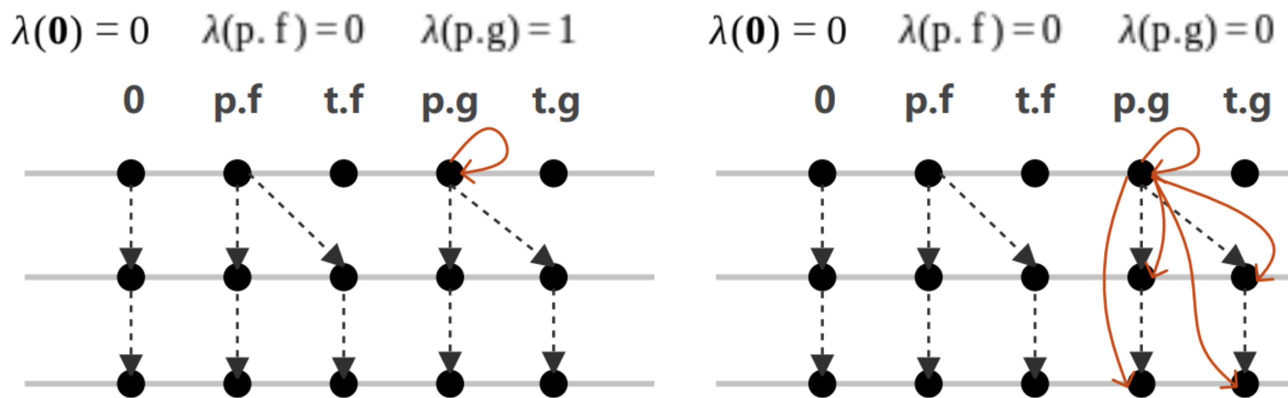

(b) after GC

(c) after the analysis

Maximum #PathEdge maintained: 9.

Limitation 2: Redundant Computations

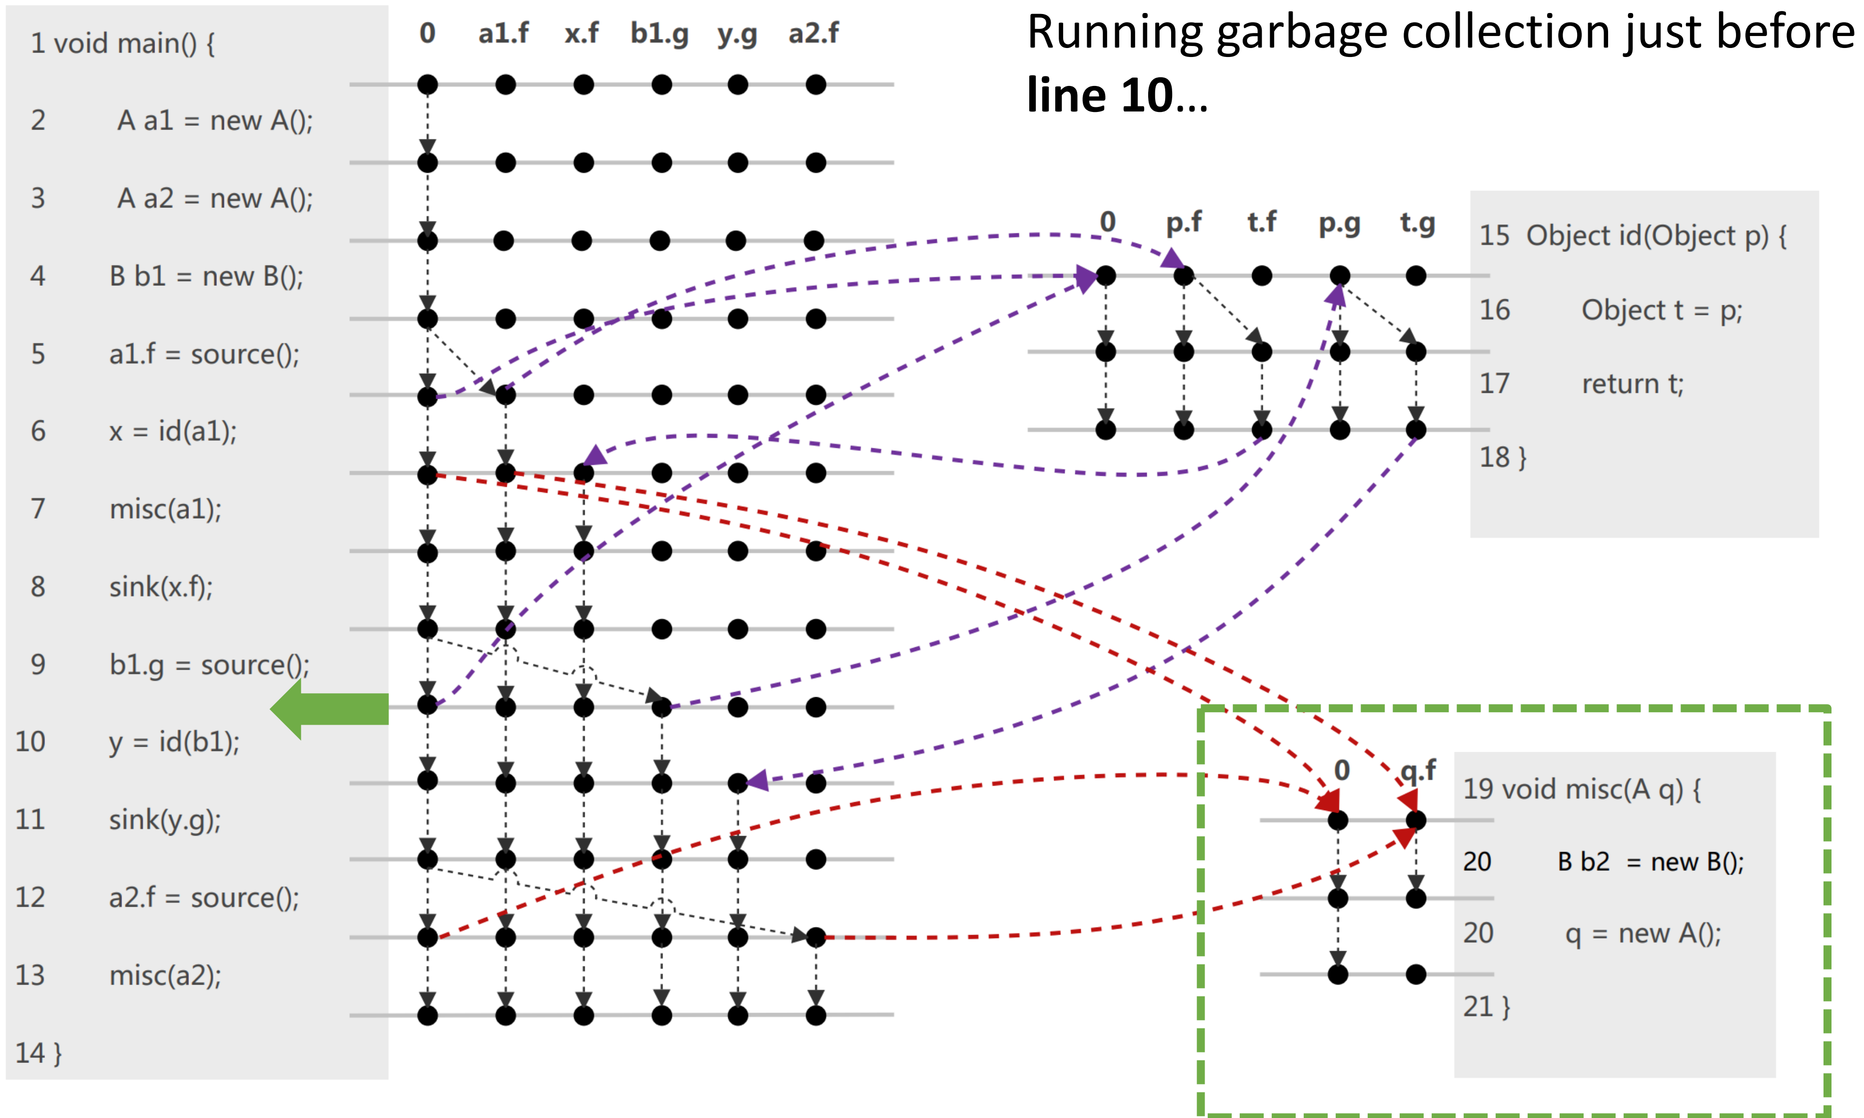

Limitation 2: Redundant Computations

```
19 void misc(A q) {
20        B b2  = new B();
20         q = new A();
21 }
```

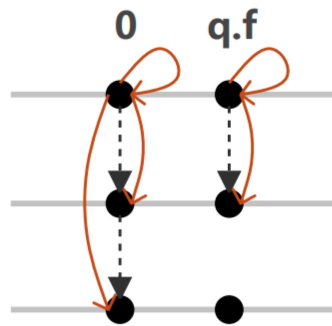

(a) before GC

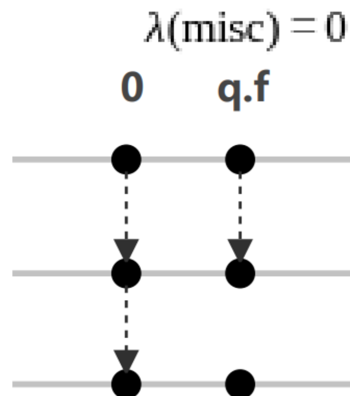

(b) after GC

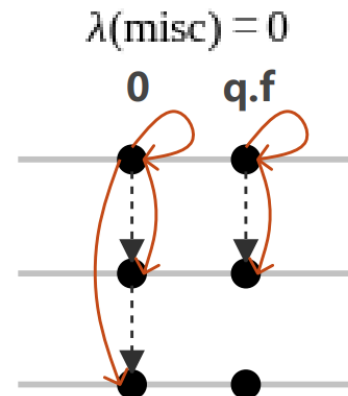

(c) after the analysis

The method `misc()` is analyzed redundantly.

Avoiding Redundant Computations

```
19 void misc(A q) {
20         B b2  = new B();
20          q = new A();
21 }
```

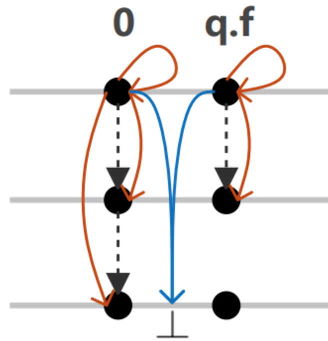

(a) before GC

Observation 2: All path edges sharing the same anchor site are generated from the same self-loop path edge.

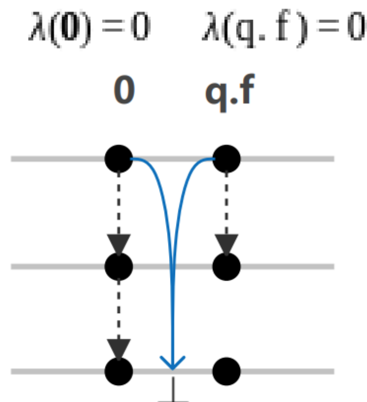

(b) after GC

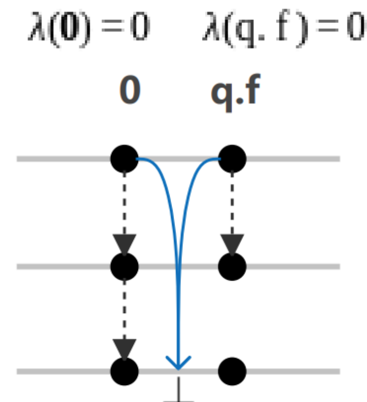

(c) after the analysis

With redundancy-avoiding edges, the method `misc()` is analyzed only once.

Data-Fact-Level Path-Edge GC Algorithm

```
1  Algorithm IFDS(G#_IP = (N#, E#))
2    InitPECollector()
3    PathEdge ← W ← S ← ∅
4    Propagate(⟨s_main, 0⟩ → ⟨s_main, 0⟩)
5    while W ≠ ∅ do
6      Pop ⟨s_m, d1⟩ → ⟨n, d2⟩ from W
7      if n is a call node then
8        Let m′ be the method called at n and n′ be the return node of n
9        for d3 such that ⟨n, d2⟩ → ⟨s_m′, d3⟩ ∈ E# do
10         ADG = ADG ∪ {⟨s_m, d1⟩ → ⟨s_m′, d3⟩}
11         Propagate(⟨s_m′, d3⟩ → ⟨s_m′, d3⟩)
12         for ⟨s_m′, d3⟩ → ⟨e_m′, d4⟩ ∈ S ∧ ⟨e_m′, d4⟩ → ⟨n′, d5⟩ ∈ E# do
13           Propagate(⟨s_m, d1⟩ → ⟨n′, d5⟩)
14       for d3 such that ⟨n, d2⟩ → ⟨n′, d3⟩ ∈ E# do
15         Propagate(⟨s_m, d1⟩ → ⟨n′, d3⟩)
16     if n is an exit node then
17       if ⟨s_m, d1⟩ → ⟨n, d2⟩ ∉ S then
18         Insert ⟨s_m, d1⟩ → ⟨n, d2⟩ into S
19         for each call site c that calls m do
20           Let m″ (n″) be the containing method (the return node) of c
21           for ⟨s_m″, d3⟩ → ⟨c, d4⟩ ∈ PathEdge ∧ ⟨c, d4⟩ →
               ⟨s_m, d1⟩ ∈ E# ∧ ⟨n, d2⟩ → ⟨n″, d5⟩ ∈ E# do
22             Propagate(⟨s_m″, d3⟩ → ⟨n″, d5⟩)
23     if n is a normal node or a return node then
24       for ⟨n′, d3⟩ such that ⟨n, d2⟩ → ⟨n′, d3⟩ ∈ E# do
25         Propagate(⟨s_m, d1⟩ → ⟨n′, d3⟩)
26     OnEdgeProcessed(⟨s_m, d1⟩ → ⟨n, d2⟩)
27 Procedure Propagate(⟨s_m, d1⟩ → ⟨n, d2⟩)
28   if ⟨s_m, d1⟩ → ⟨n, d2⟩ ∉ PathEdge then
29     if ⟨s_m, d1⟩ = ⟨n, d2⟩ then                    // a self-loop edge
30       if ⟨s_m, d1⟩ → ⊥ ∈ RAEdges then
31         return // avoid redundant re-computations
32       else
33         RAEdges = RAEdges ∪ {⟨s_m, d1⟩ → ⊥}
34     OnEdgeScheduled(⟨s_m, d1⟩ → ⟨n, d2⟩)
35     Insert ⟨s_m, d1⟩ → ⟨n, d2⟩ into both PathEdge and W
```

```
1  Procedure InitPECollector()
2    λ = {α ↦ 0}
3    RAEdges = C = ∅
4  Procedure OnEdgeScheduled(⟨s_m, d1⟩ → ⟨n, d2⟩)
5    Consume(⟨s_m, d1⟩ → ⟨n, d2⟩) // do some analysis-specific task
6    λ = { α ↦ λ(α) + 1   if α = ⟨s_m, d1⟩
         { α ↦ λ(α)       otherwise
7    C = C ∪ {α}
8  Procedure OnEdgeProcessed(⟨s_m, d1⟩ → ⟨n, d2⟩)
9    λ = { α ↦ λ(α) − 1   if α = ⟨s_m, d1⟩
         { α ↦ λ(α)       otherwise
10 Procedure RunFineGrainedPECollector()
11   foreach α ∈ C do
12     if ∀α′ ∈ ADGTC(α, ADG): λ(α′) = 0 then
         // Remove path edges with α as their anchor site
13       foreach e : ⟨s_m, d1⟩ → ⟨n, d2⟩ ∈ PathEdge do
14         if α = ⟨s_m, d1⟩ then
15           PathEdge = PathEdge\{e}
16       C = C\{α}
```

Our Approach: FPC

Using a fine-grained path edge garbage collector

- Data-fact-level
- Light-weight
- Efficient
- No redundancy

Implementation

- Implemented in **FlowDroid (PLDI'14)** to compare with **CleanDroid (ICSE'21)**
- In about **600 lines** of Java code
- Has been **merged** into FlowDroid
- Available at https://github.com/DongjieHe/FPC

Evaluation

- Benchmark
  - 58 apps from 2 previous papers: **ASE’19** and **CGO’21**
- Evaluate on 2 metrics
  - Memory usage
  - Analysis time
- Evaluate under varying GC intervals
  - Default: 1-second

RQ1: Memory Usage

Memory budget: 200 GB per app

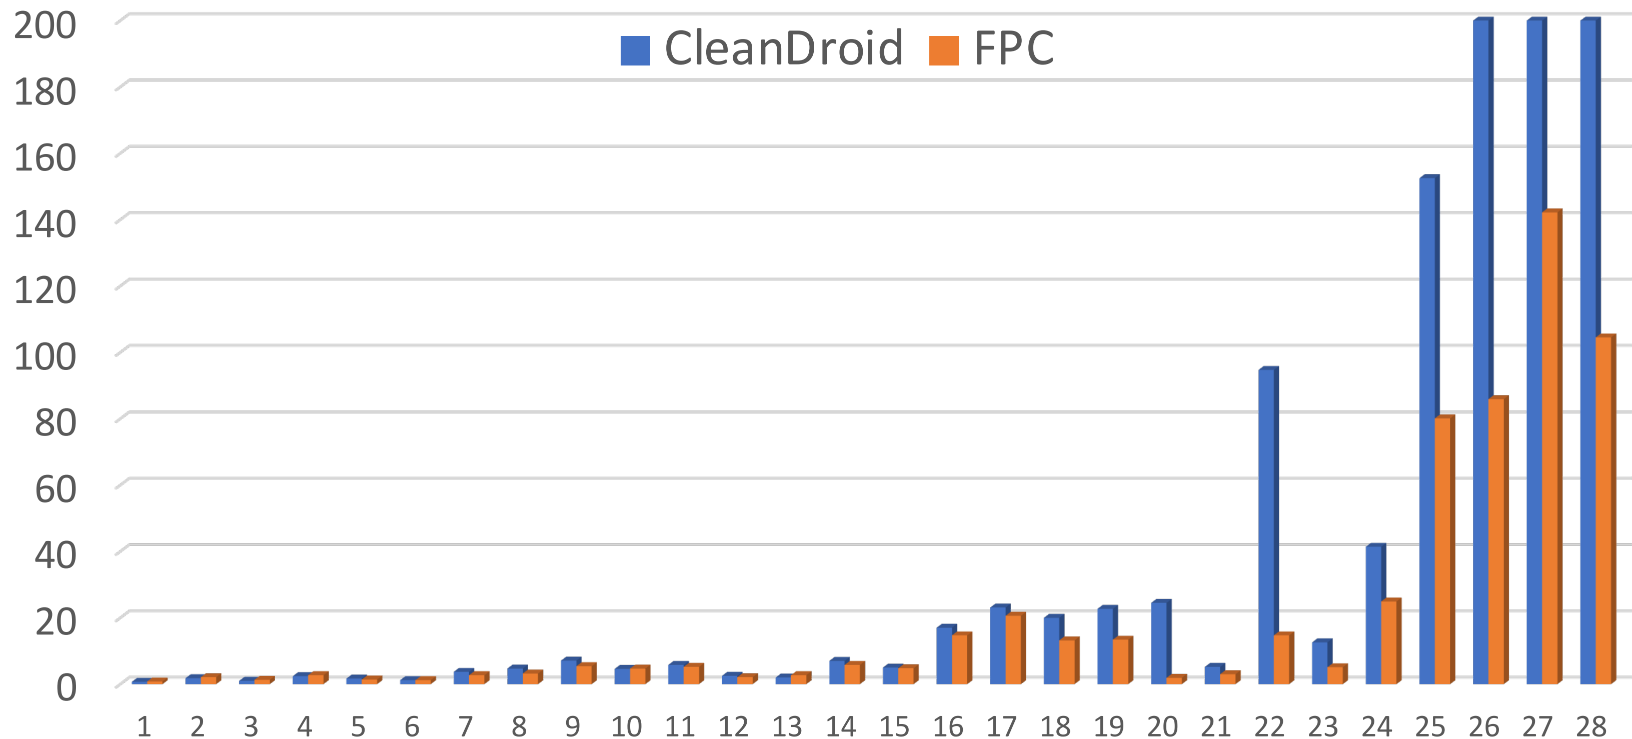

FPC can reduce the peak memory usage by 1.4× on average and can scale **3 more** apps than baseline.

RQ2: Analysis Time

Time budget: 3 hours per app

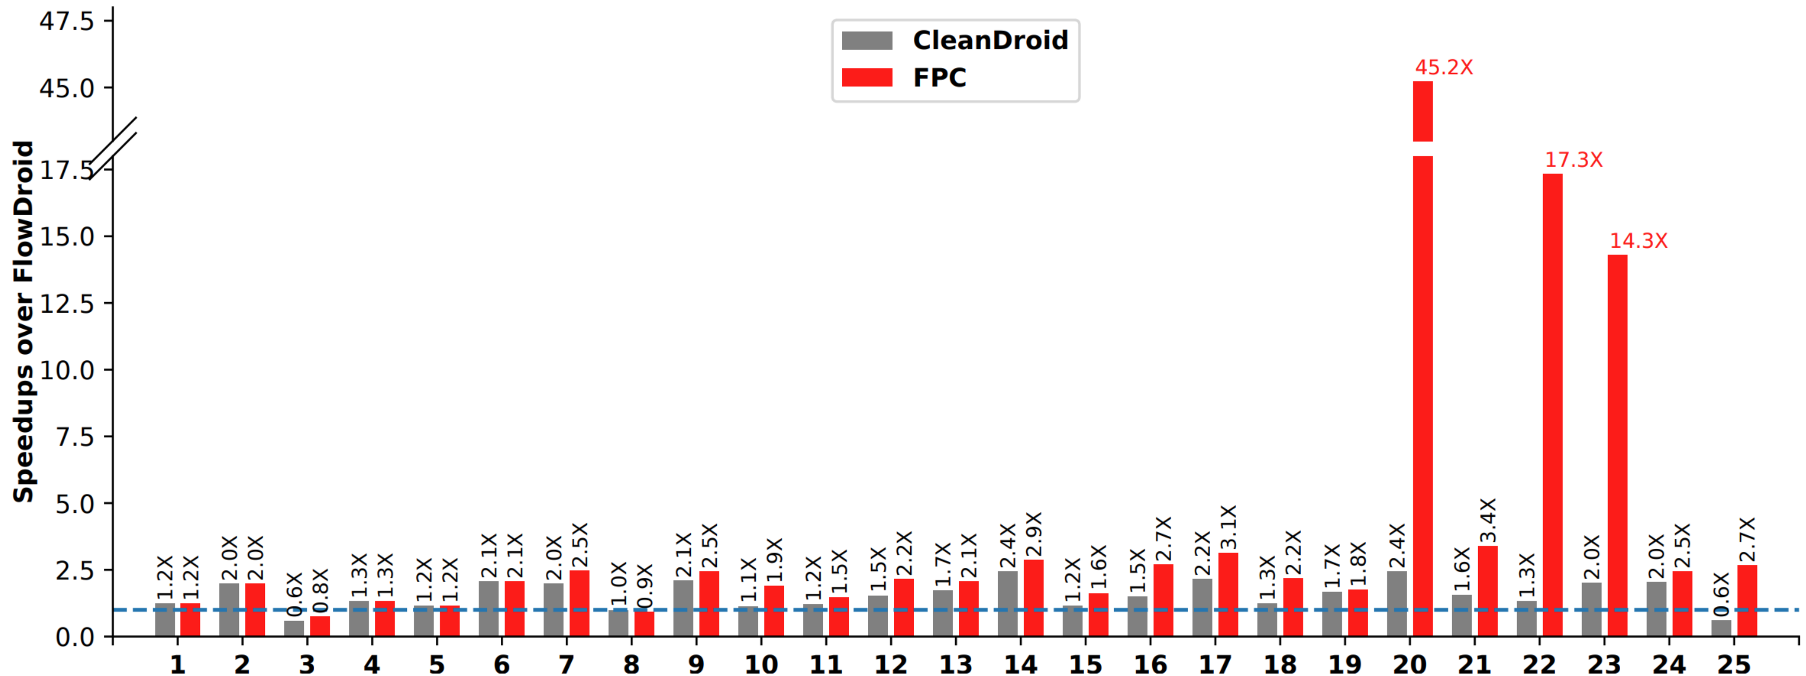

Figure 9: Comparing the speedups of FPC and CLEANDROID with FLOWDROID as the baseline.

FPC can improve both the scalability and efficiency of CleanDroid, the speedups range from 0.9× to 18.5× with an average of 1.7×.

RQ3: Varying GC Intervals

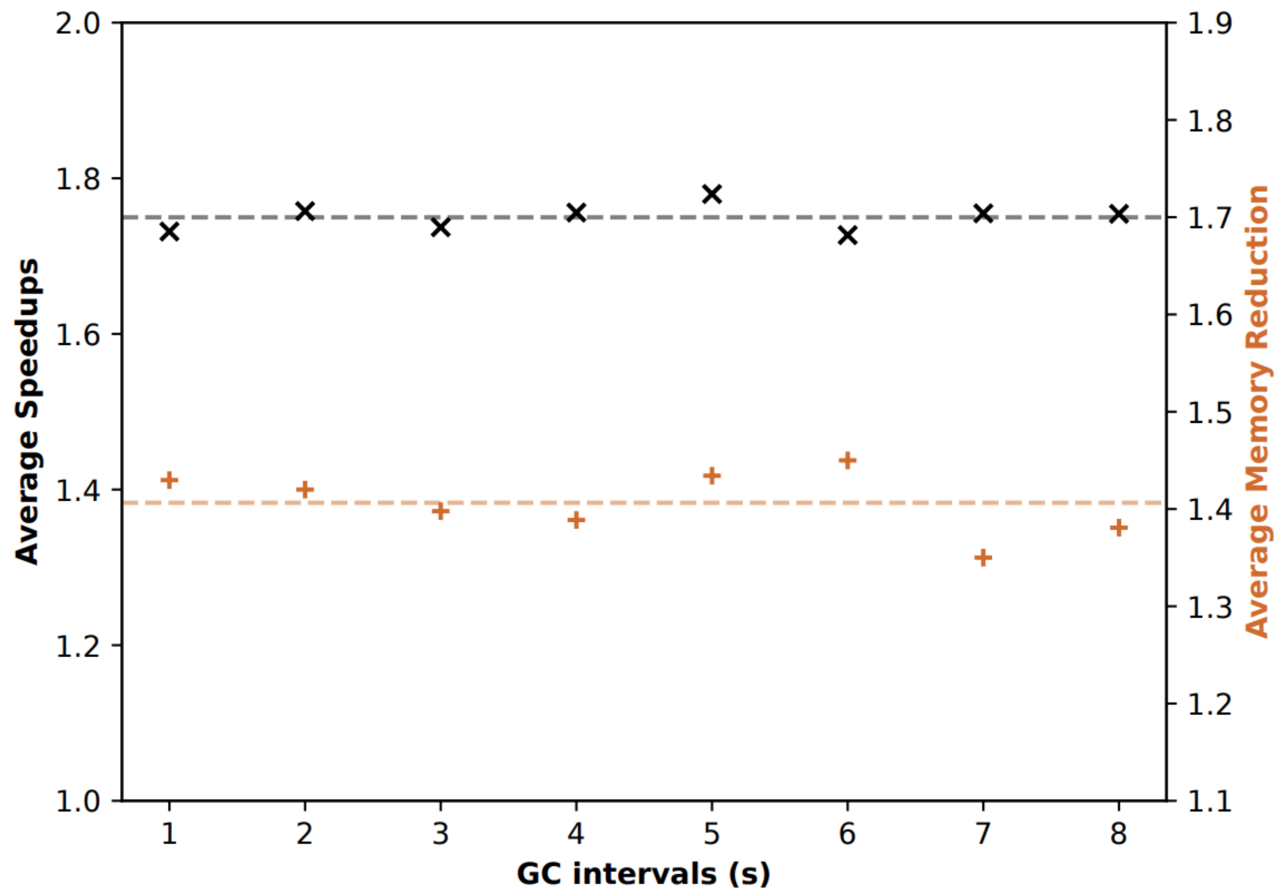

FPC has improved CleanDroid by an average of 1.40× ±0.03 for the memory usage, and 1.74× ±0.02 for the analysis time. The result obtained using 1-second GC interval is **reliable**.

Q & A

Thank you!

Dongjie He: dongjieh@cse.unsw.edu.au
Yujiang Gui: yujiang.gui@unsw.edu.au