

X-ray: Root-cause Diagnosis of Performance

Anomalies in Production Software

Mona Attariyan, Michael Chow, and Jason Flinn

University of Michigan-Ann Arbor




Troubleshooting the performance of complex production software is challenging

- Profilers and logging only reveal **what** events happened
- Users must manually infer **why** those events affected performance
- X-ray automates this task by attributing performance to specific root causes

X-ray uses performance summarization

Step 1: Attribute costs to specific instructions and system calls



Step 2: Determine root cause using taint-tracking (ConfAid)

```
If (X) {  
  // Execution depends on  
  // option1 & option2 with  
  // probabilities 0.5 & 0.2  
  write(100);  
}  
  
If (Y) {  
  // Execution depends on  
  // option2 with  
  // probability 0.2  
  write(200);  
}
```

Step 3: Sum the costs over specific root causes

option1 costs
 $0.5 * 100 \text{ bytes} = 50 \text{ bytes}$
option2 costs
 $0.2 * 100 \text{ bytes} = 20 \text{ bytes}$

option2 costs
 $0.2 * 200 \text{ bytes} = 40 \text{ bytes}$

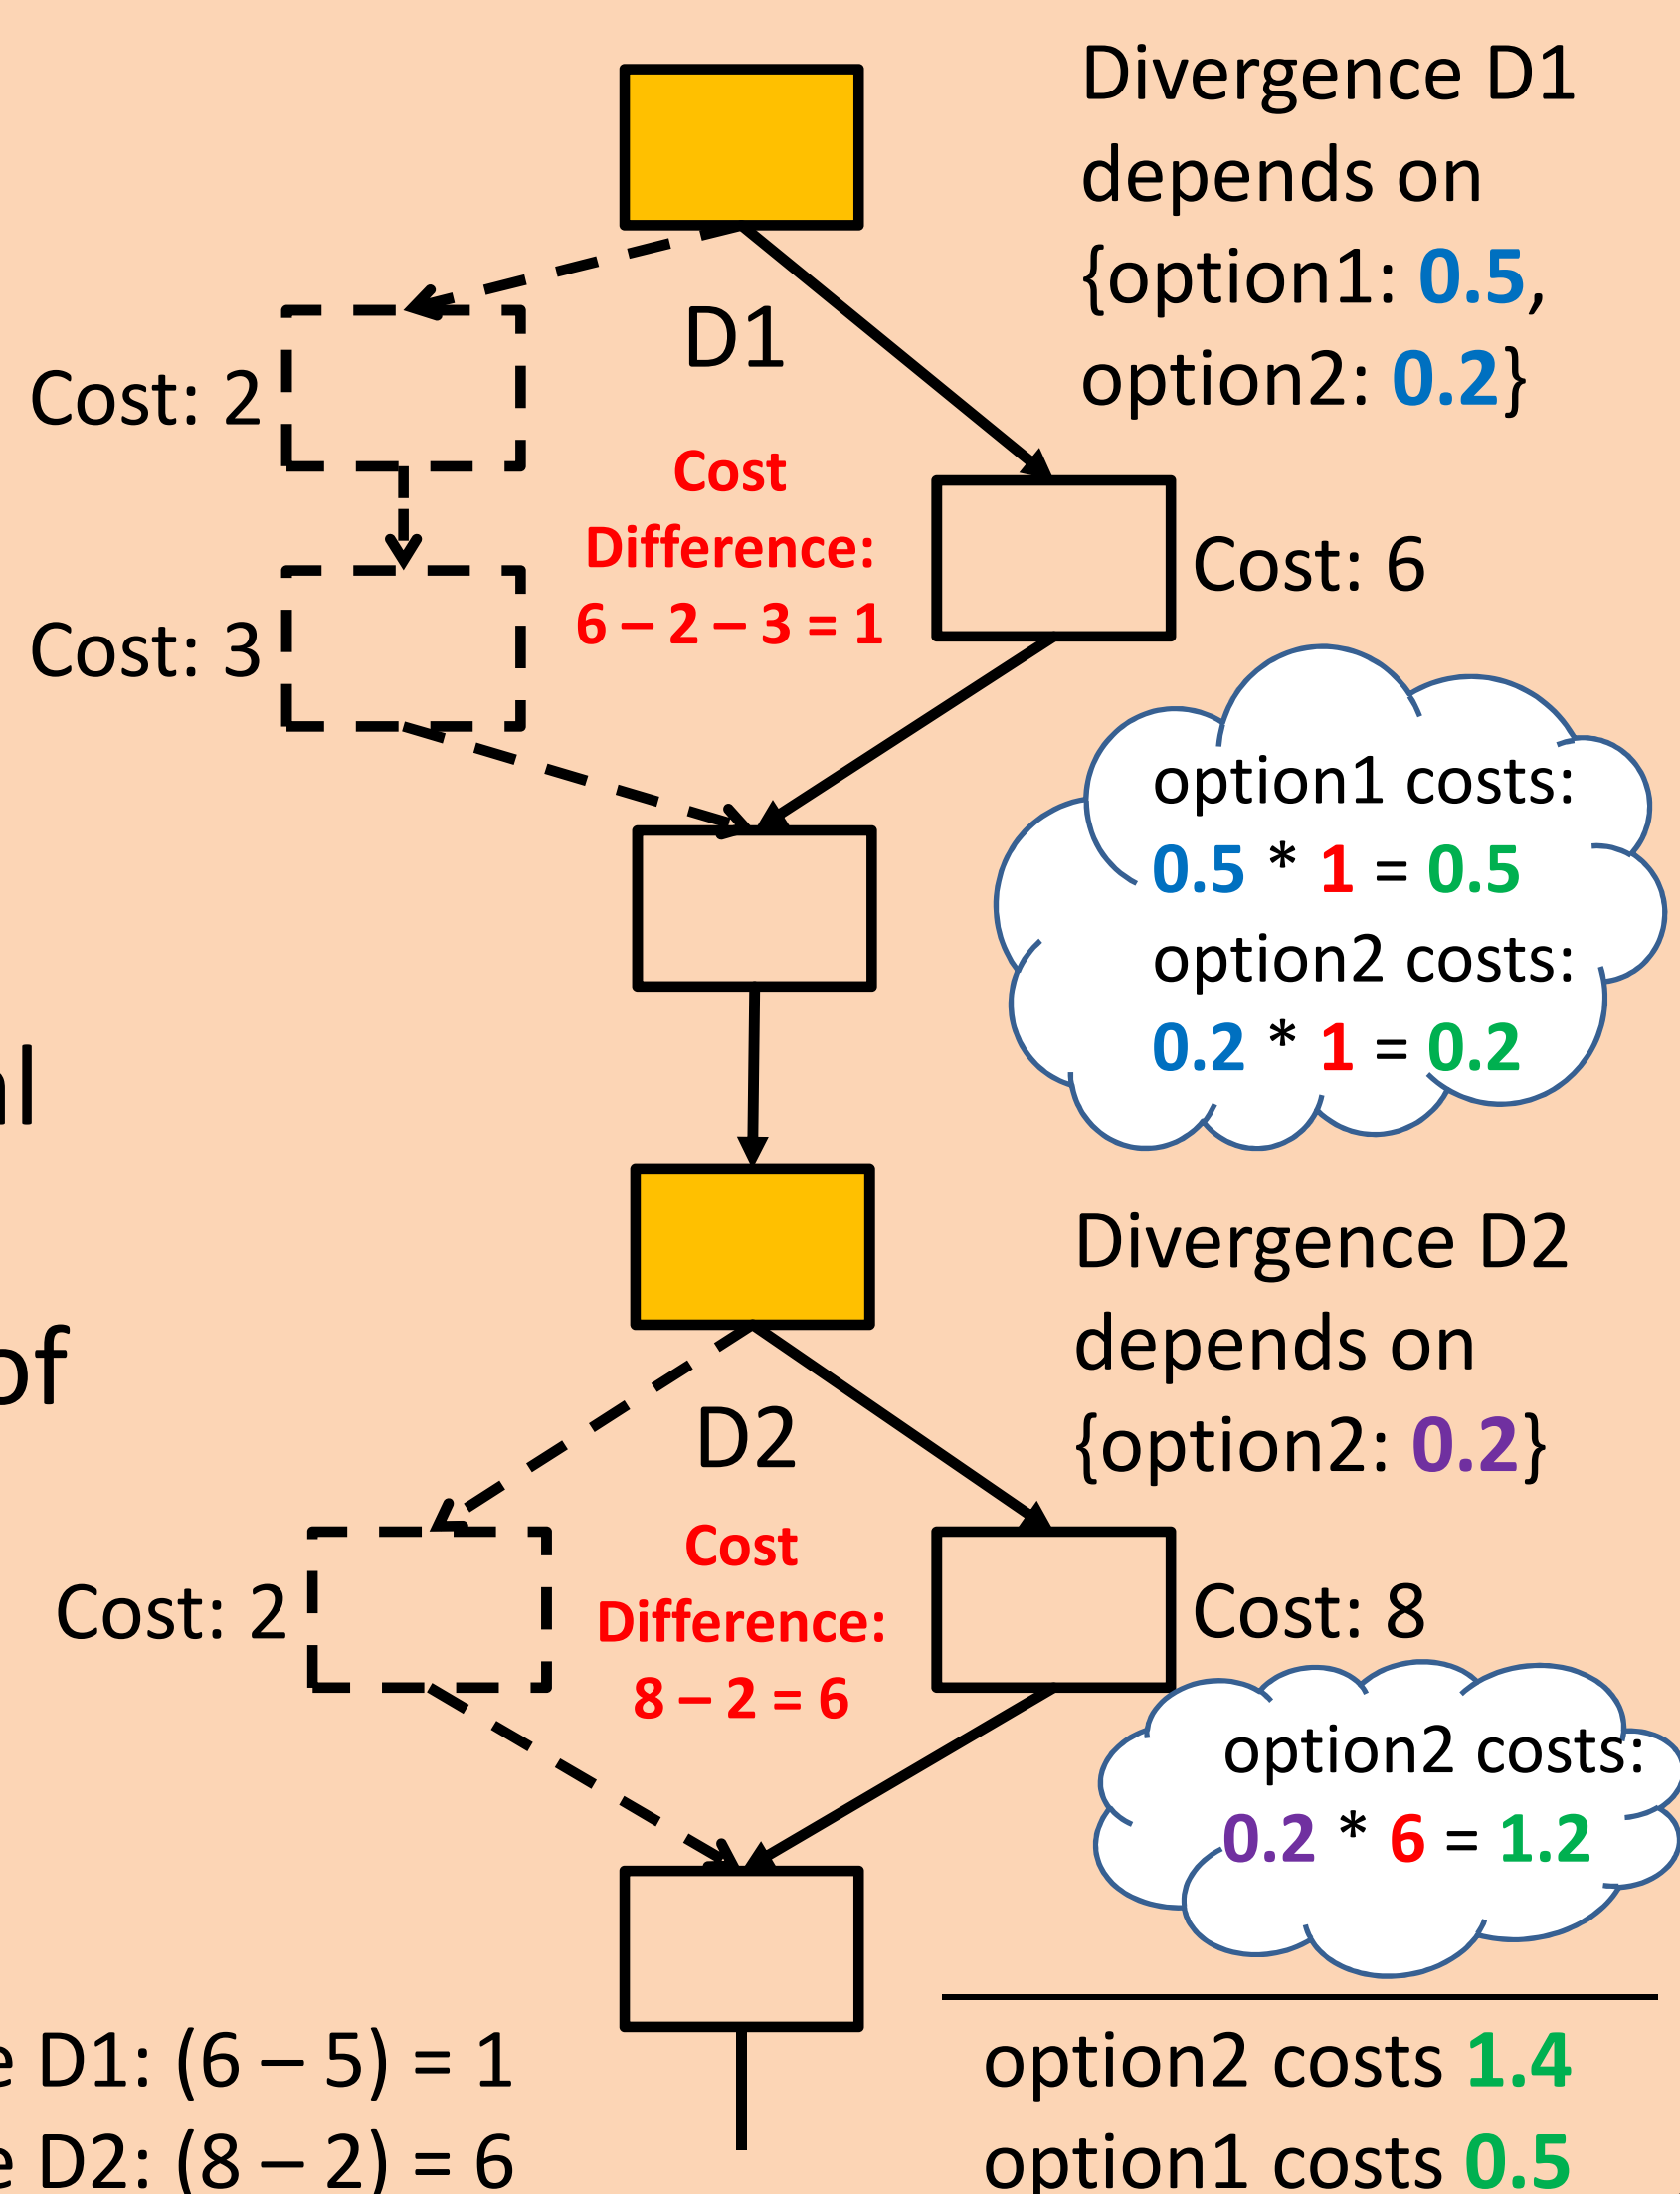
option2 costs **60 bytes**
option1 costs **50 bytes**

- Analysis is time-consuming so we use **Deterministic Replay** to do it **offline**

Differential Performance Analysis

- Identify why two requests differed in performance

- Identify where requests diverged and probabilities of root causes of the divergence
- Assign performance cost to each branch
- Calculate differential costs
- Calculate the costs of the root causes of divergences



Cost of Divergence D1: $(6 - 5) = 1$
Cost of Divergence D2: $(8 - 2) = 6$

Results

- Analyzed performance issues in Apache, Postfix, and PostgreSQL
- In 12 of 14 cases, X-ray identifies the correct root cause as the largest contributor
- Online overhead of 1 – 7% due to deterministic replay

Future Work

- Simultaneously compare thousands of requests to look for anomalies
- Better understand how workload distribution affects performance
- Graphical tools for visualizing performance issues