

On Testing Non-functional Software Properties

Sudipta Chattopadhyay
Linköping University, Sweden

Joint work with Abhijeet Banerjee, Lee Kee Chong and Abhik Roychoudhury
National University of Singapore

Ahmed Rezine and Ke Jiang
Linköping University, Sweden



Context

Developer



Write efficient software

Programming abstractions

Tools and techniques

Desktops, handheld devices etc.

Overview

Tools and Techniques to Write Efficient Software

**Software Testing
for Non-functional
Properties**



**Inefficient
Code Patterns**



**Coding Guidelines
for Building Efficient
Software**

Desktop machines, handheld devices etc.

Overview

Tools and Techniques to Write Efficient Software

Performance Testing
Energy Testing

**Inefficient
Code Patterns**

**Coding Guidelines
for Building Efficient
Software**

Desktop machines, handheld devices etc.

**Memory
Performance**
Android Devices

Overview

Tools and Techniques to Write Efficient Software

Performance Testing

**Performance-inefficient
Code Patterns**

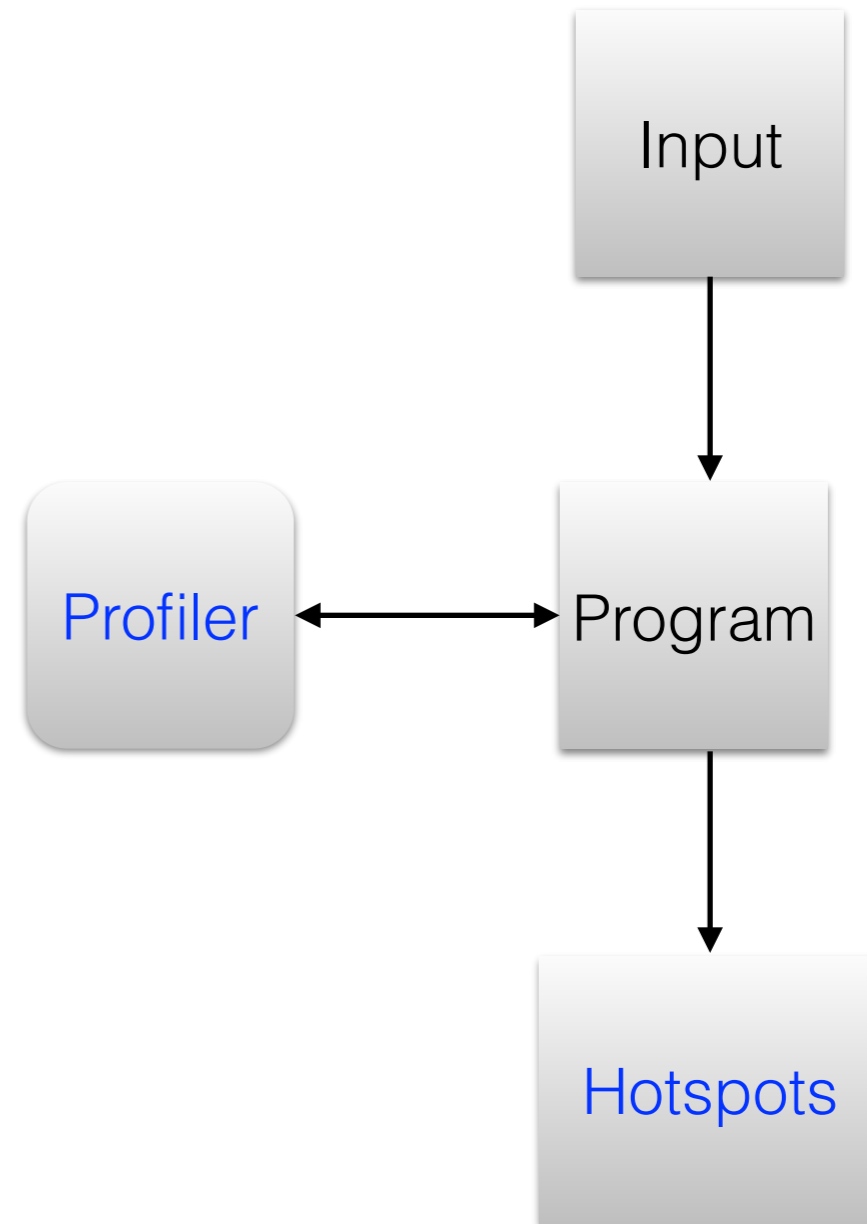
**Performance-aware
Coding Guidelines**

Desktop machines, handheld devices etc.

**Memory
Performance**

State-of-the-art in Detecting
Performance Loss

Program profiling

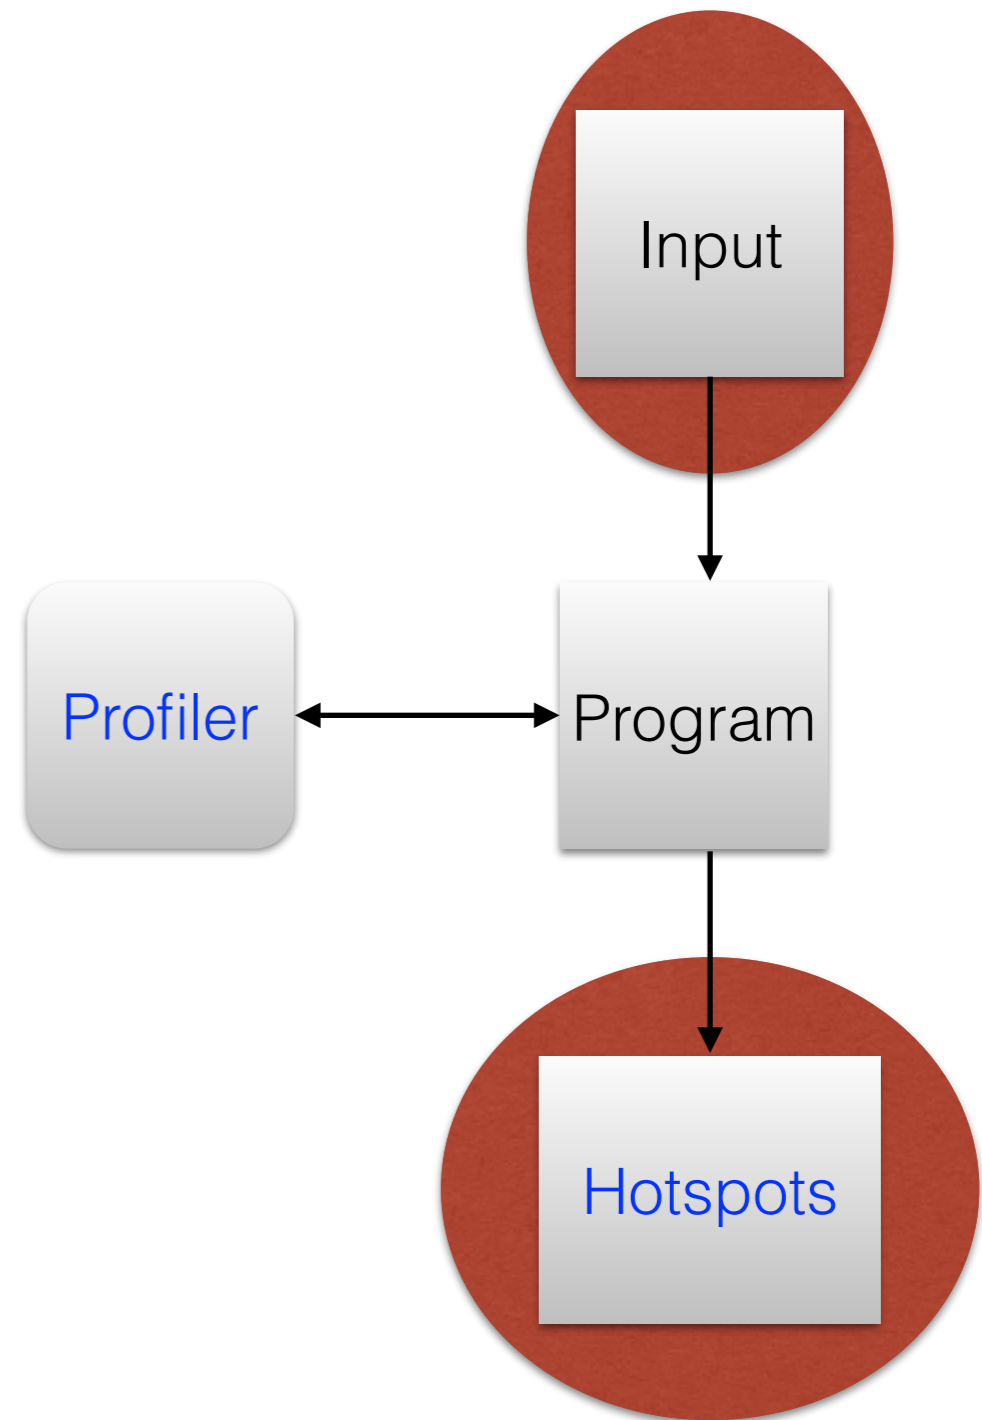


State-of-the-art in Detecting
Performance Loss

Program profiling

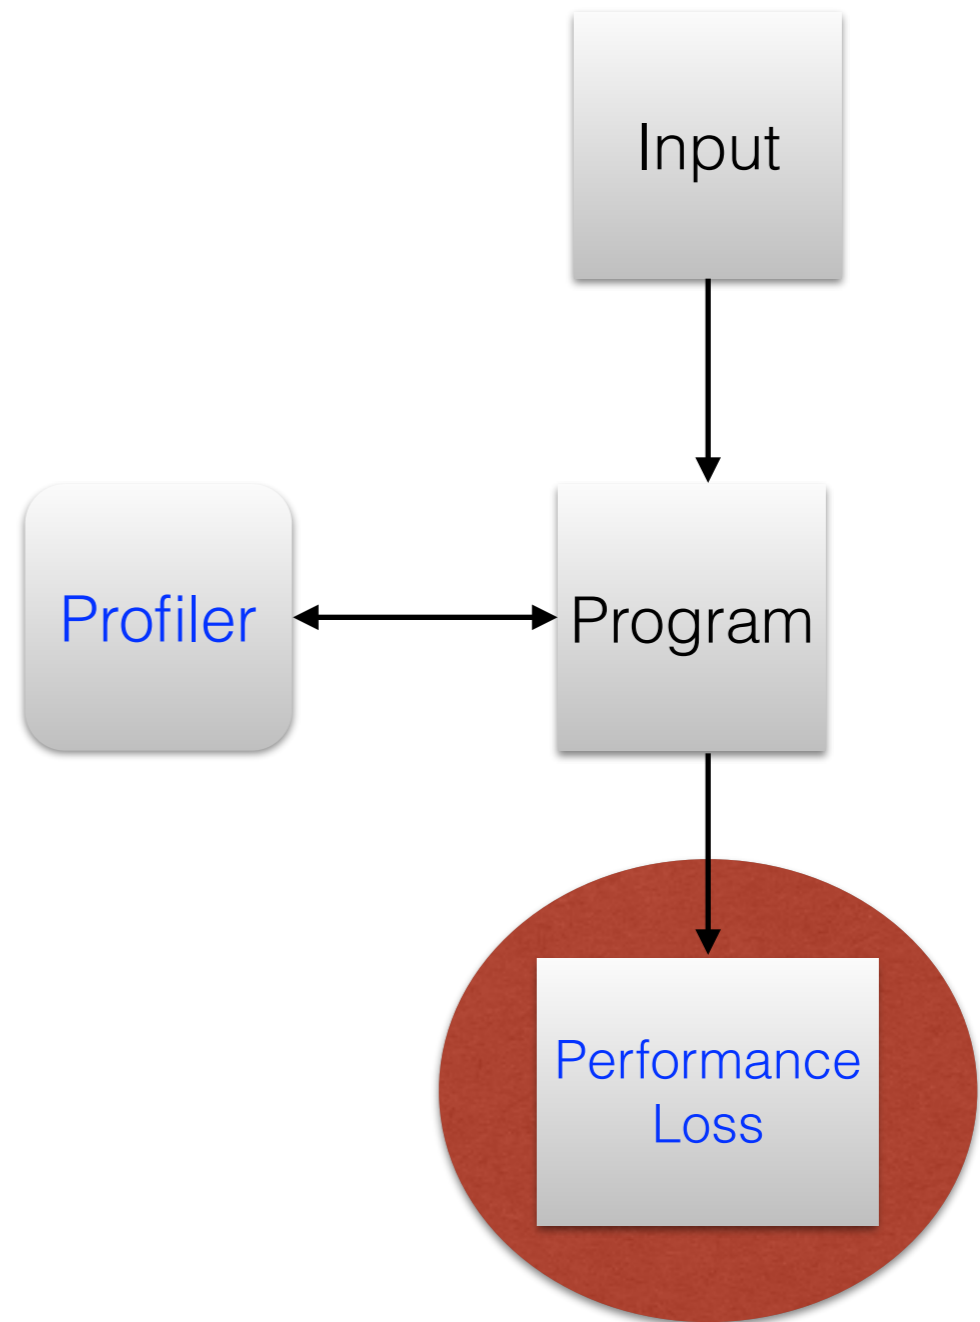
Program inputs that expose performance loss

Detecting performance loss




State-of-the-art in Detecting
Performance Loss

Detecting performance loss



Memory Performance



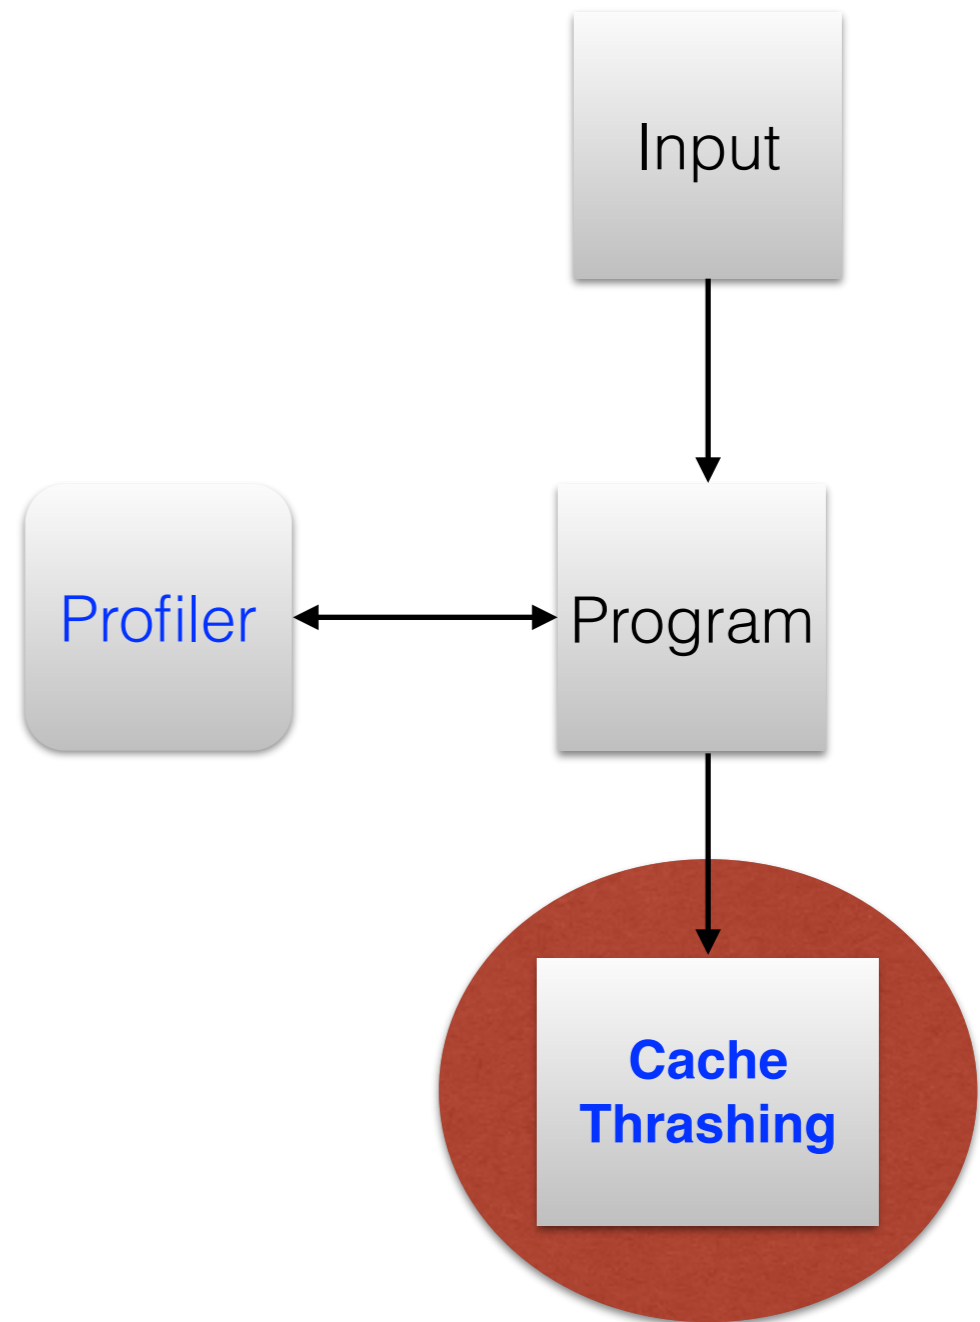
Typically, DRAM is several magnitudes slower than caches

Performance Loss

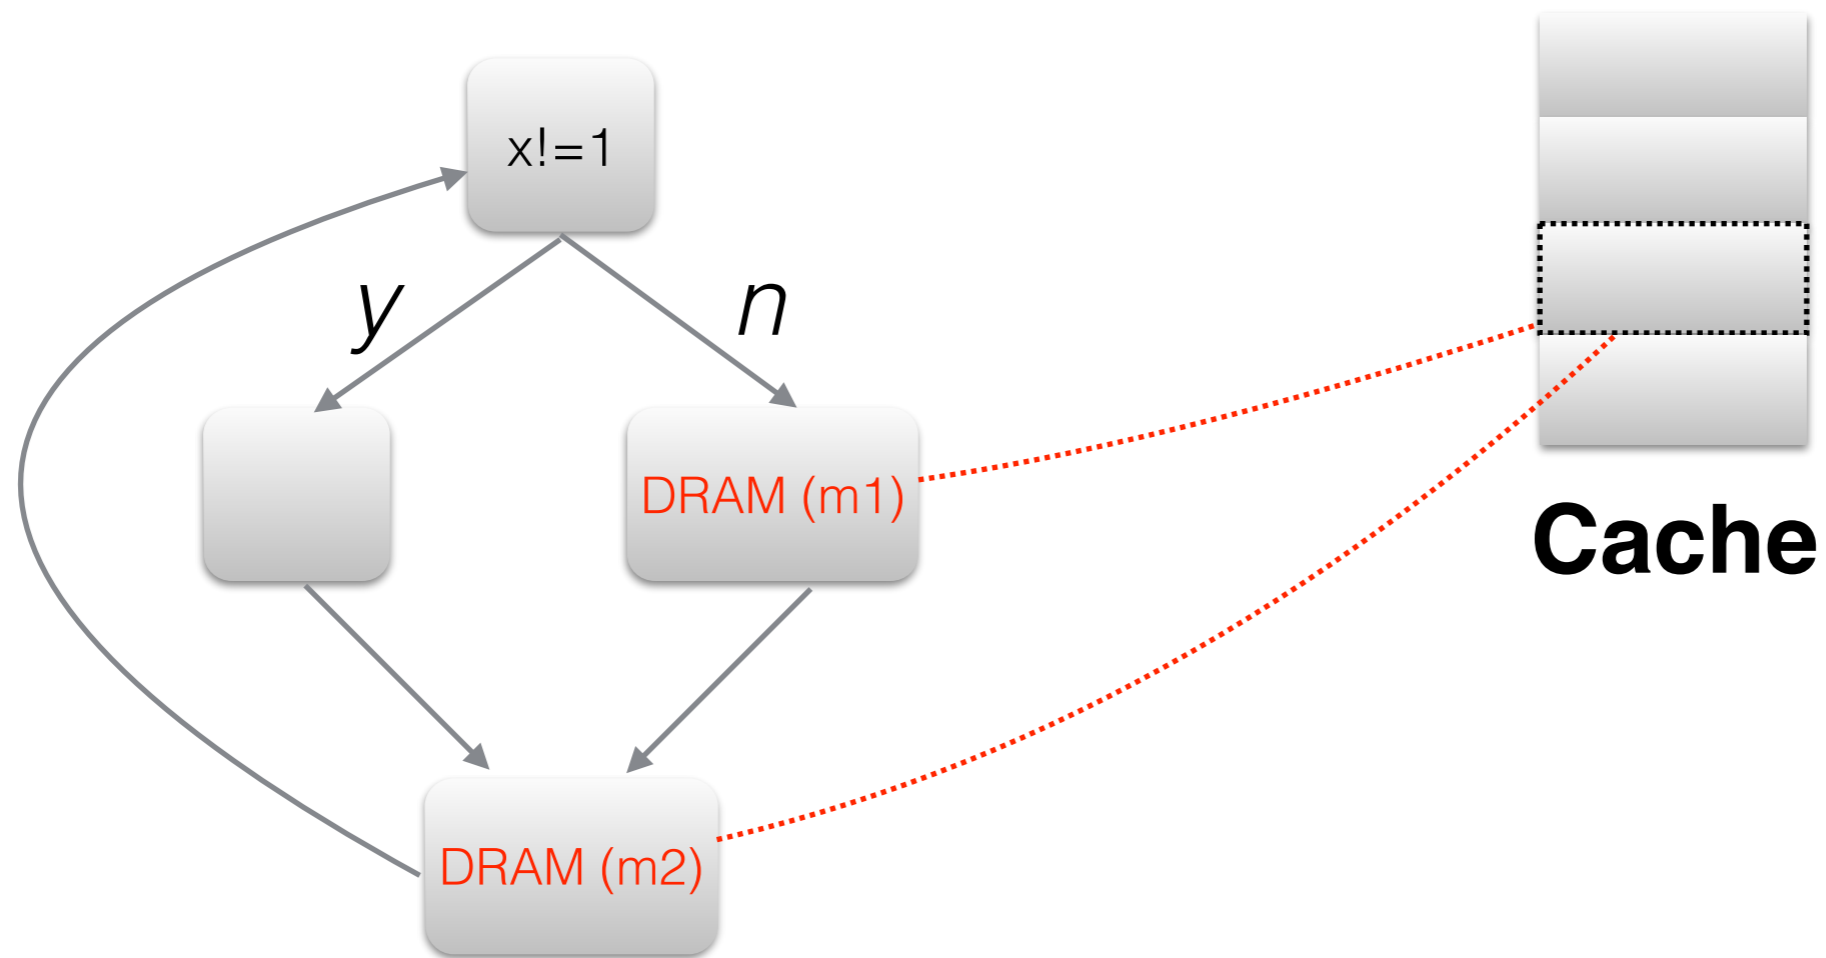
- Memory Performance
 - How do we define?
 - Many cache misses (How many is bad enough?)
- Our approach
 - Detect ***cache thrashing***

State-of-the-art in Detecting
Performance Loss

Detecting performance loss

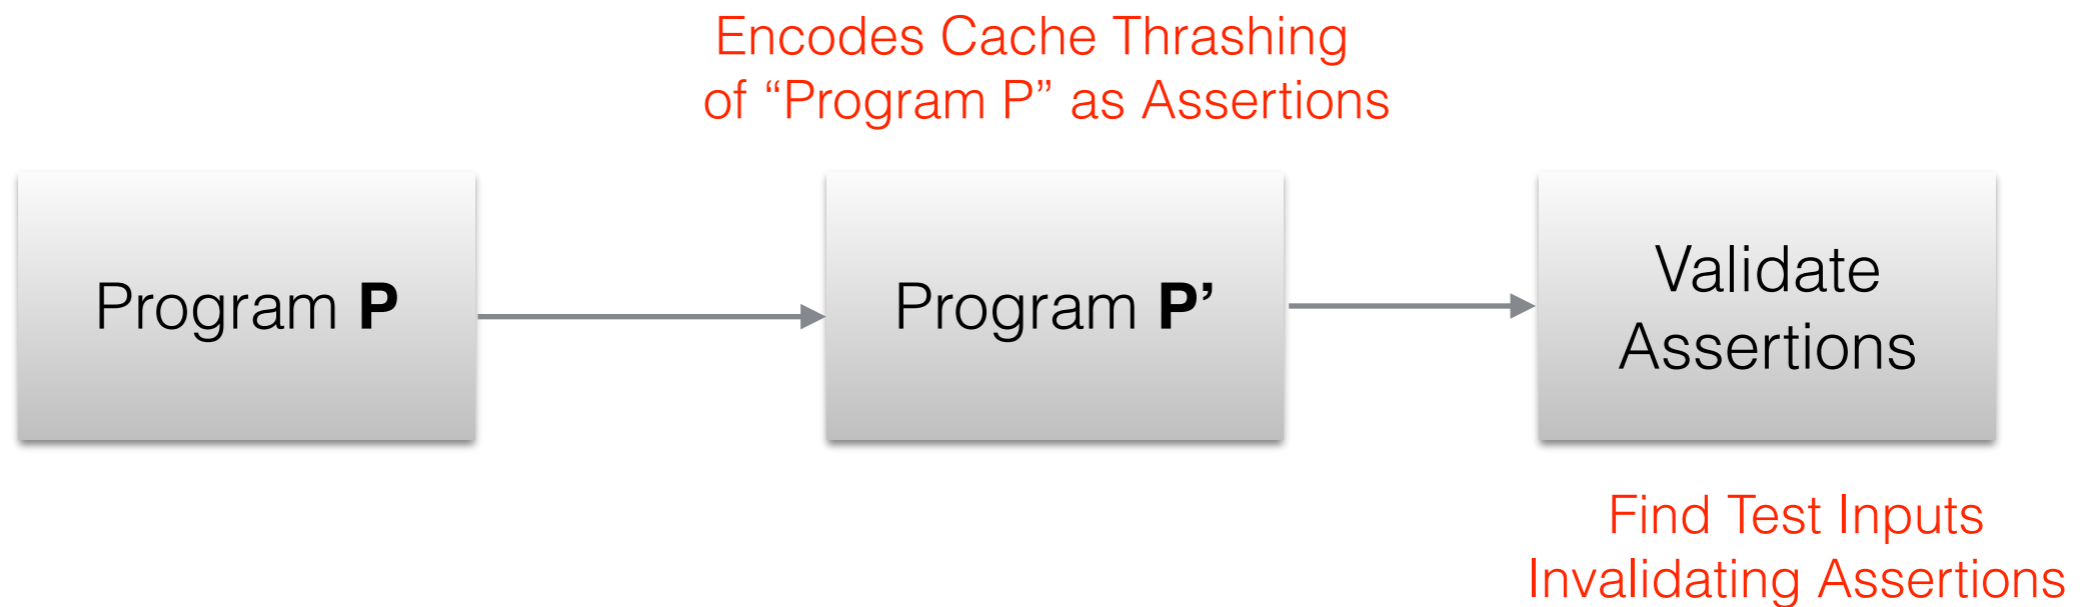


Cache Thrashing Scenario

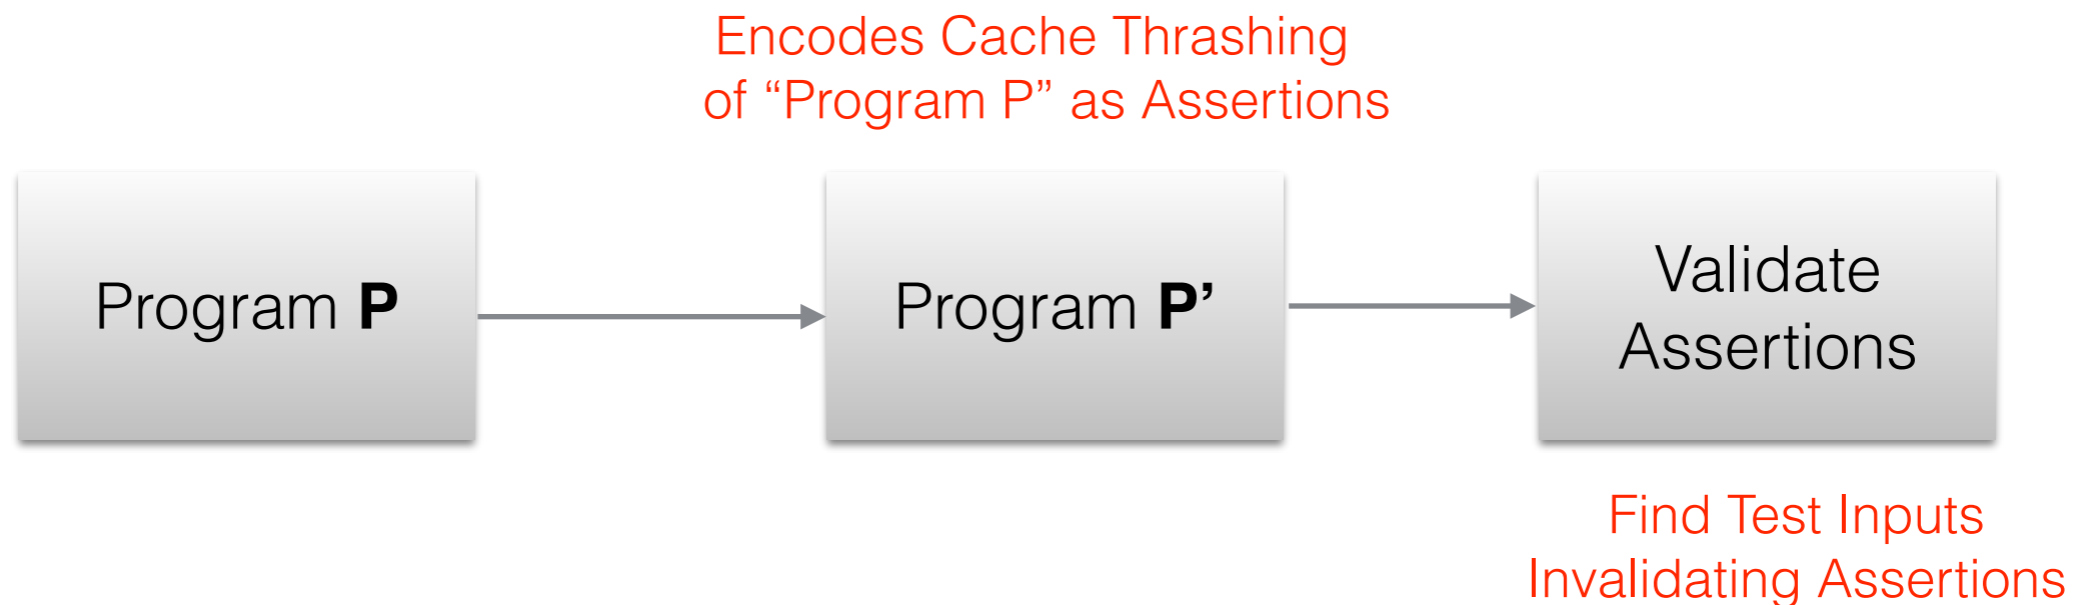


m1 replaces m2 from the cache and vice versa

Test Generation




Test Generation



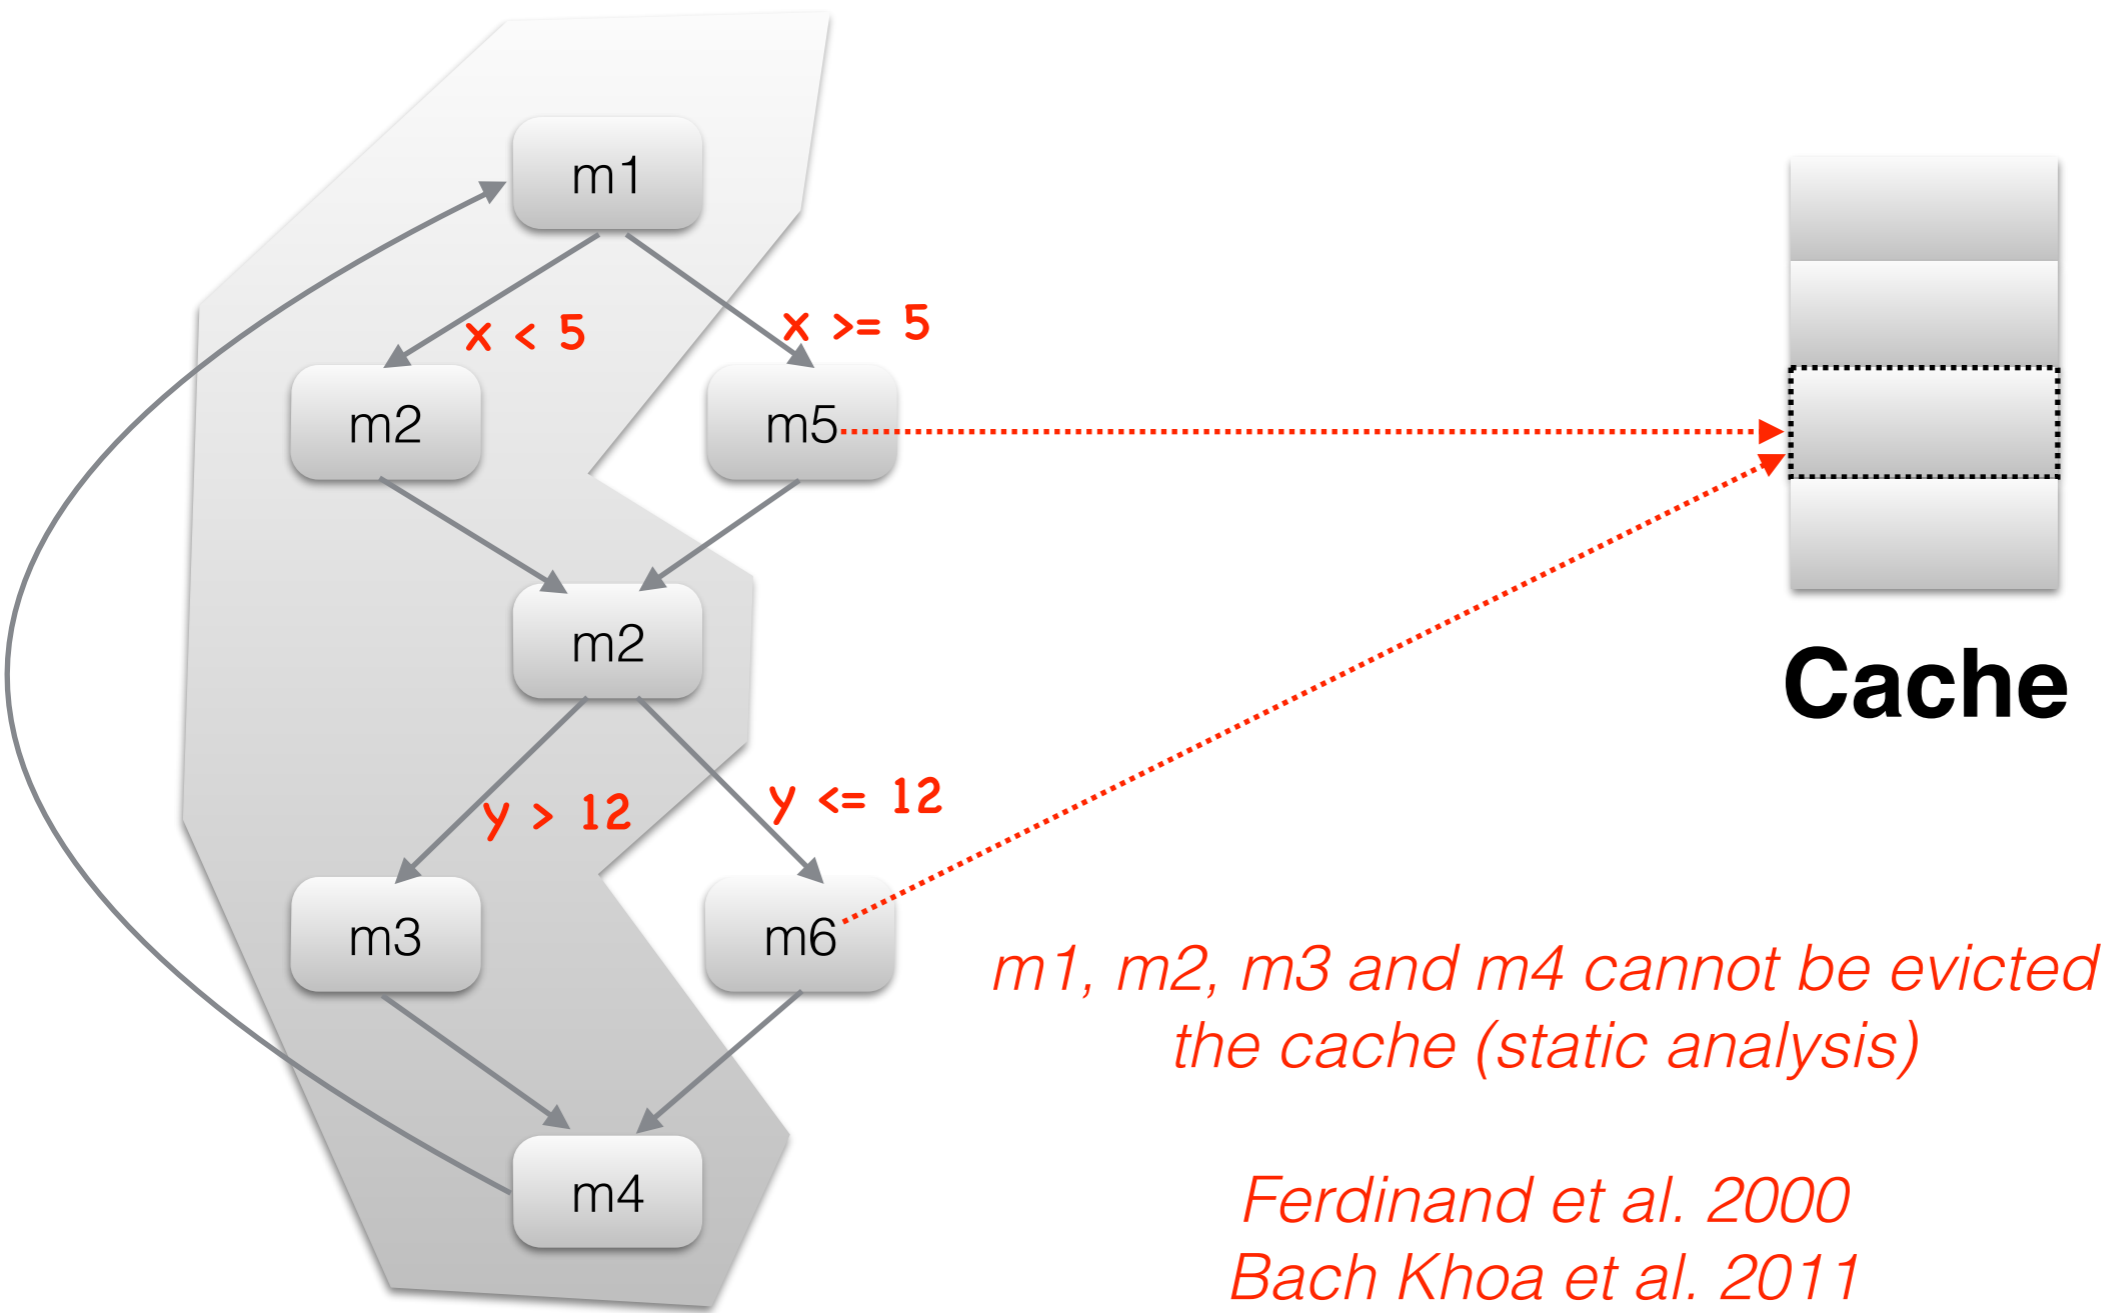
In other words, we reduce **memory performance testing** into an equivalent **functionality testing** problem

Encoding Cache Thrashing Scenario



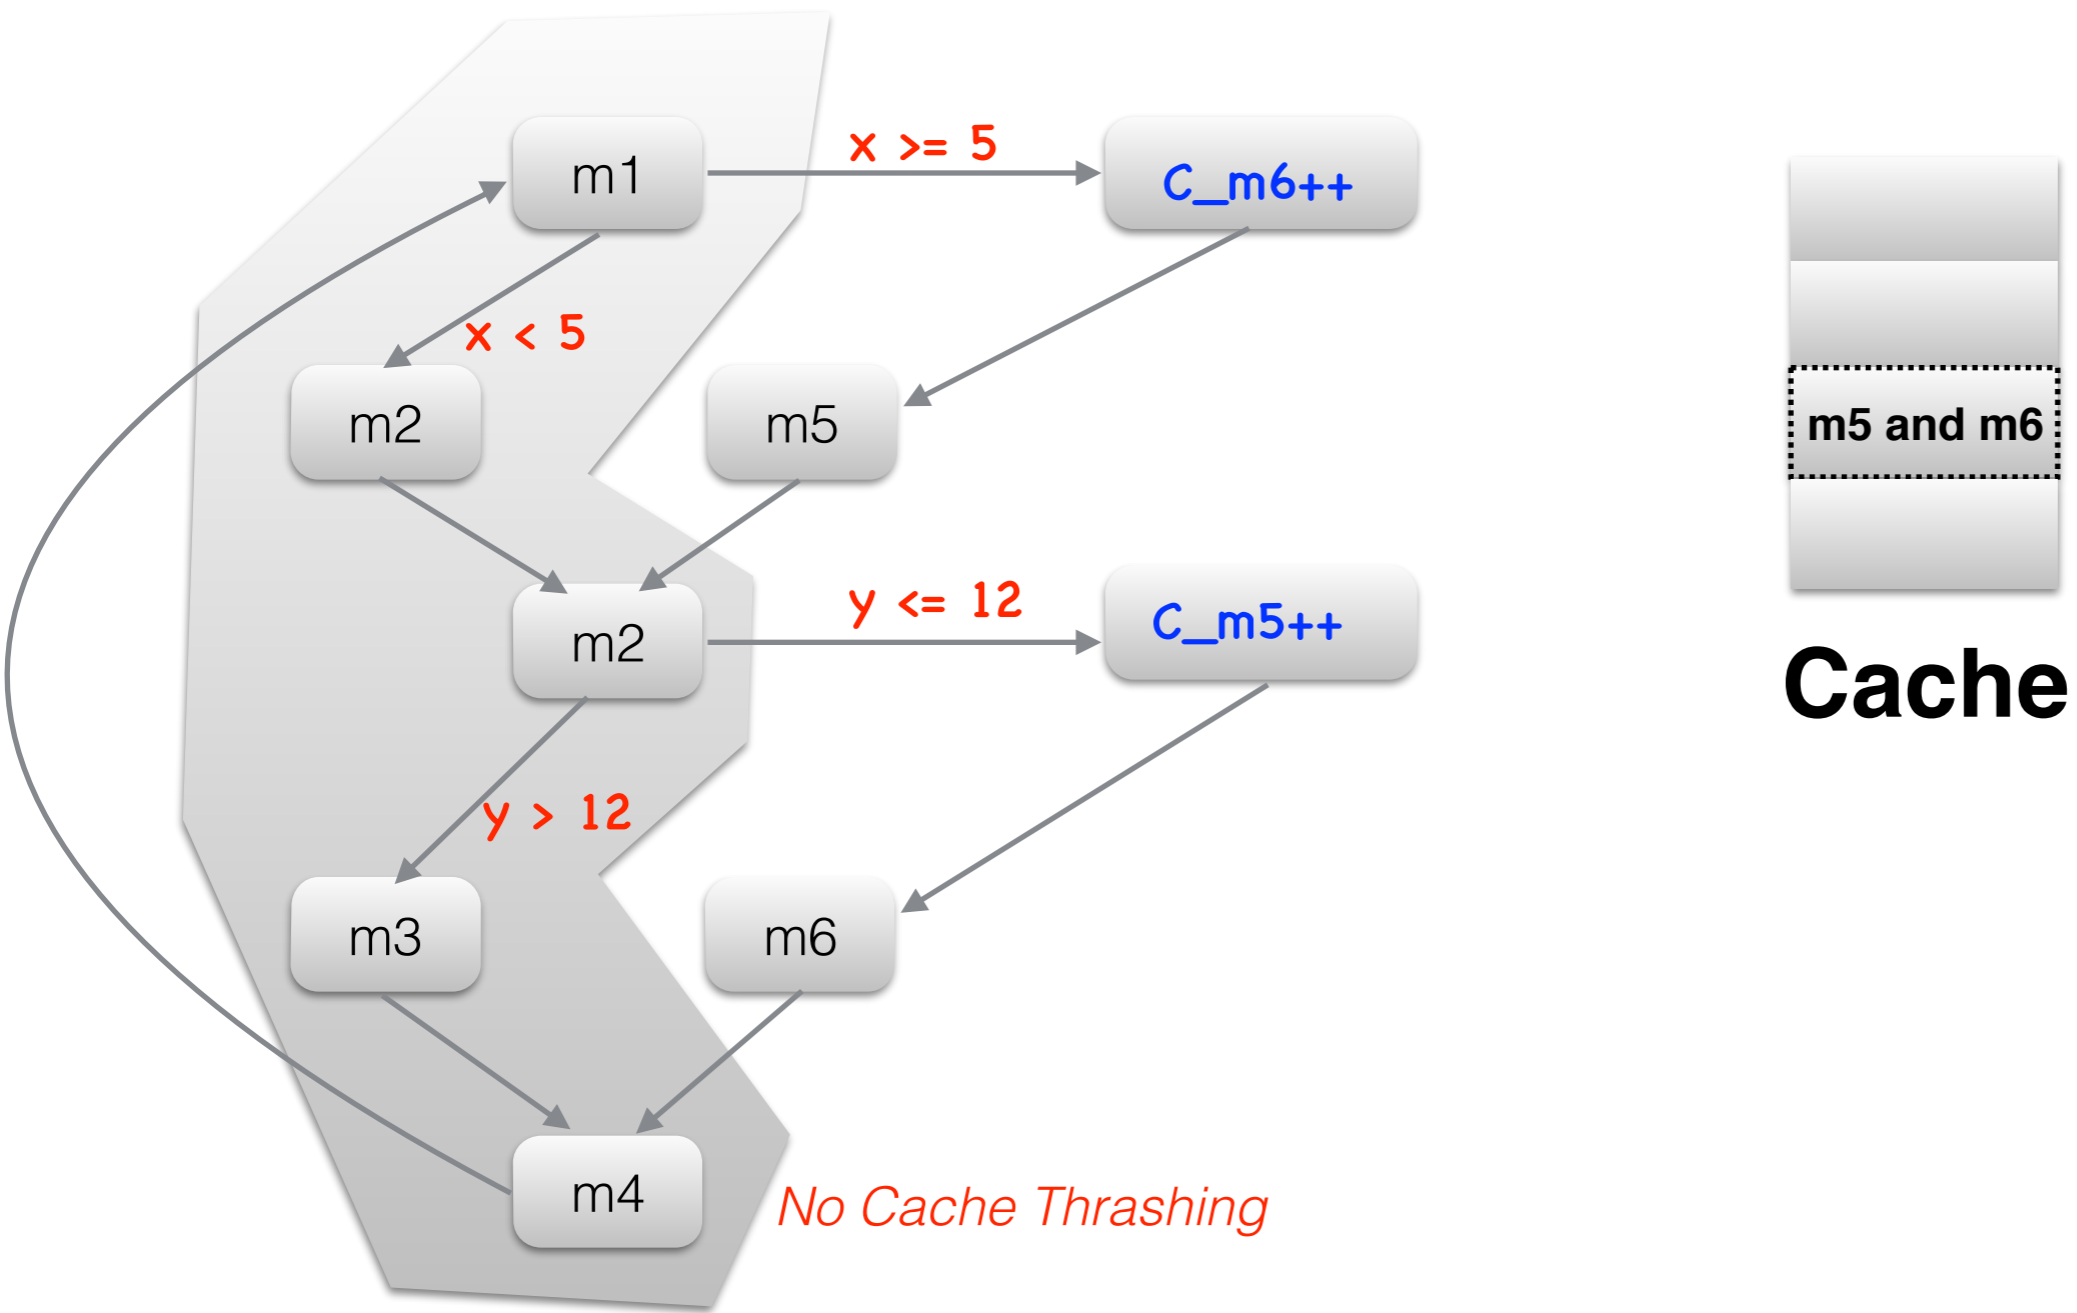
m5 and m6 map to the same cache set

Encoding Cache Thrashing Scenario



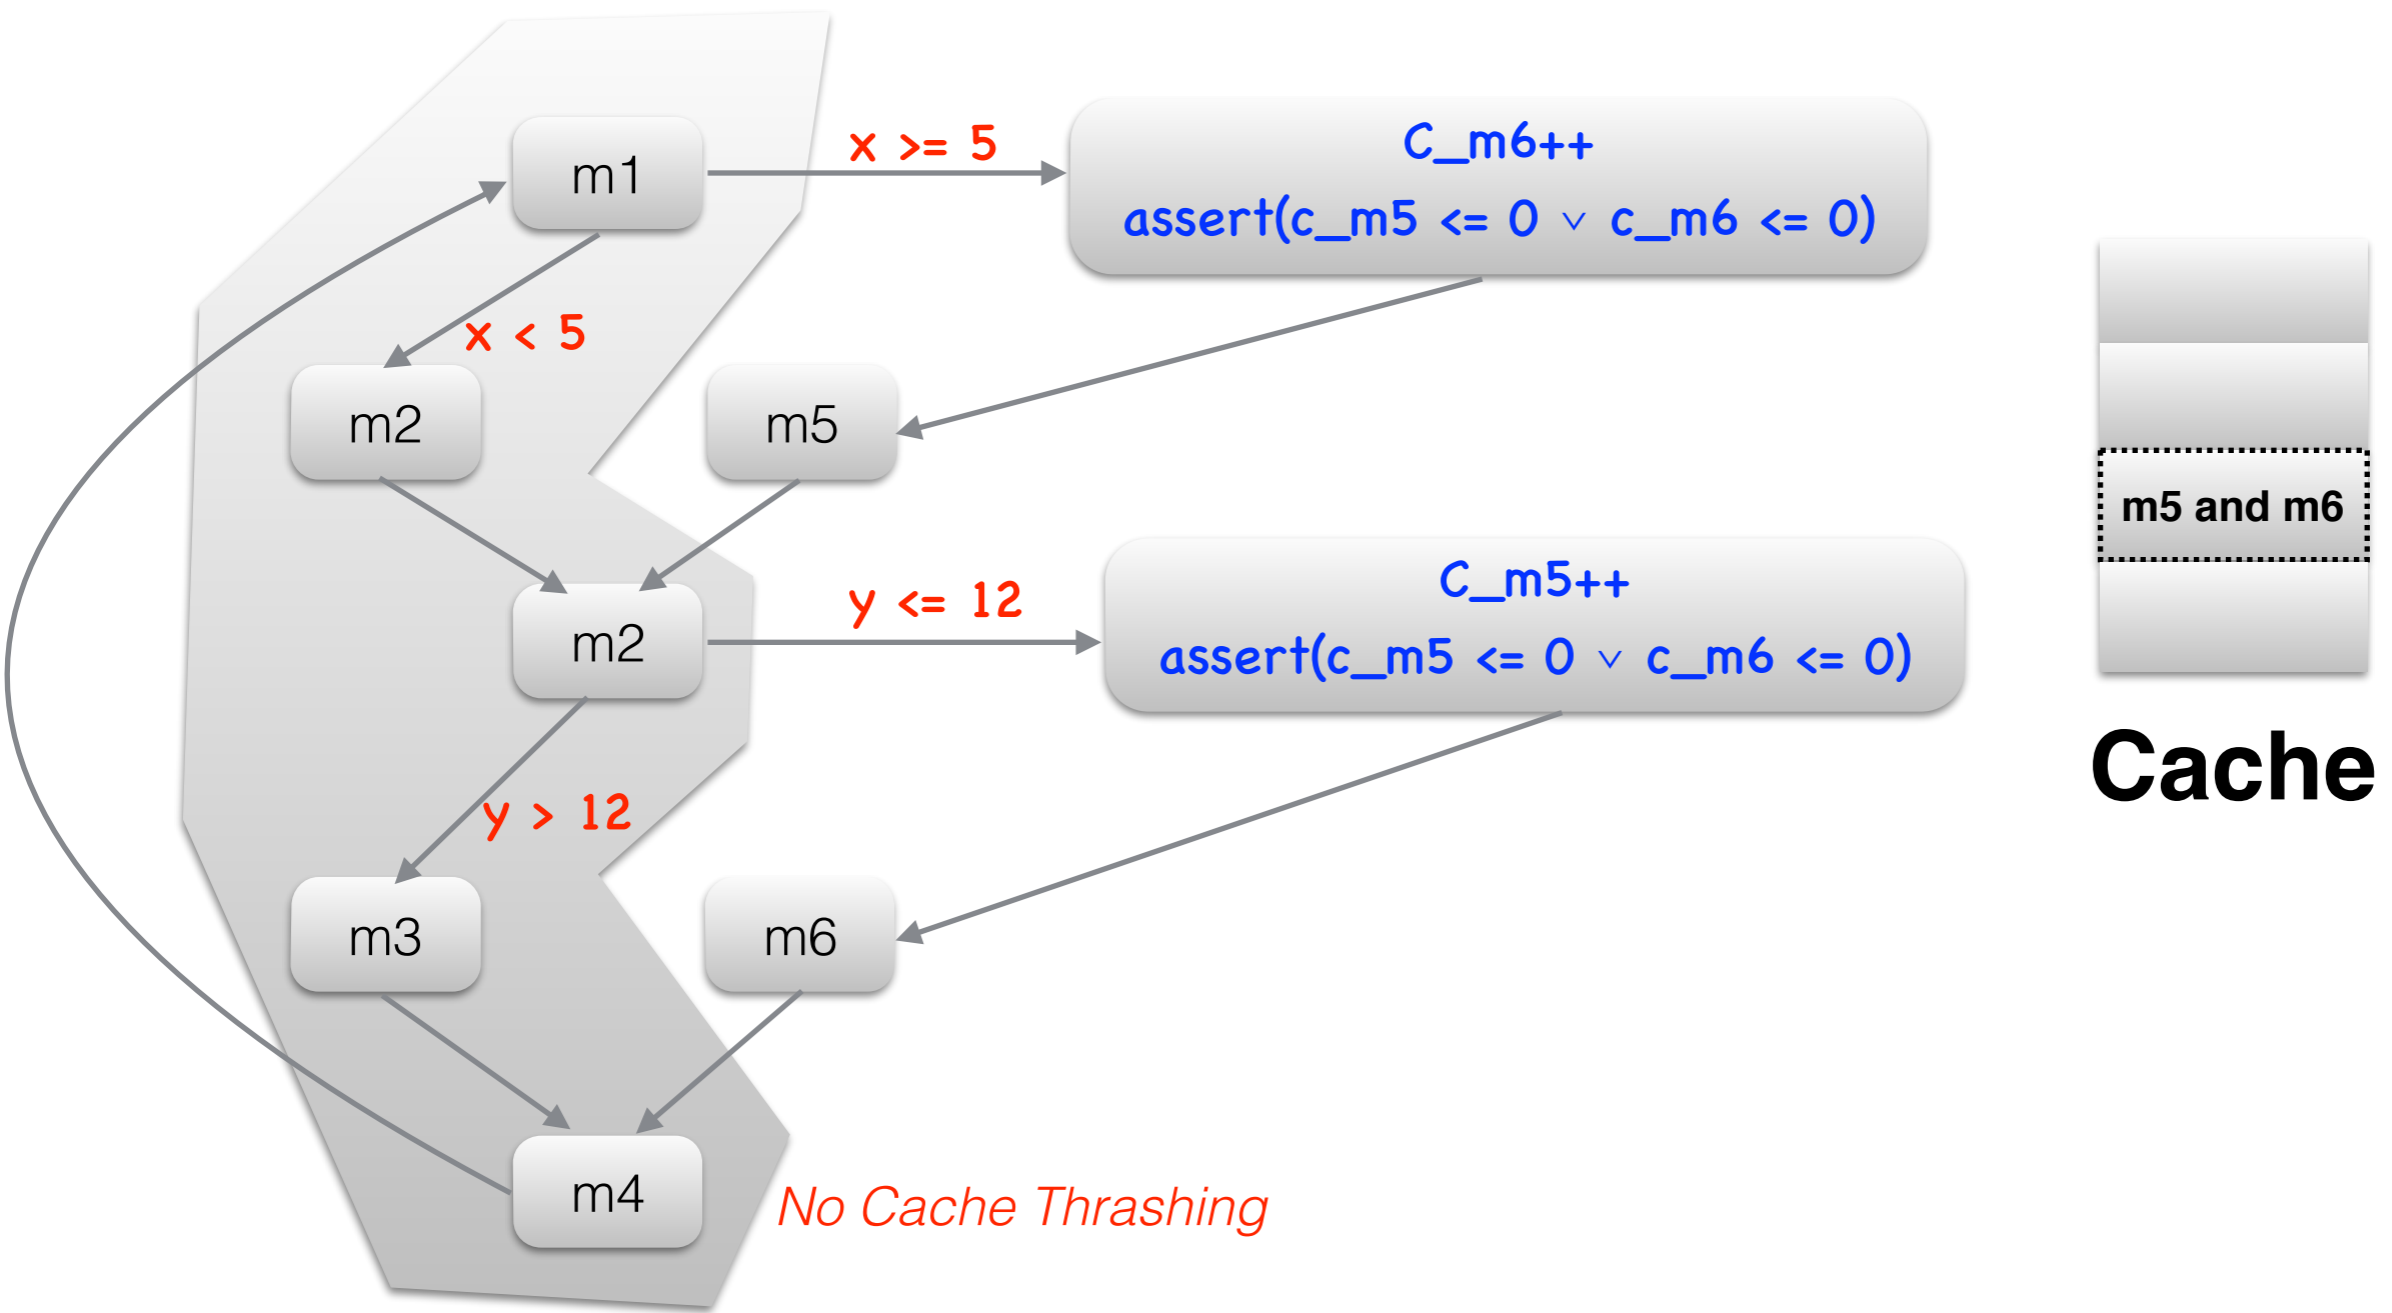
m5 and m6 map to the same cache set

Encoding Cache Thrashing Scenario




m5 and m6 map to the same cache set

Encoding Cache Thrashing Scenario



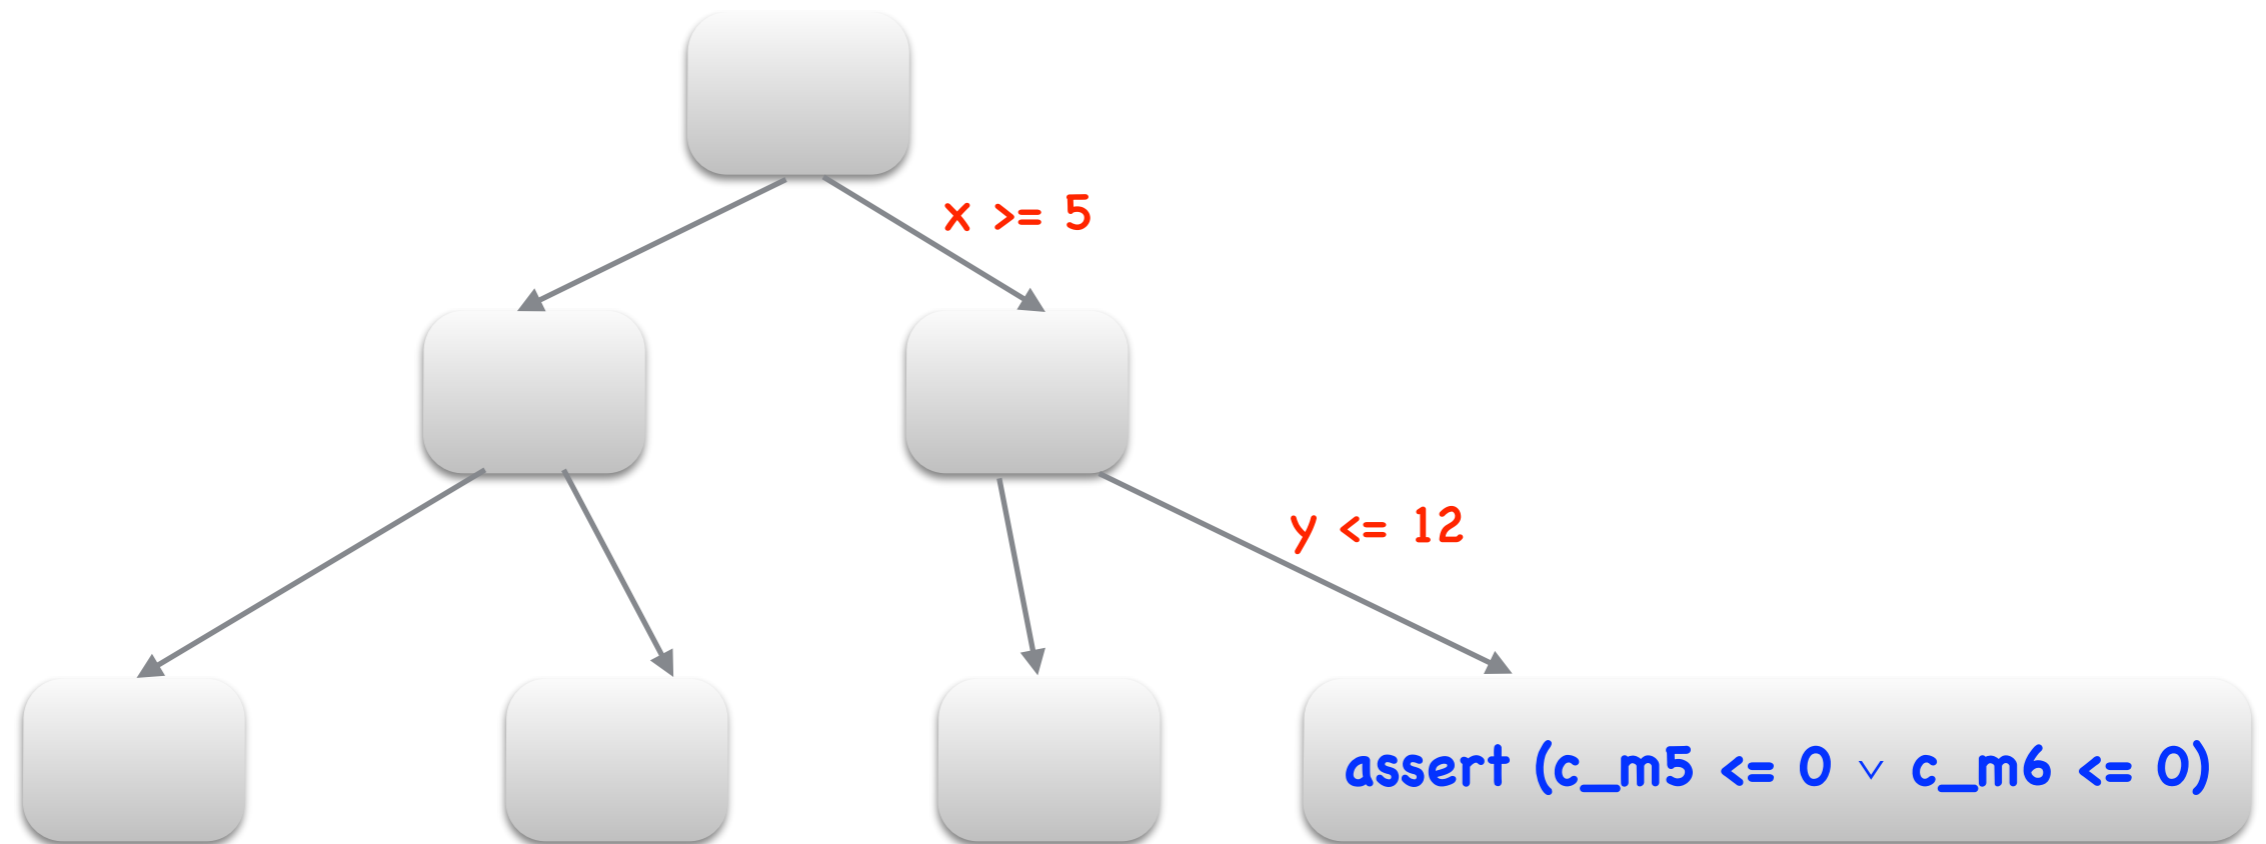
m5 and m6 map to the same cache set

Encoding Cache Thrashing Scenario



Test Generation is Performed on the modified program

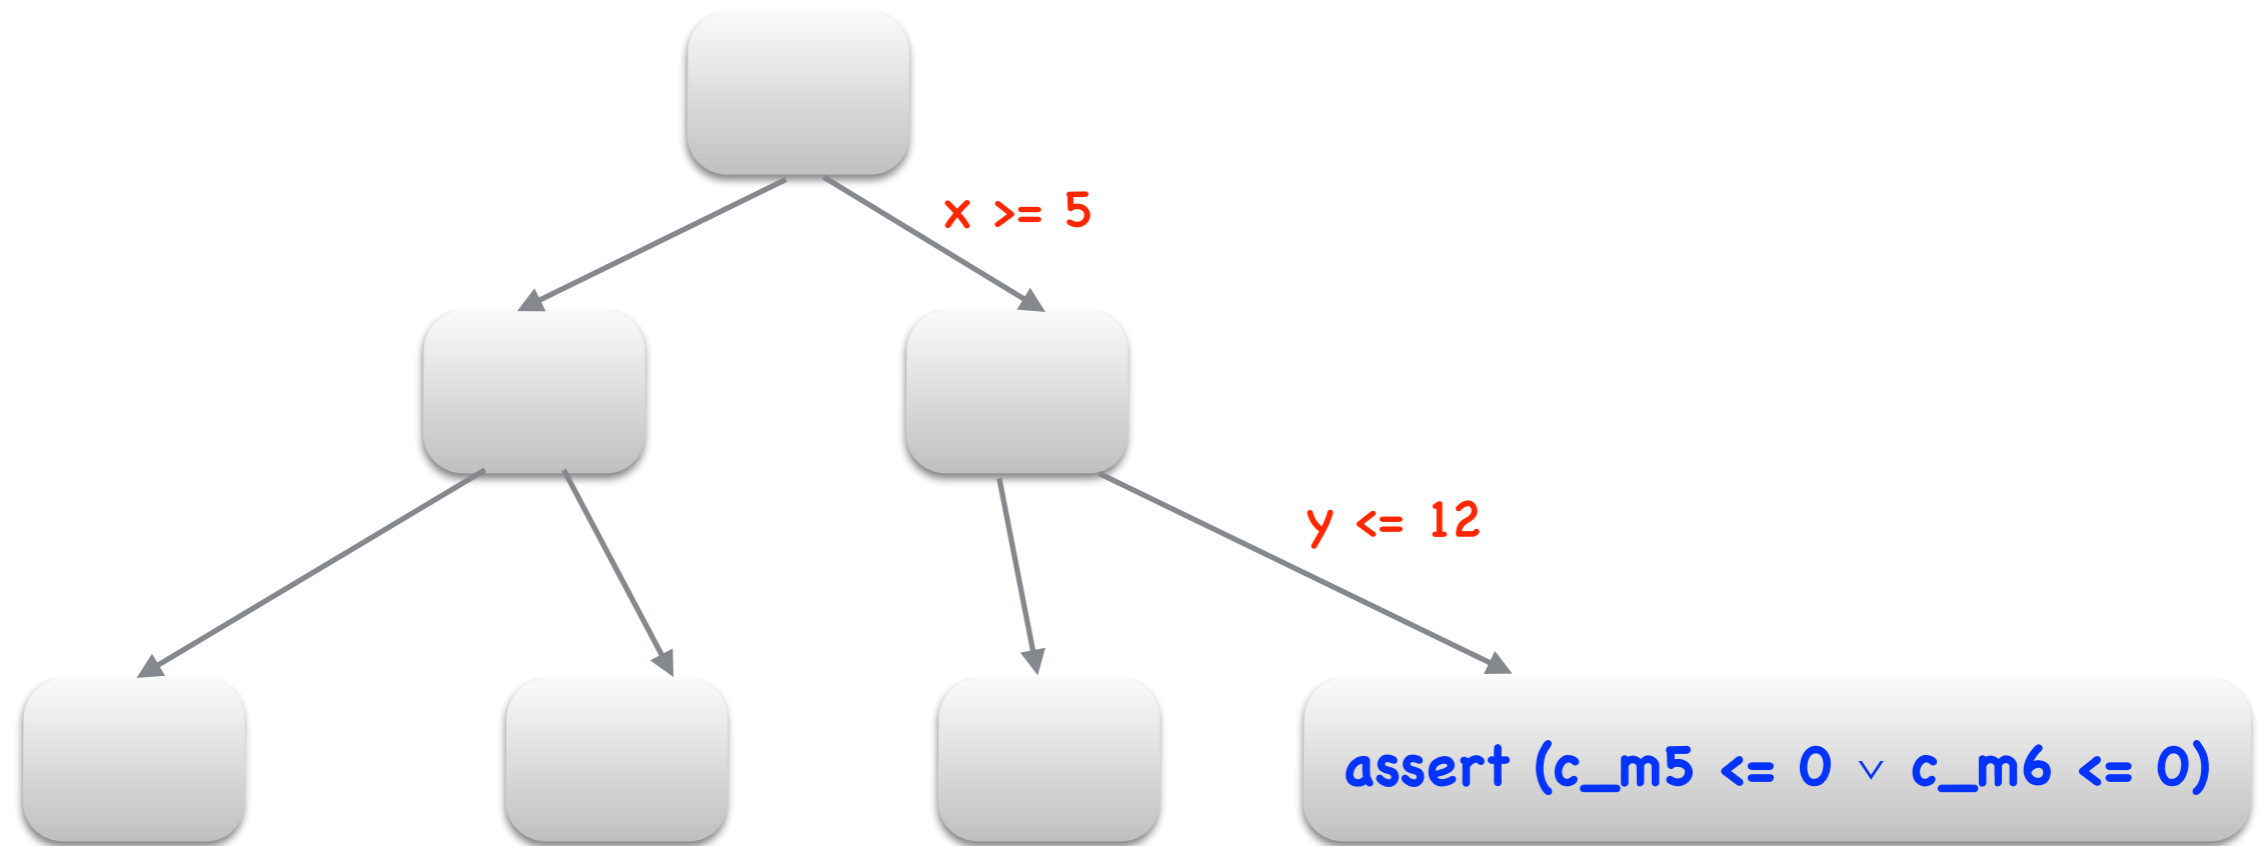
Test Generation Approach



Dynamic Symbolic Execution
Guidance via Control Dependency Graph

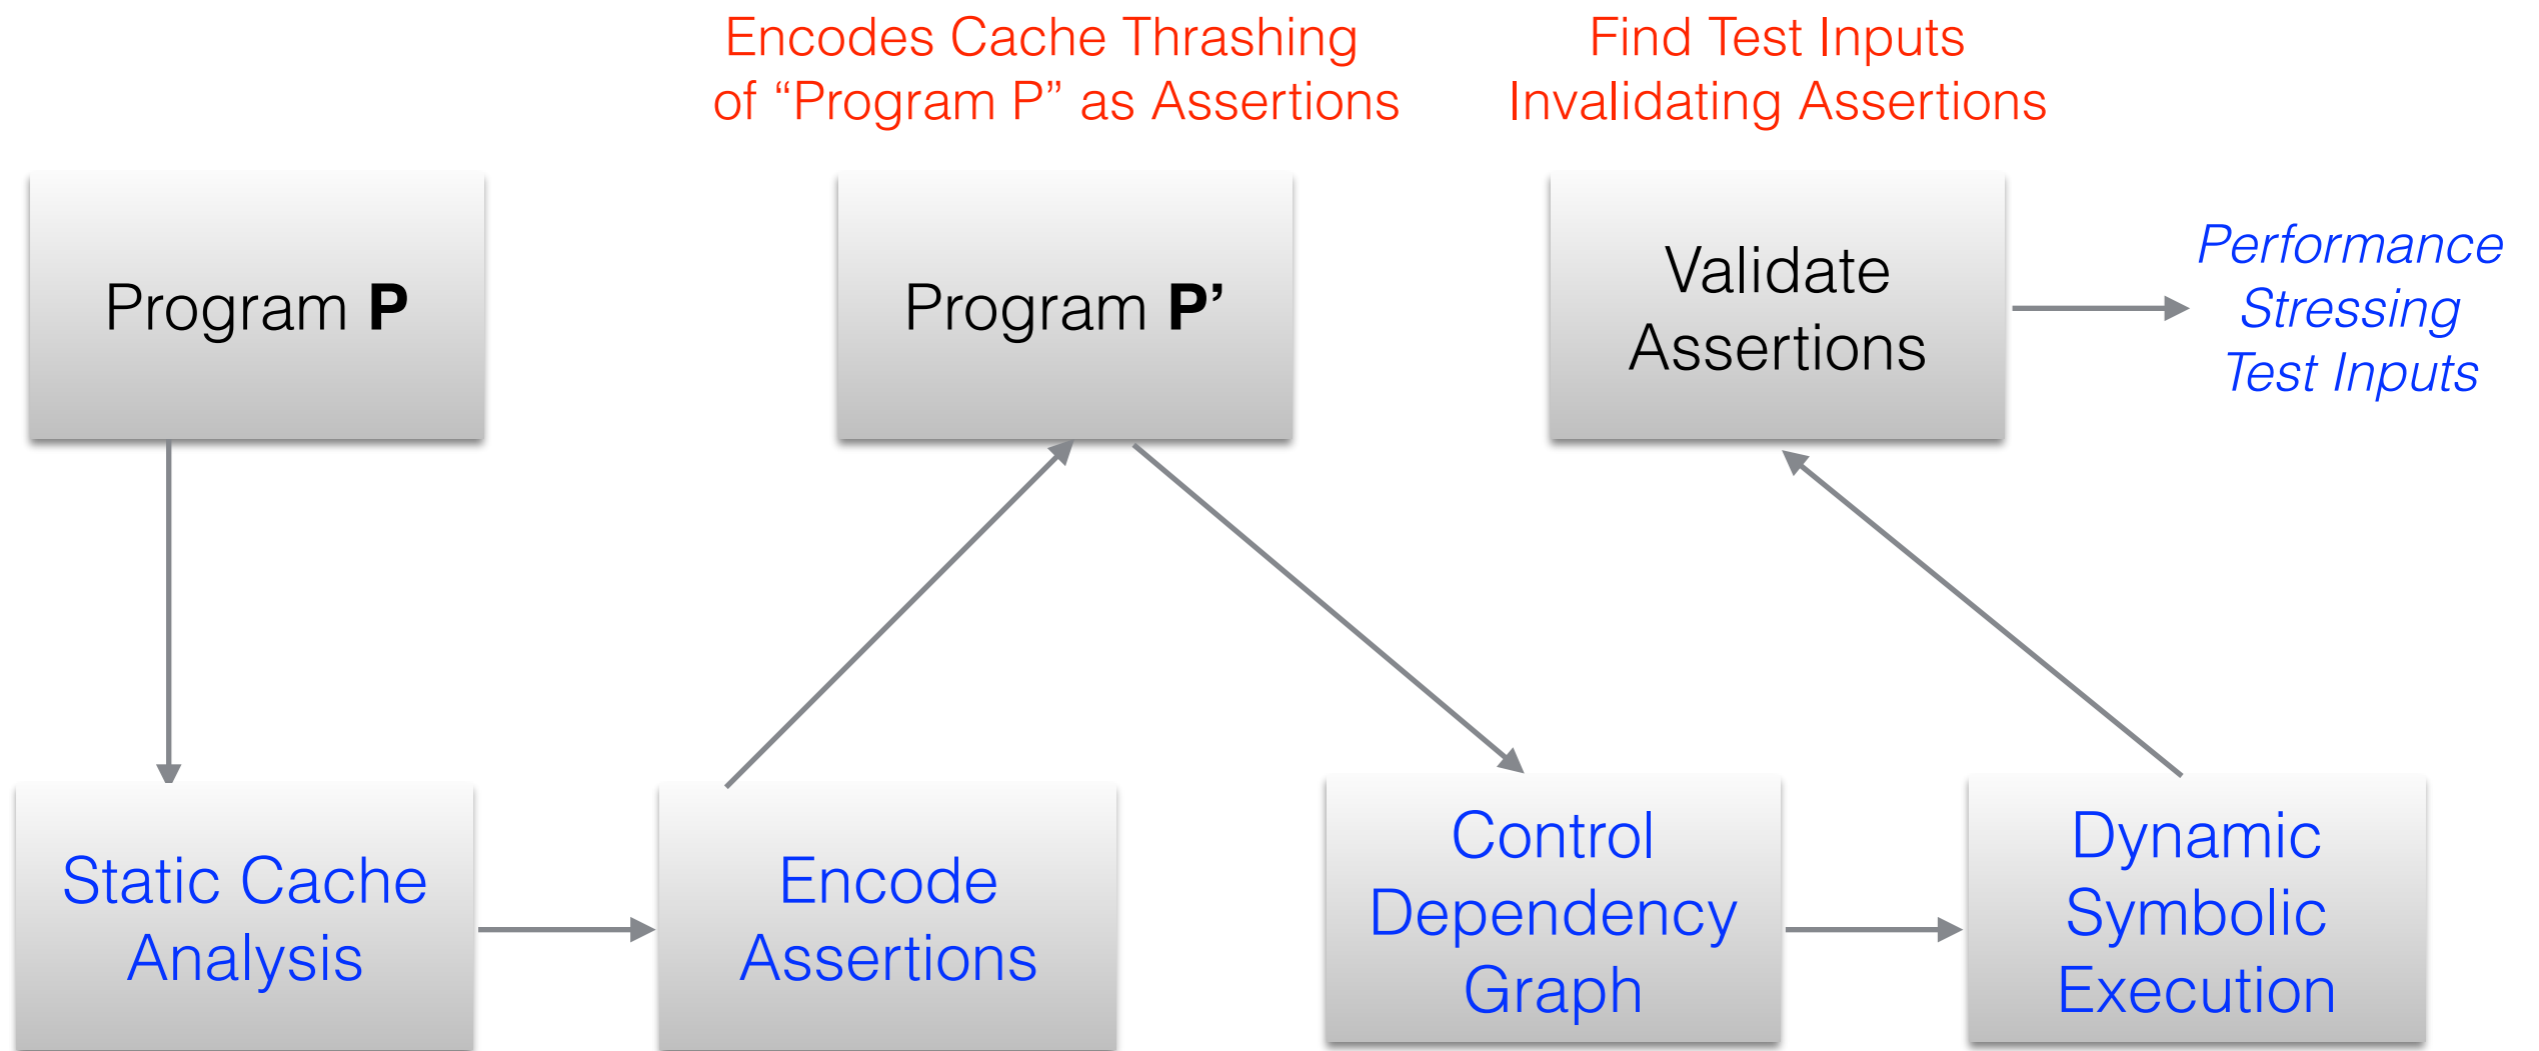
(reaching path to Assertions)

Test Generation Approach

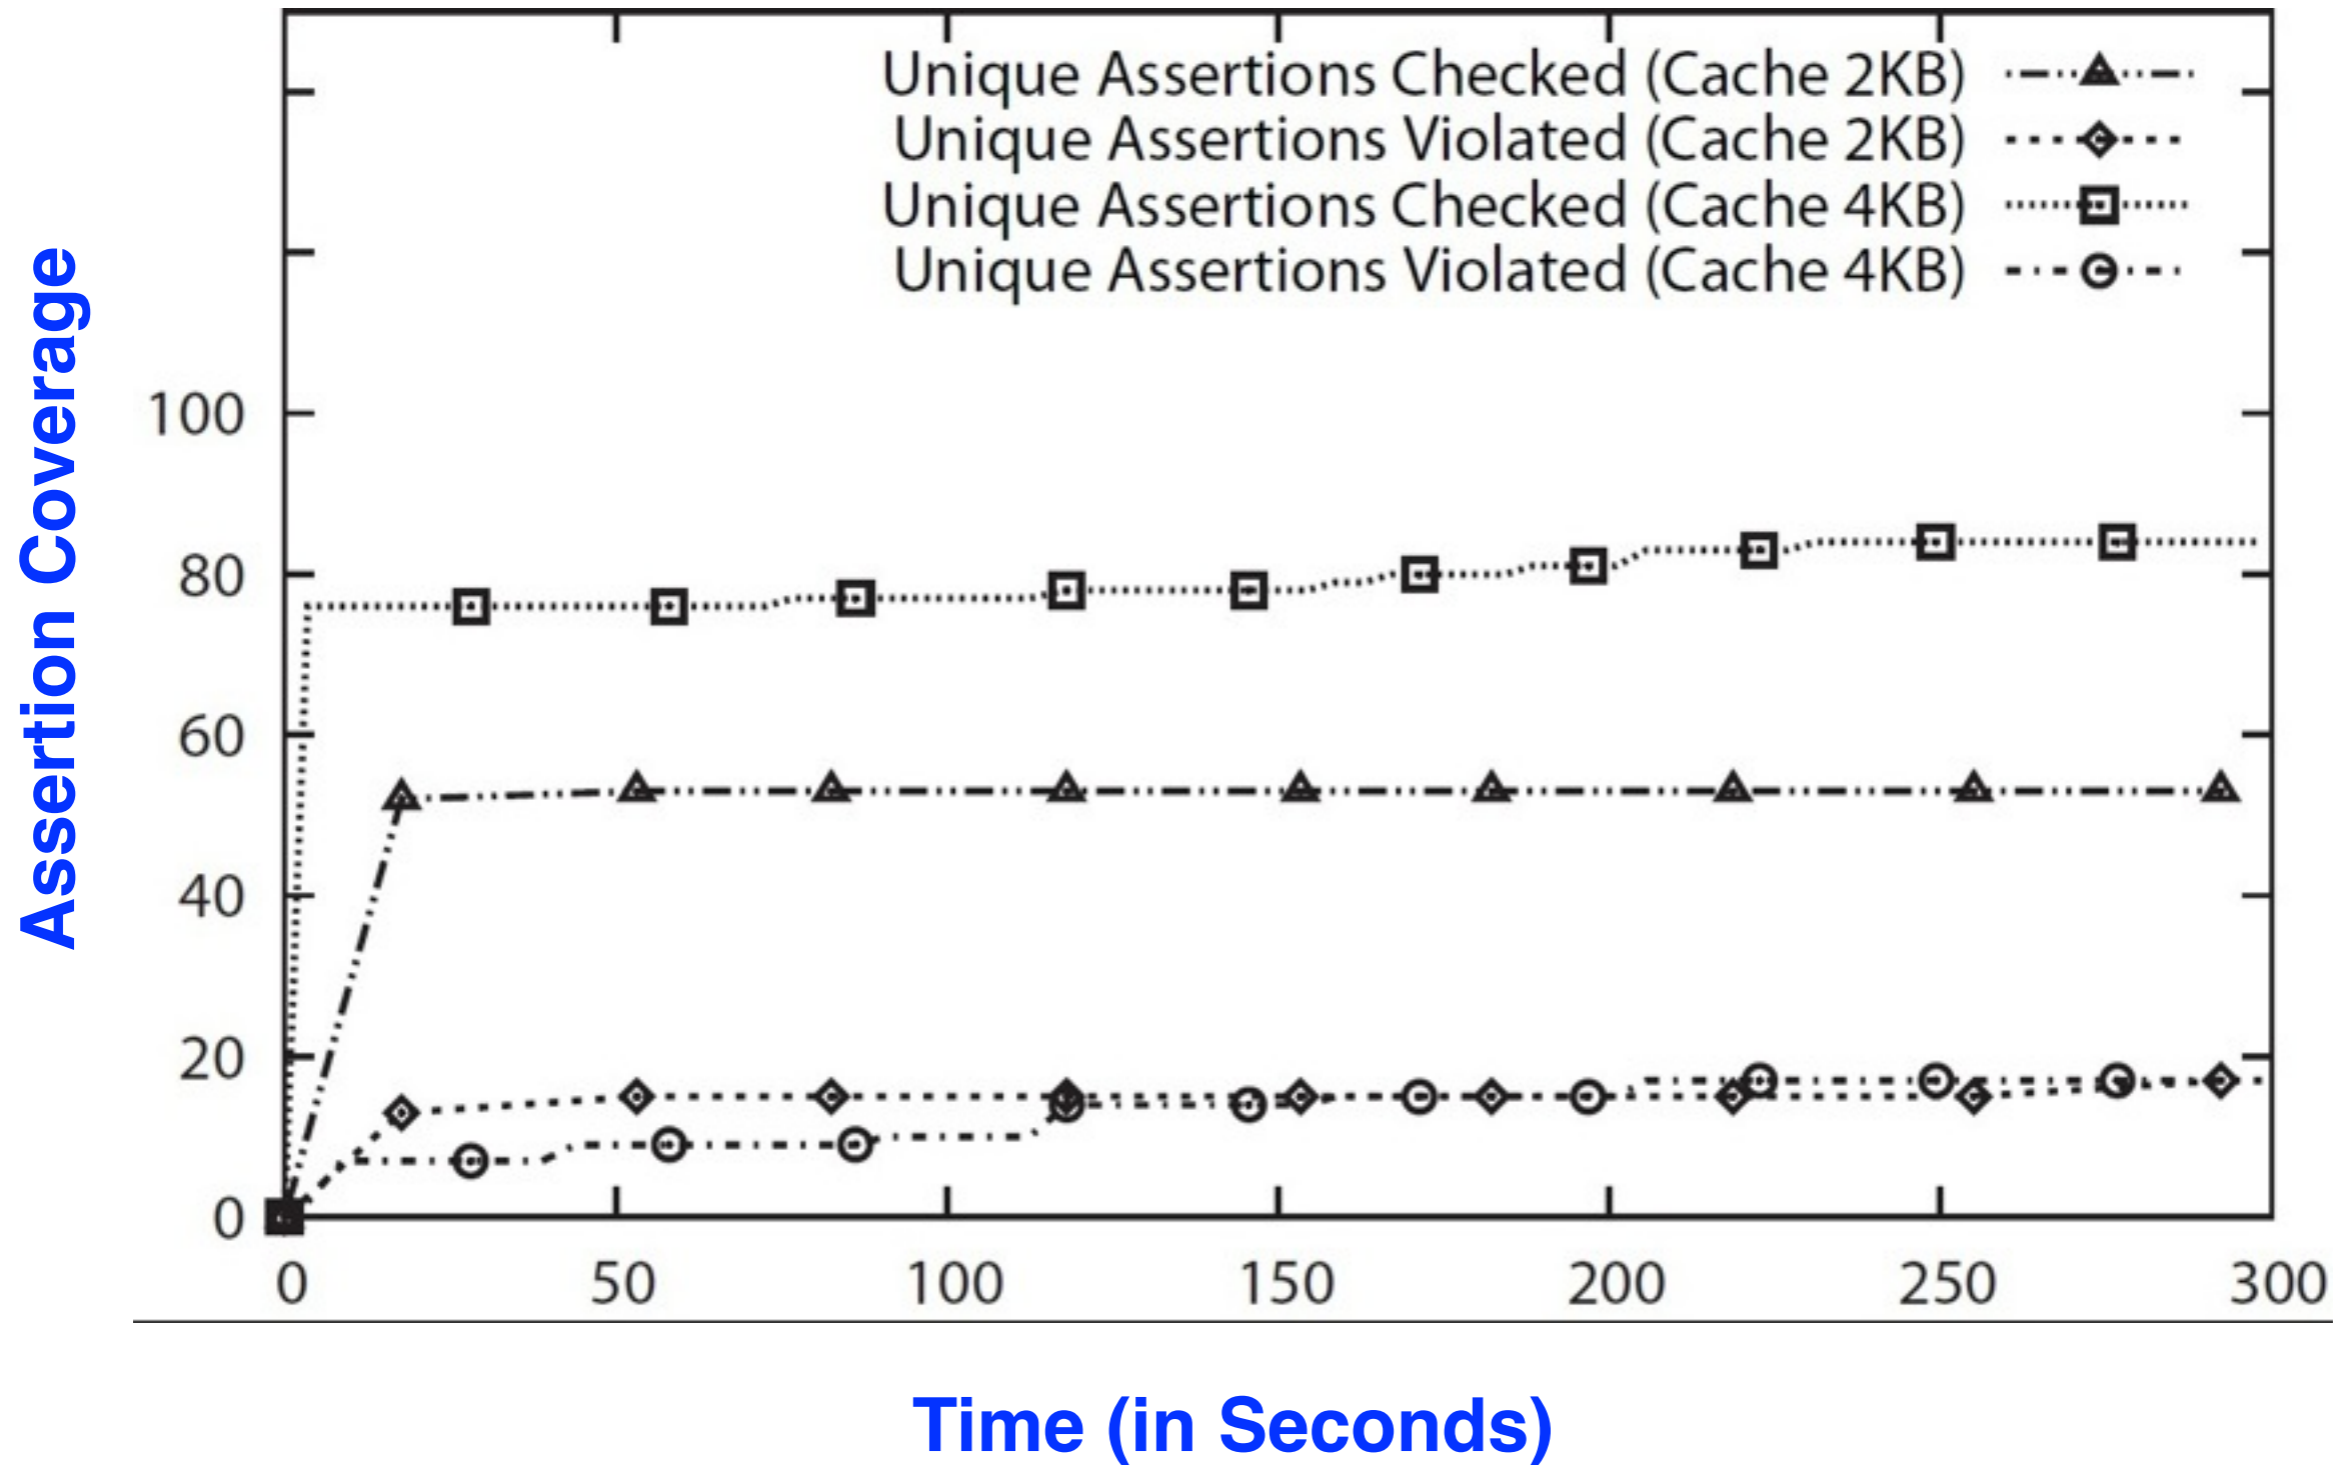


Generate inputs satisfying $(x \geq 5 \wedge y \leq 12)$

Summary



Evaluation



Overview

Tools and Techniques to Write Energy-efficient Software

Energy Testing

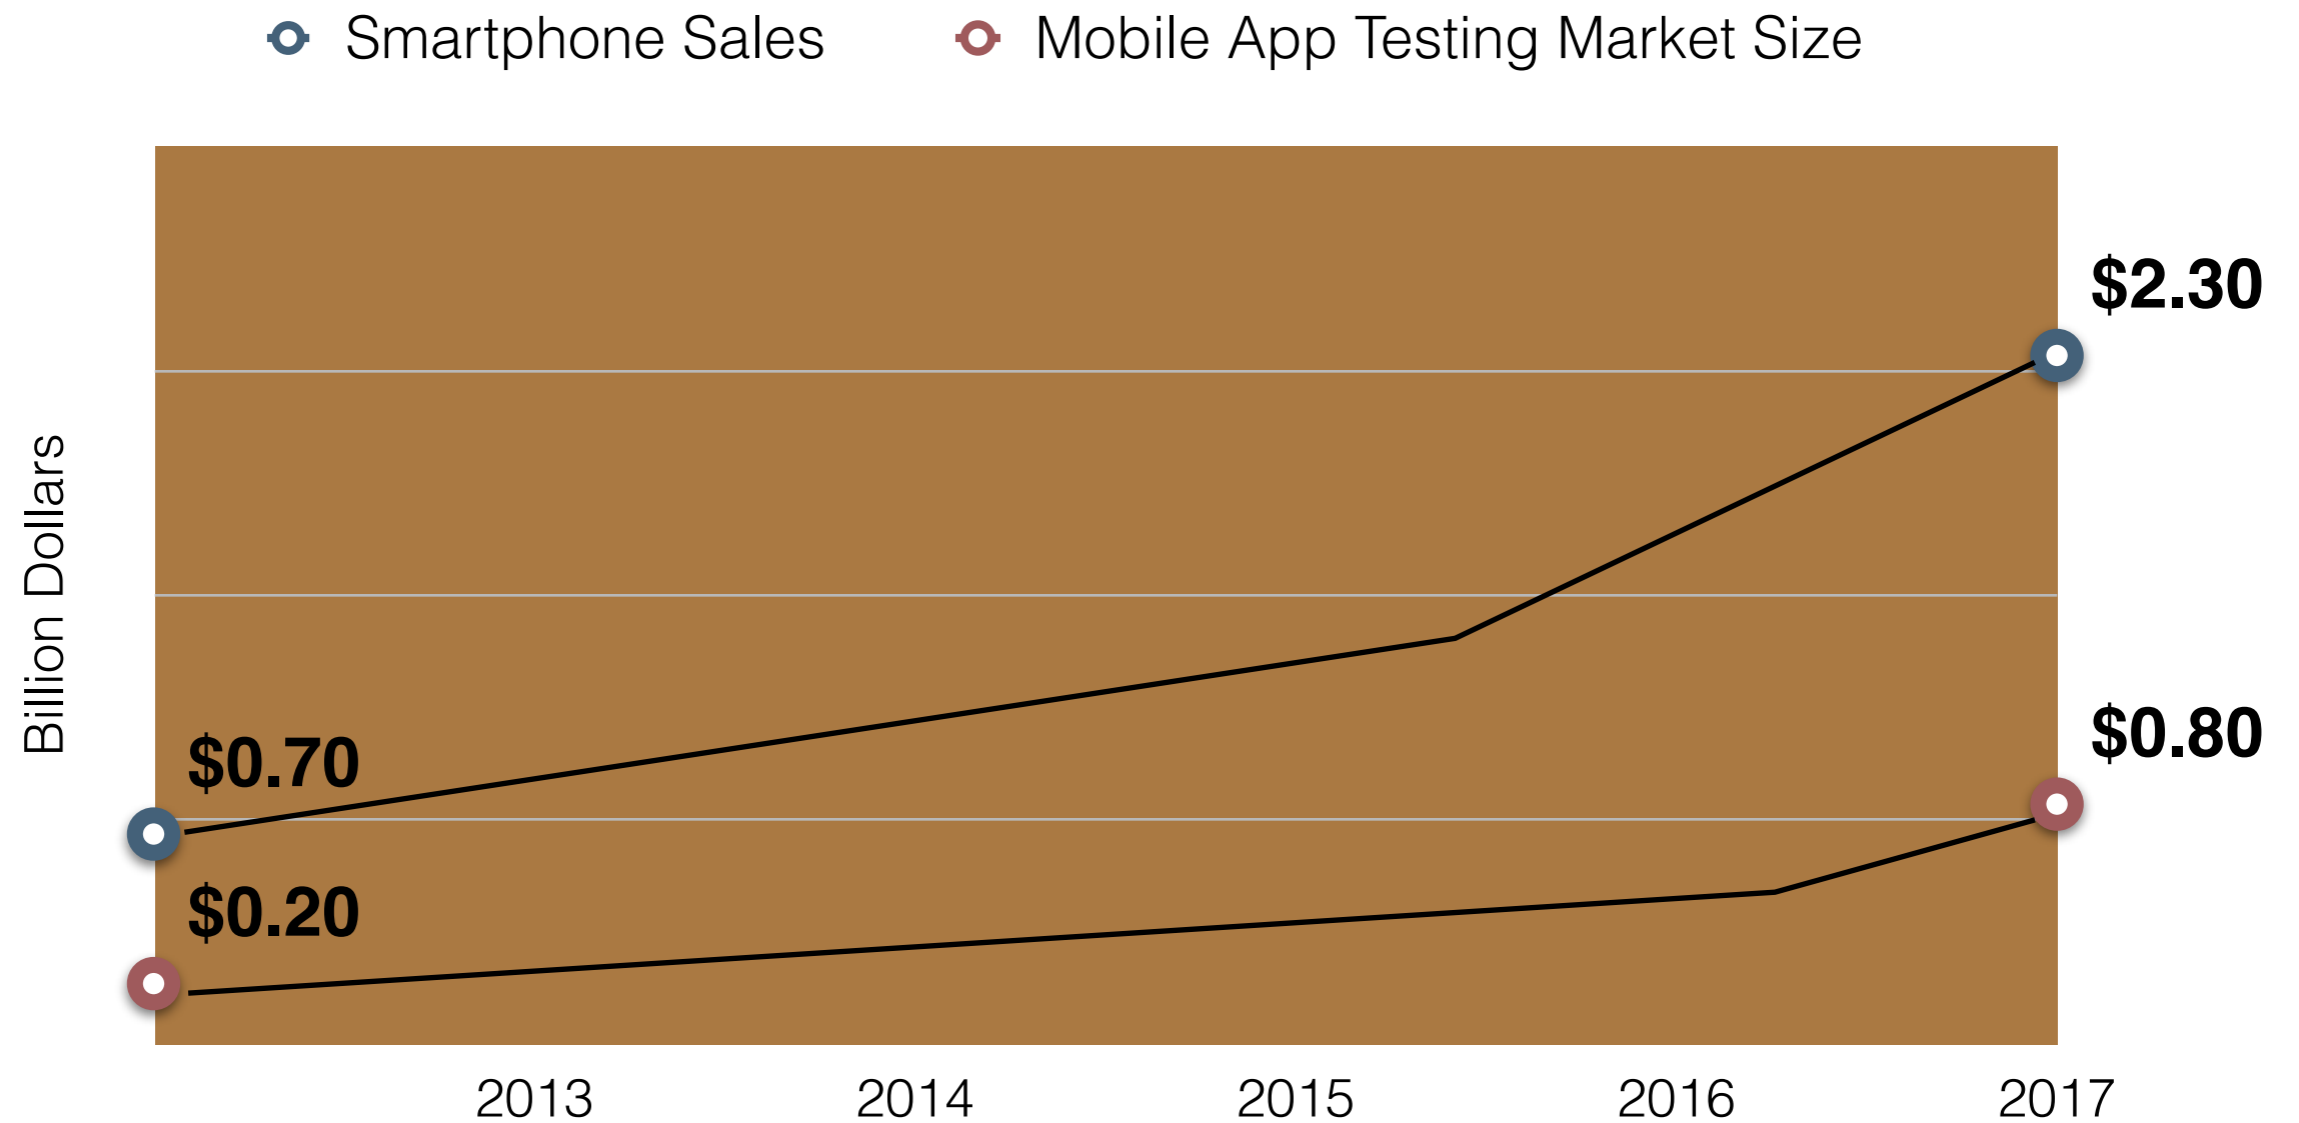
**Energy-hungry
Code Patterns**

**Energy-aware
Coding Guidelines**

Desktop machines, handheld devices etc.

Android Devices

Smartphone Market



Data obtained from IDC, Gartner and ABI Research

Energy Inefficiency

- How do we quantify energy inefficiency?
 - High energy consumption, what is *high?*
- High energy consumption
 - High utilization of hardware components
 - *Low utilization of hardware components*
- Ratio Energy/Utilization

Energy Inefficiency

Cause/Source		
Hardware components	Resource leak	Suboptimal resource binding
Sleep state transition	Wakelock bug	Tail Energy hotspot
Background Service	Vacuous background service	Expensive background service
Defective Functionality	Immortality bug	Loop energy hotspot

Energy Inefficiency



Suboptimal resource binding



Vacuous background service

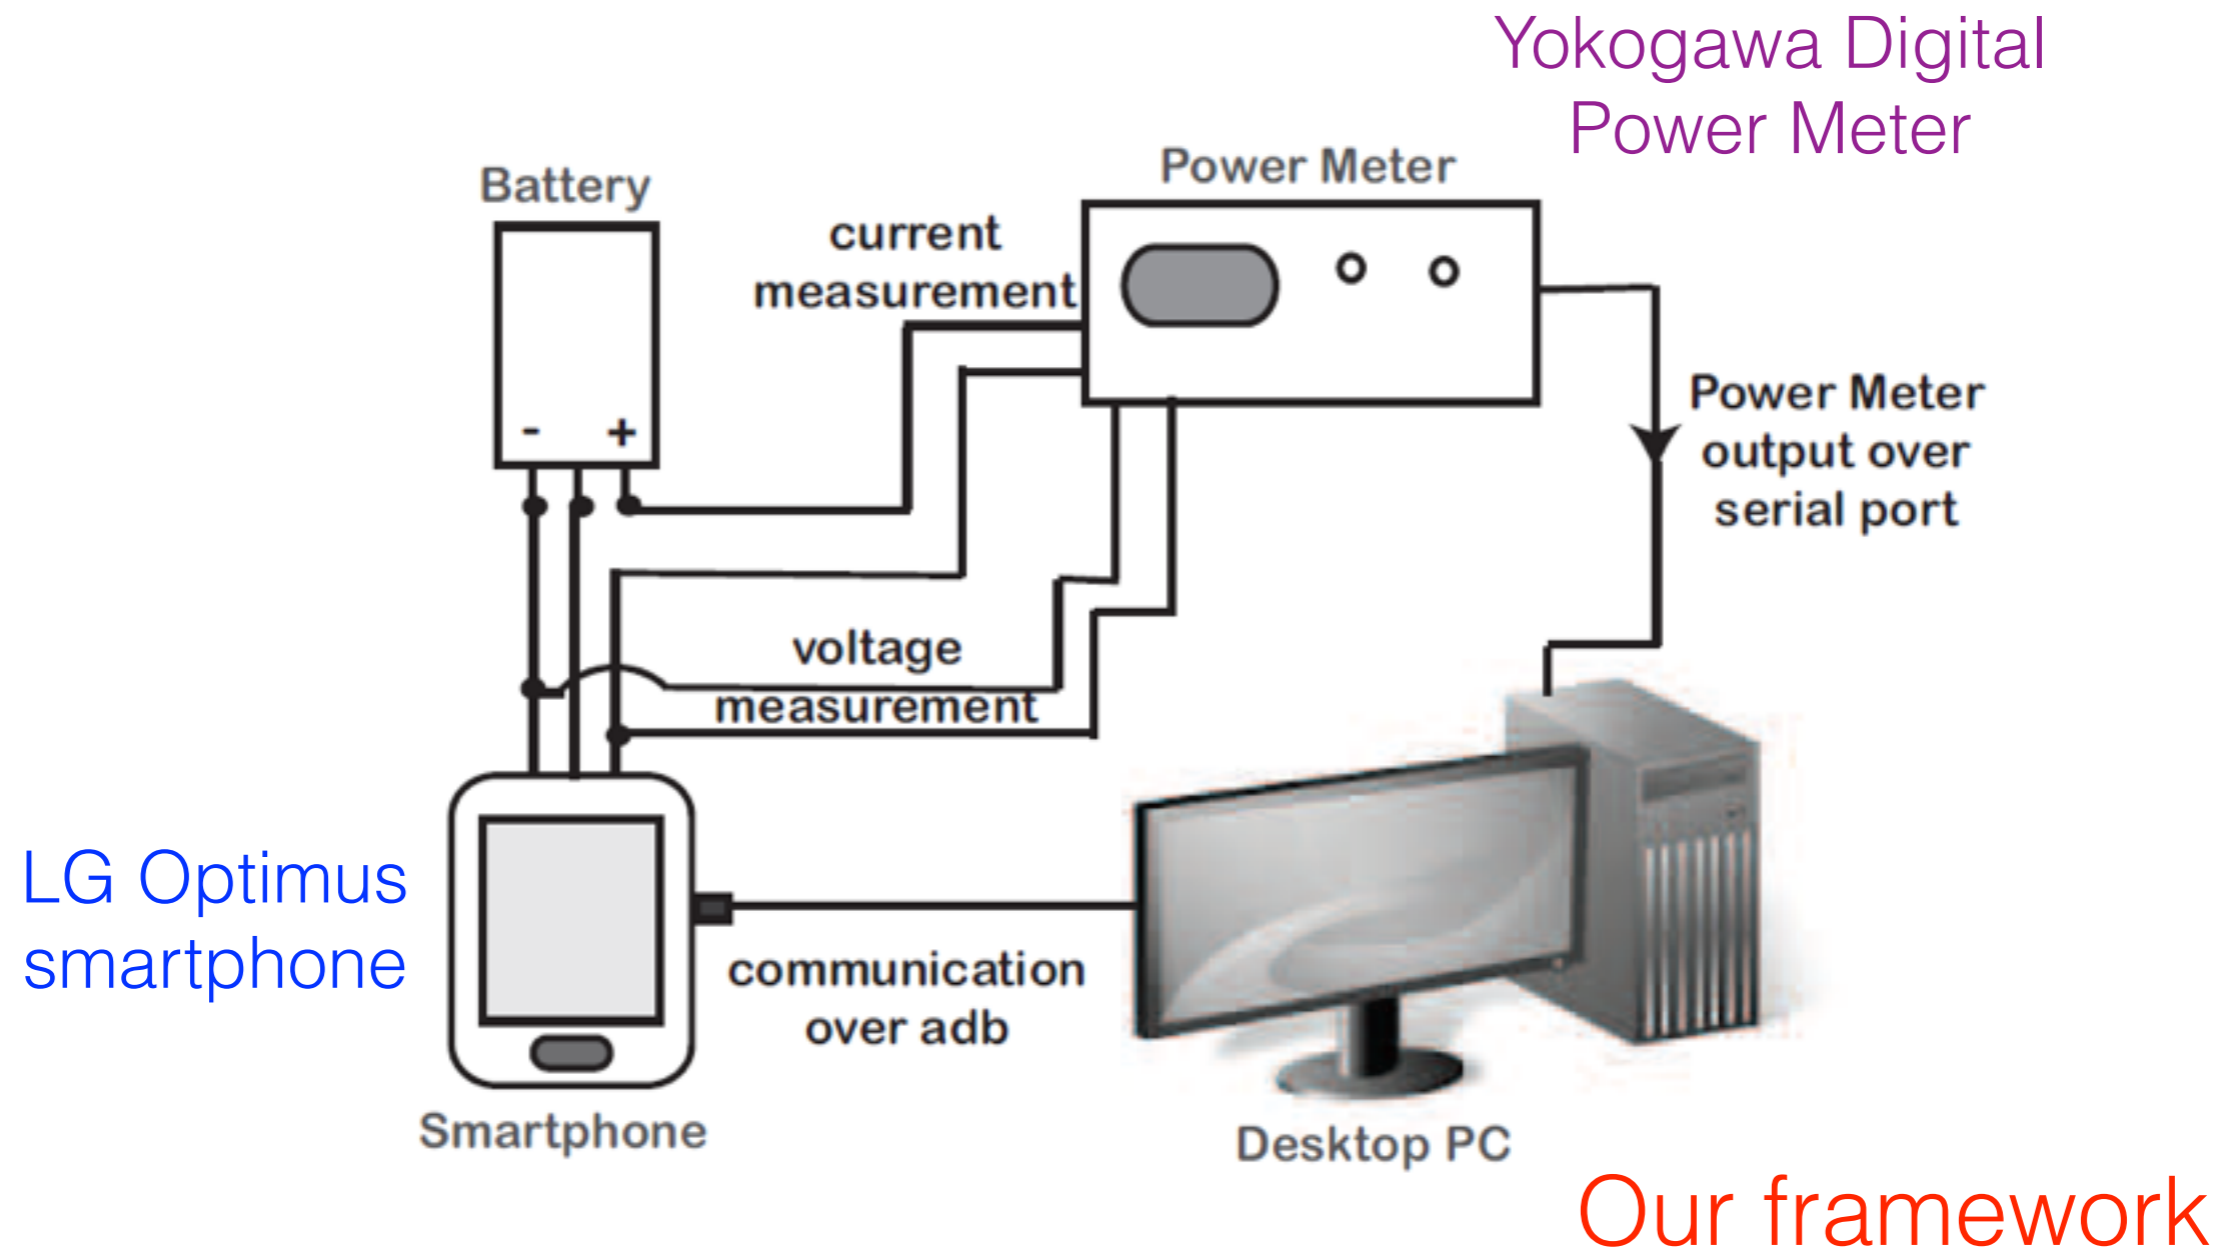
A Broader Categorization

Cause/Source	Energy Bugs	Energy Hotspots
Hardware components	Resource leak	Suboptimal resource binding
Sleep state transition	Wakelock bug	Tail Energy hotspot
Background Service	Vacuous background service	Expensive background service
Defective Functionality	Immortality bug	Loop energy hotspot

Device does not return to idle

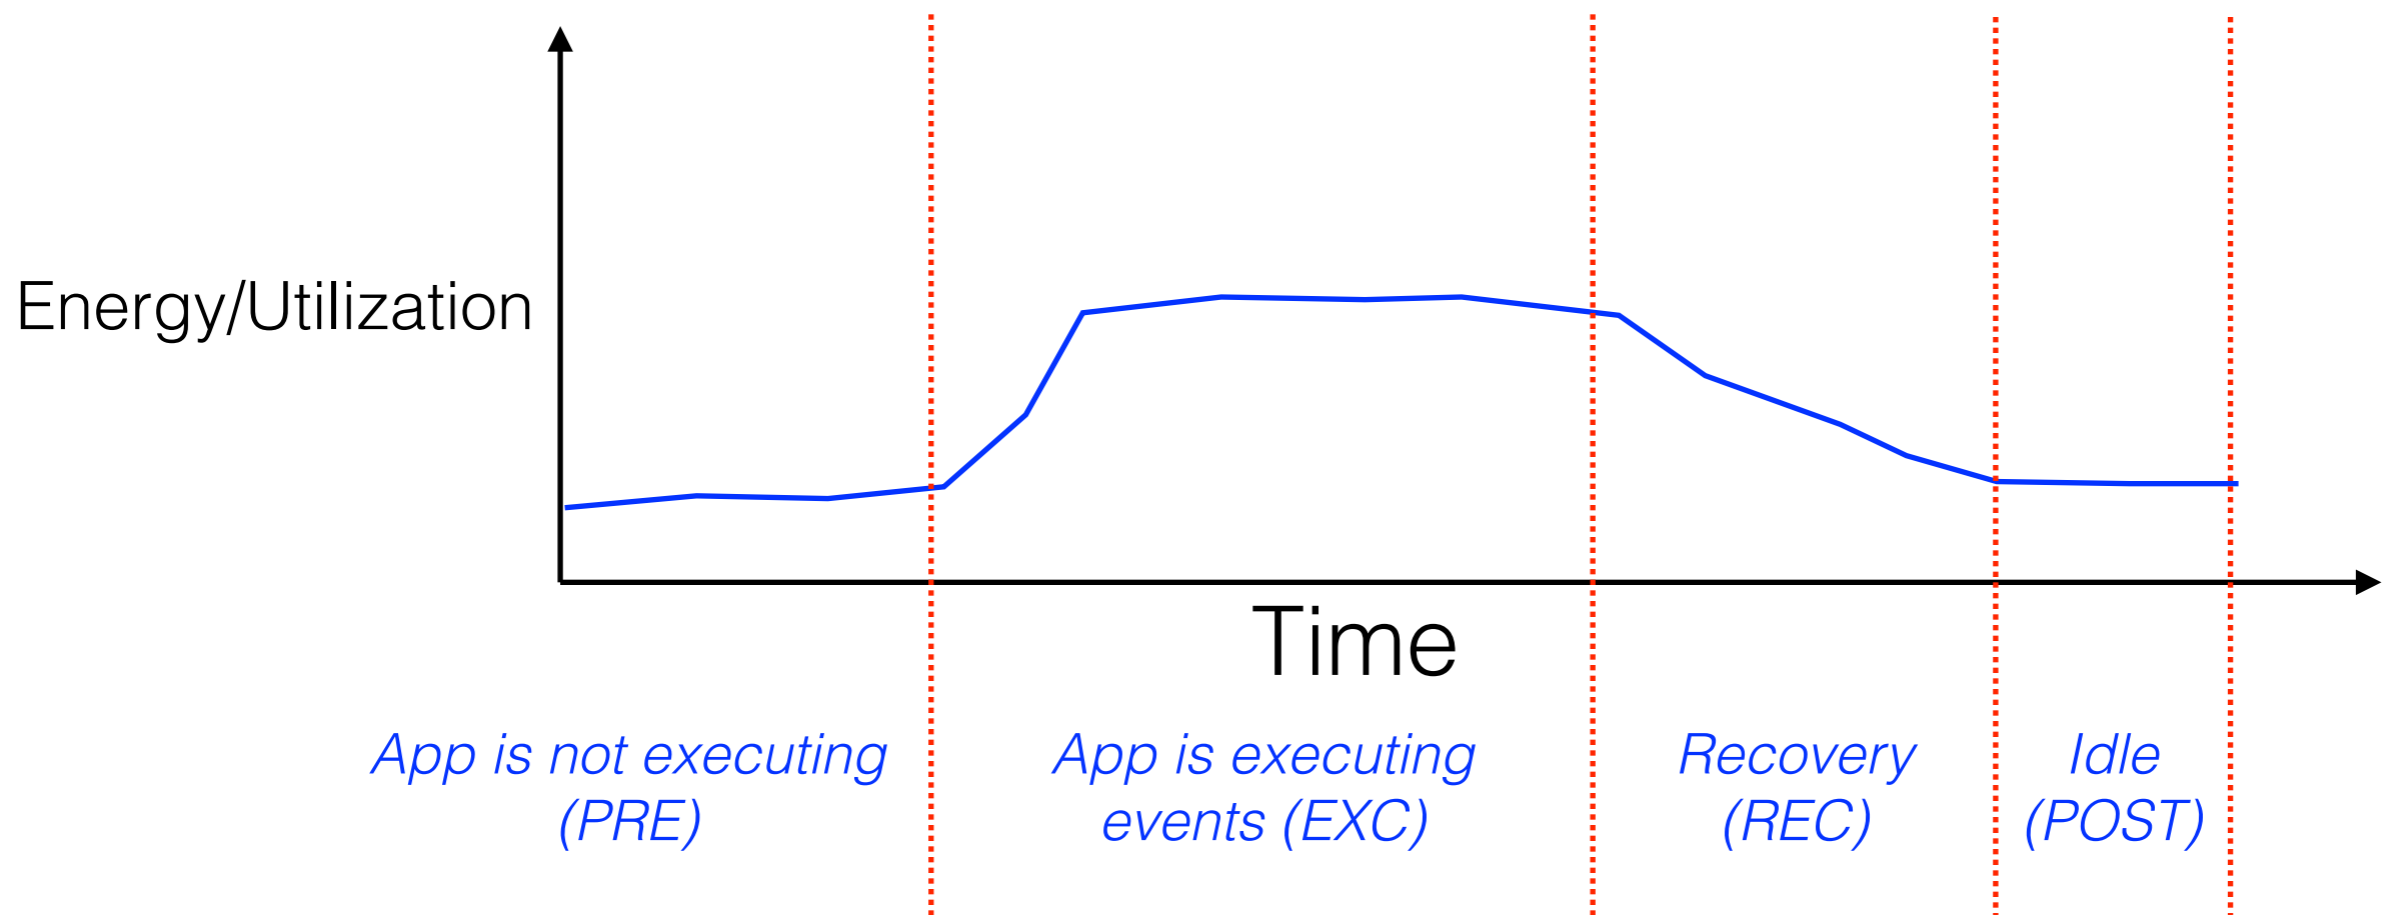
High energy consumption + low utilization

Measurement

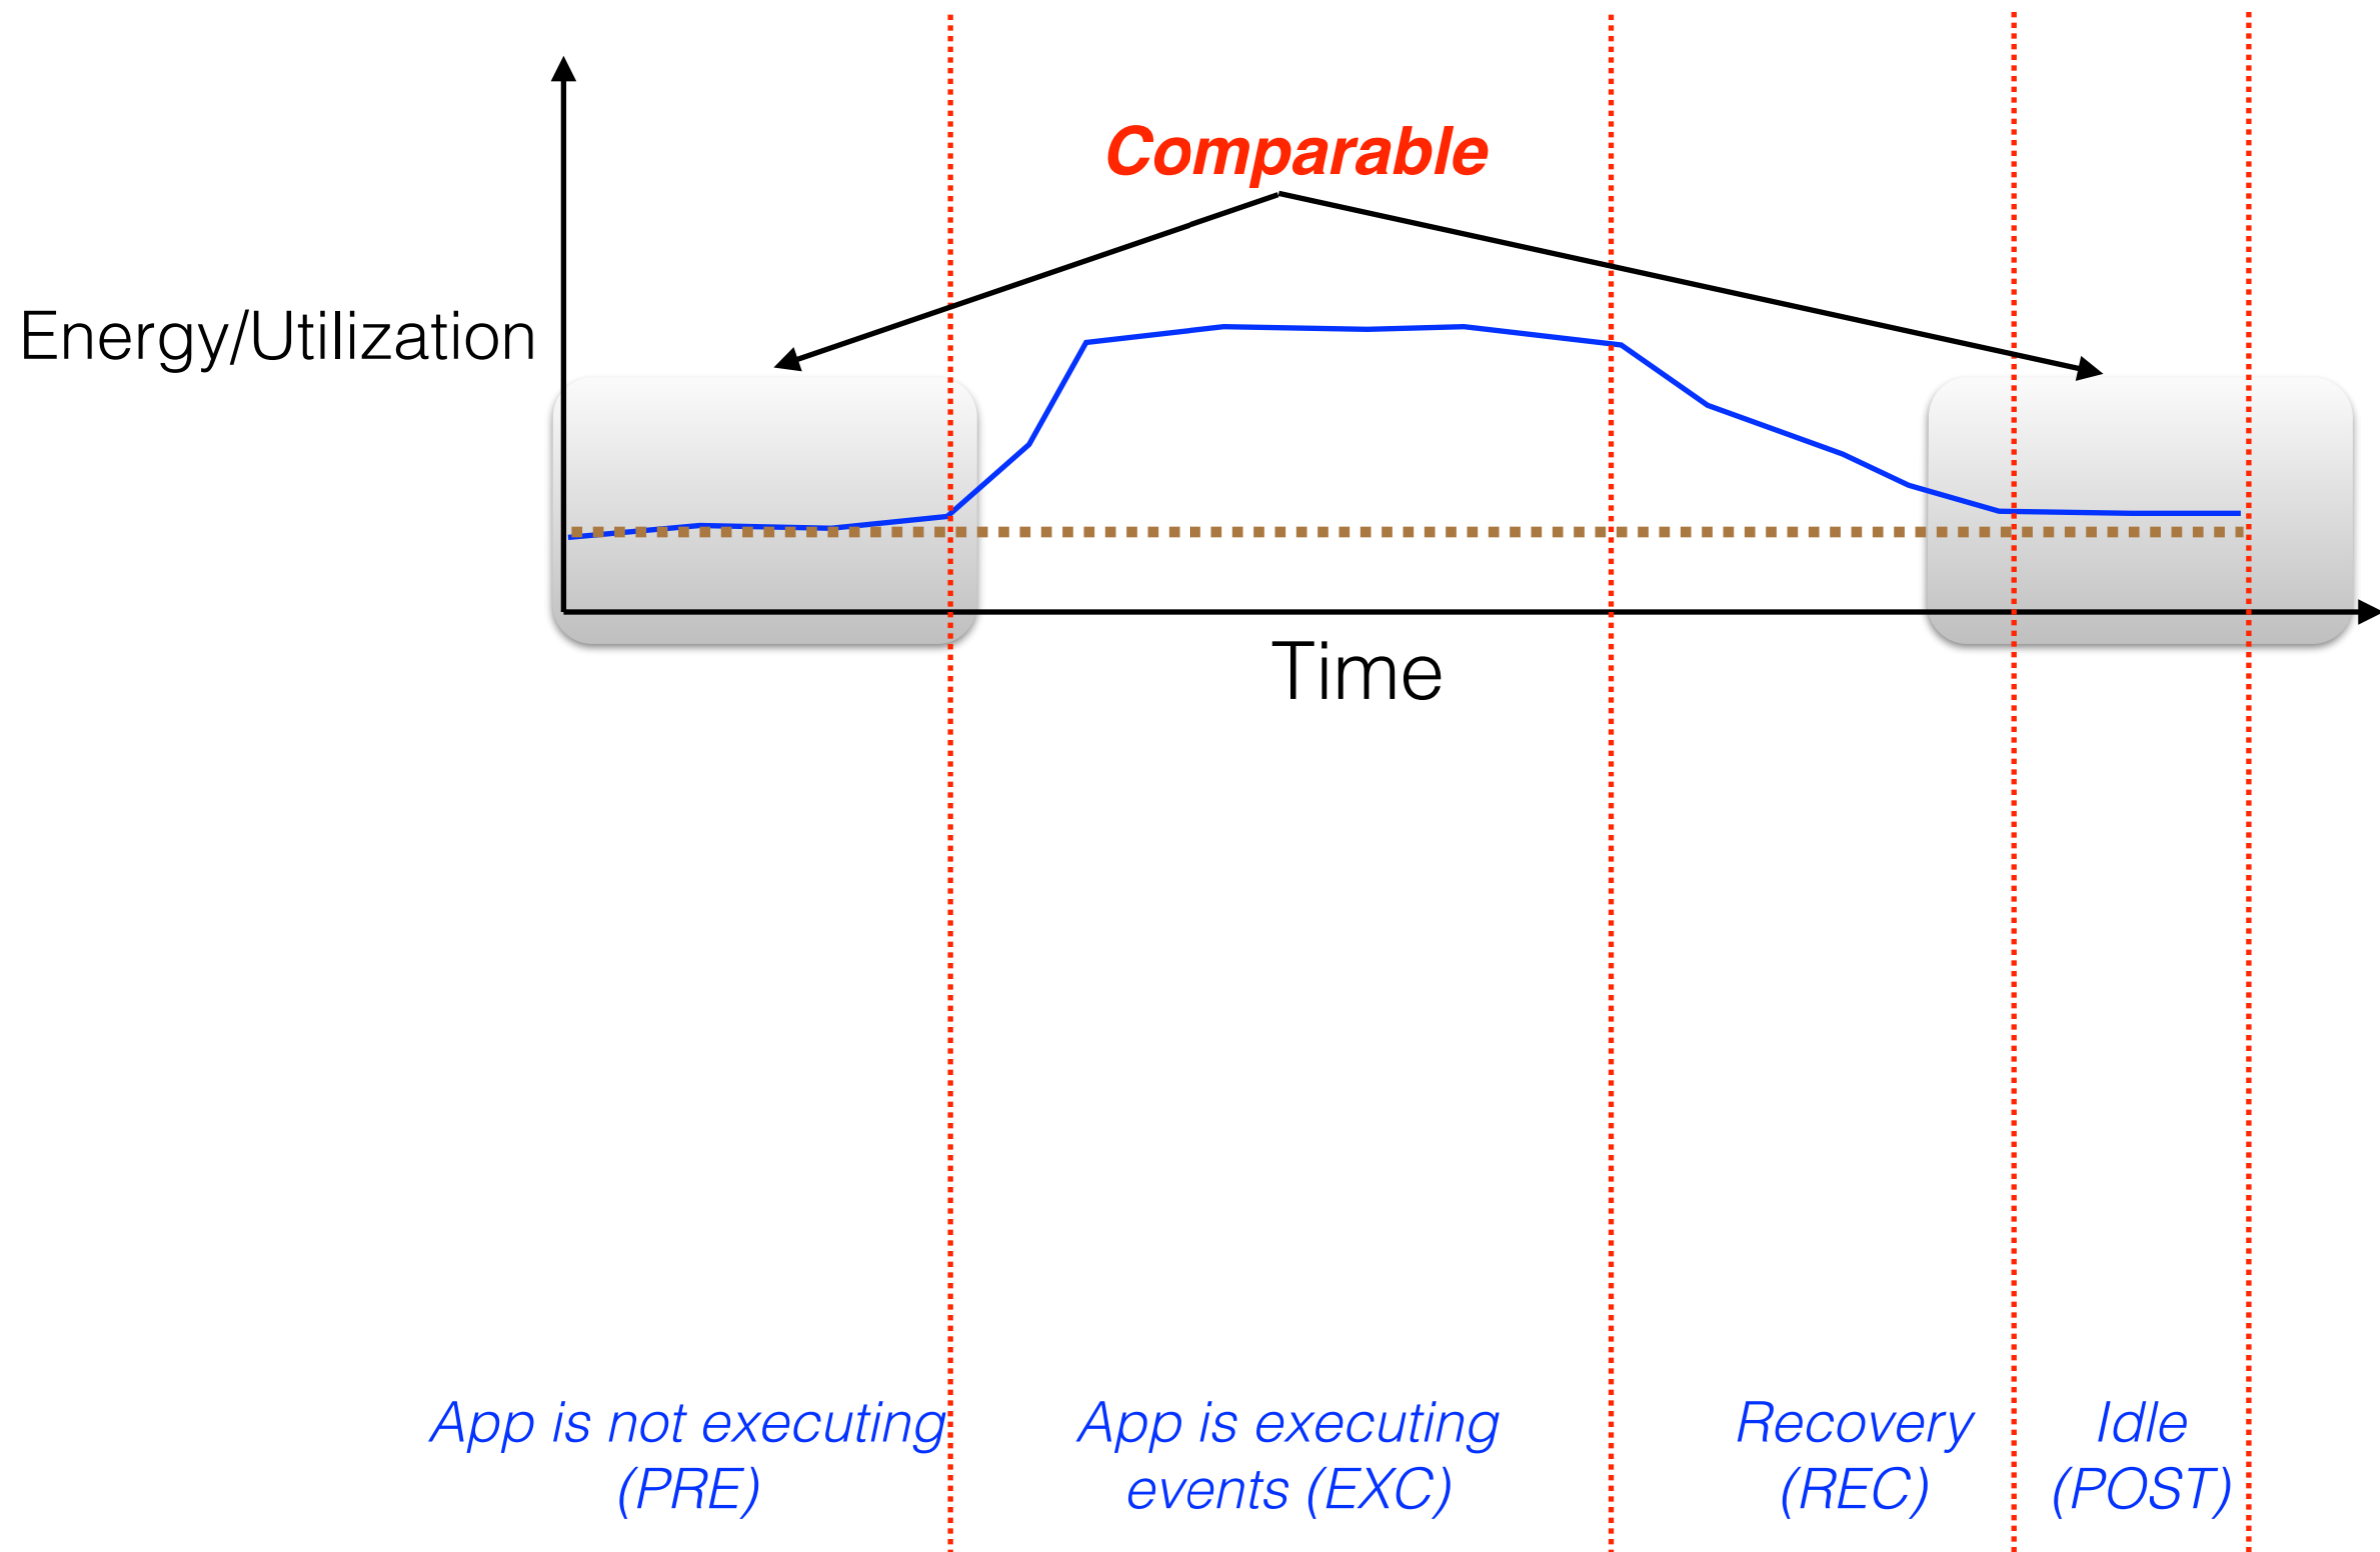


Measurement

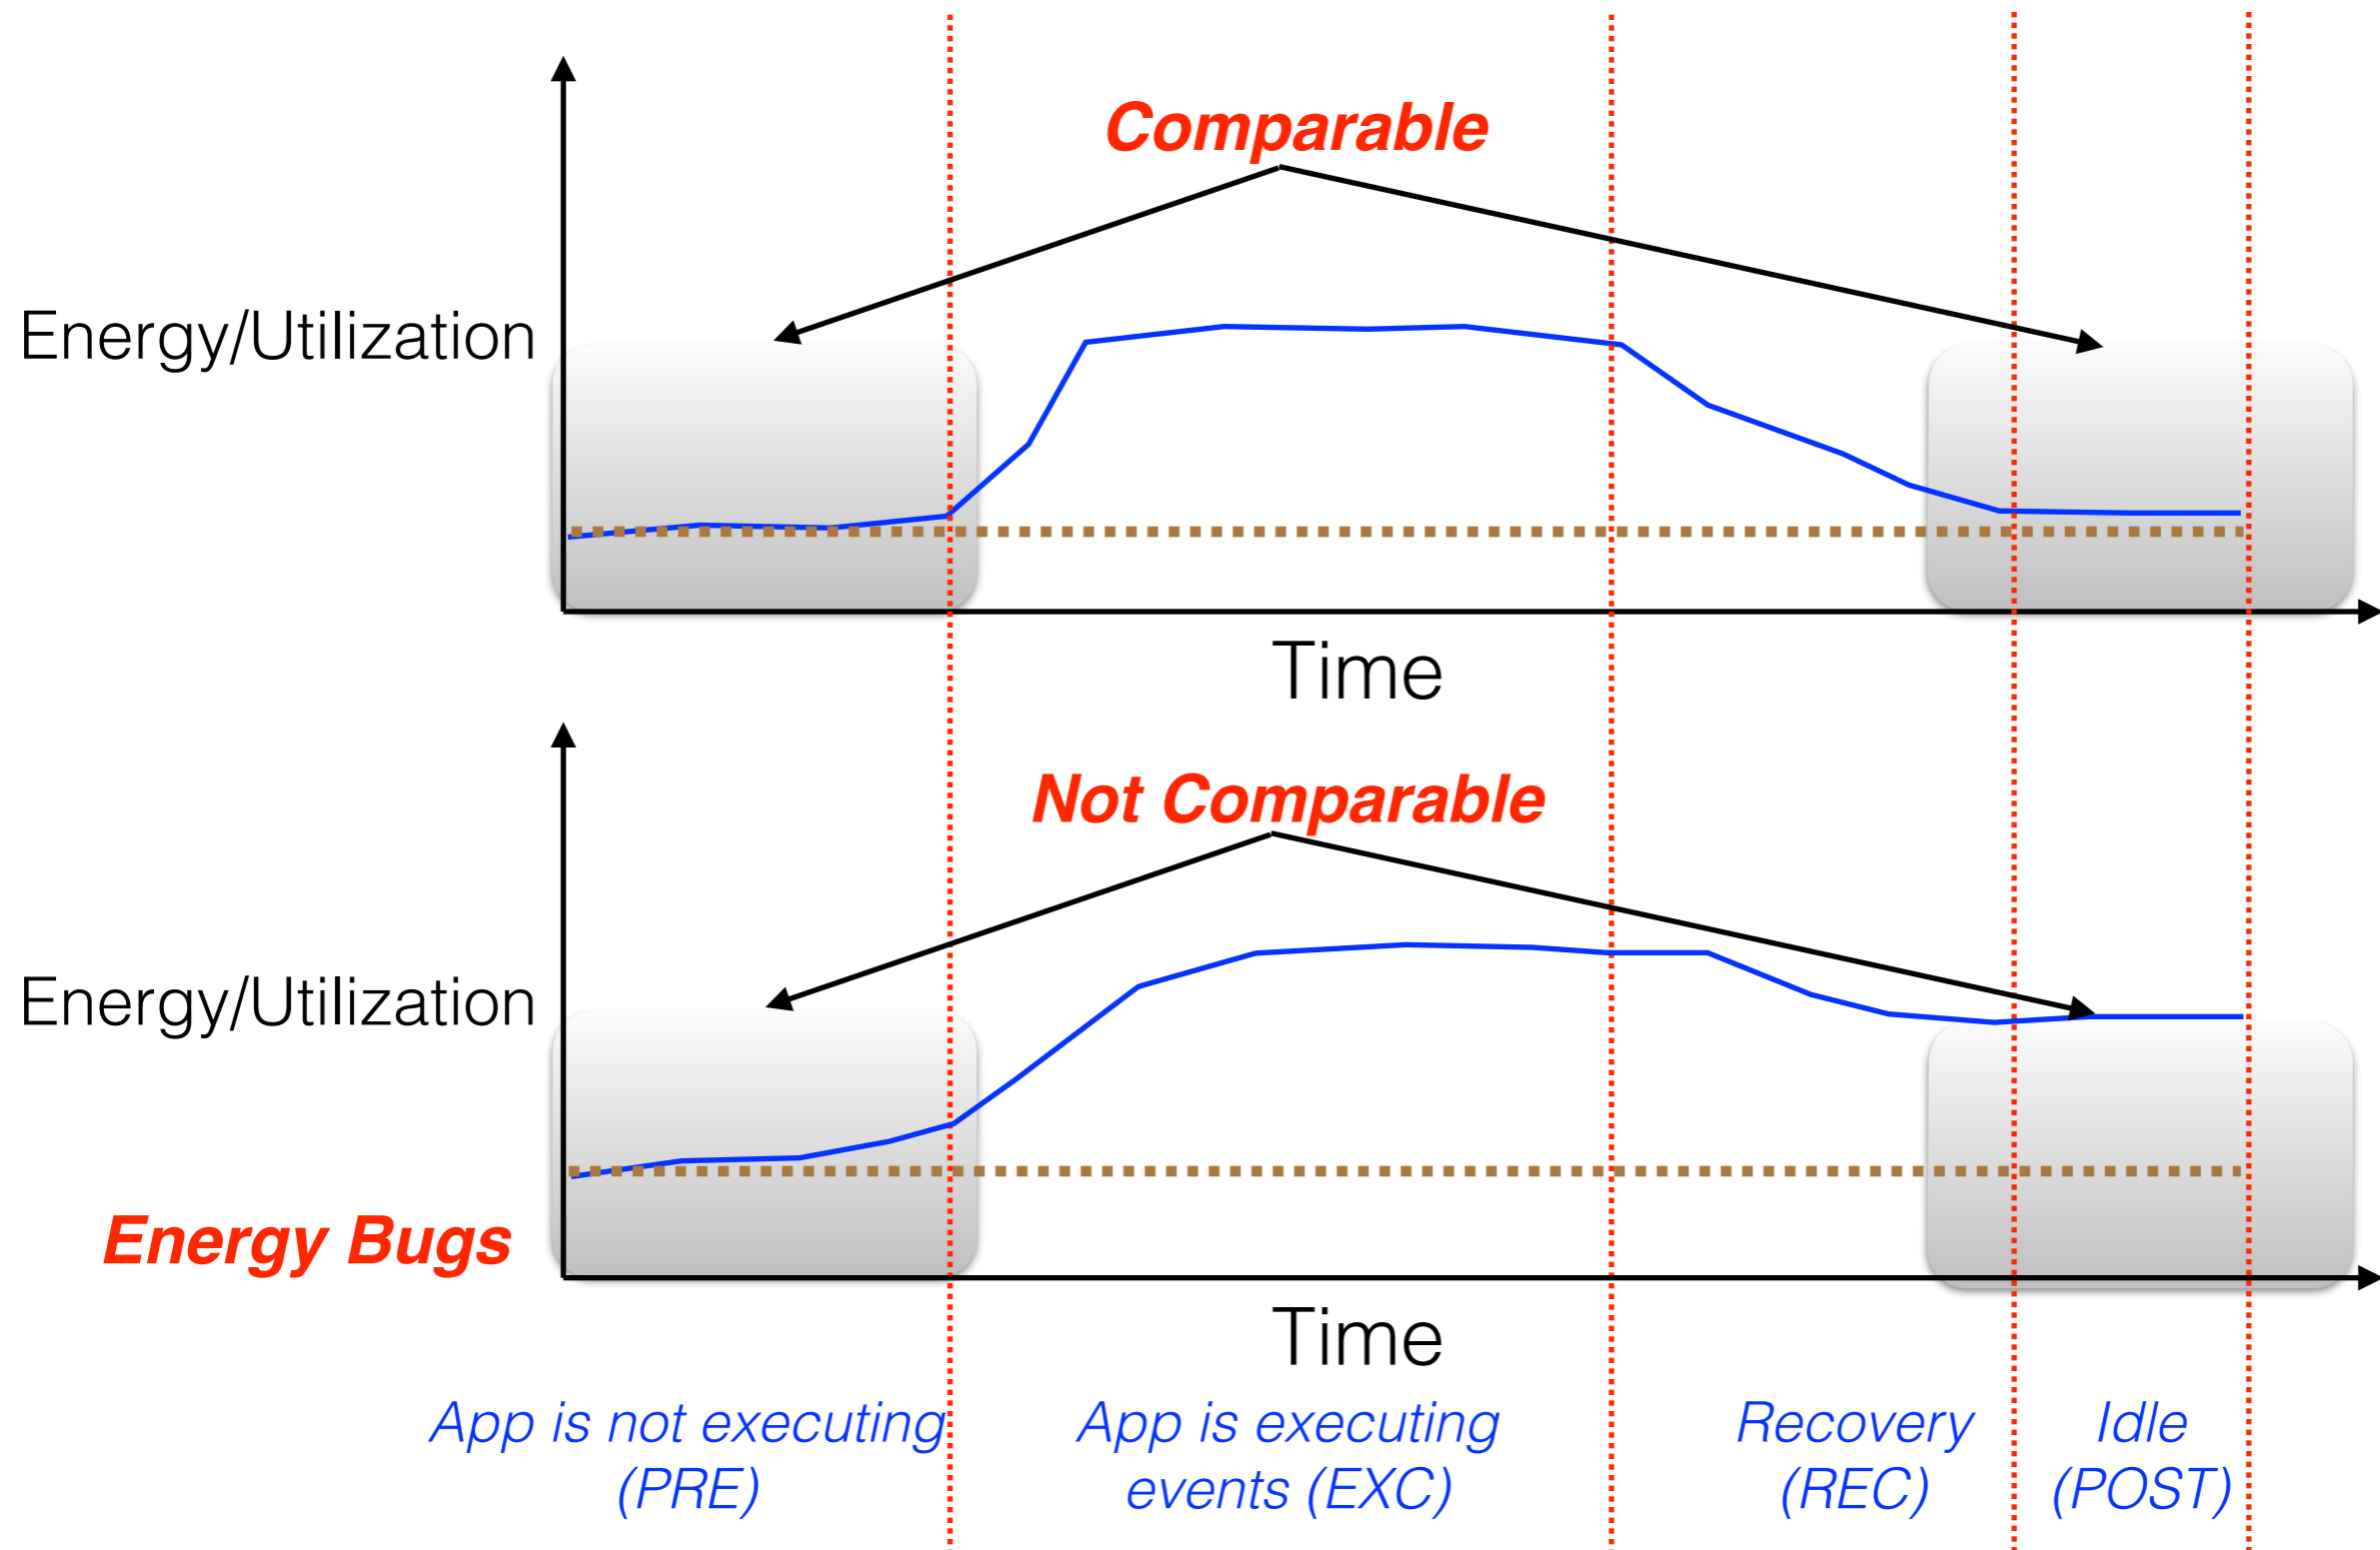
- Measuring Energy/Utilization ratio for an application



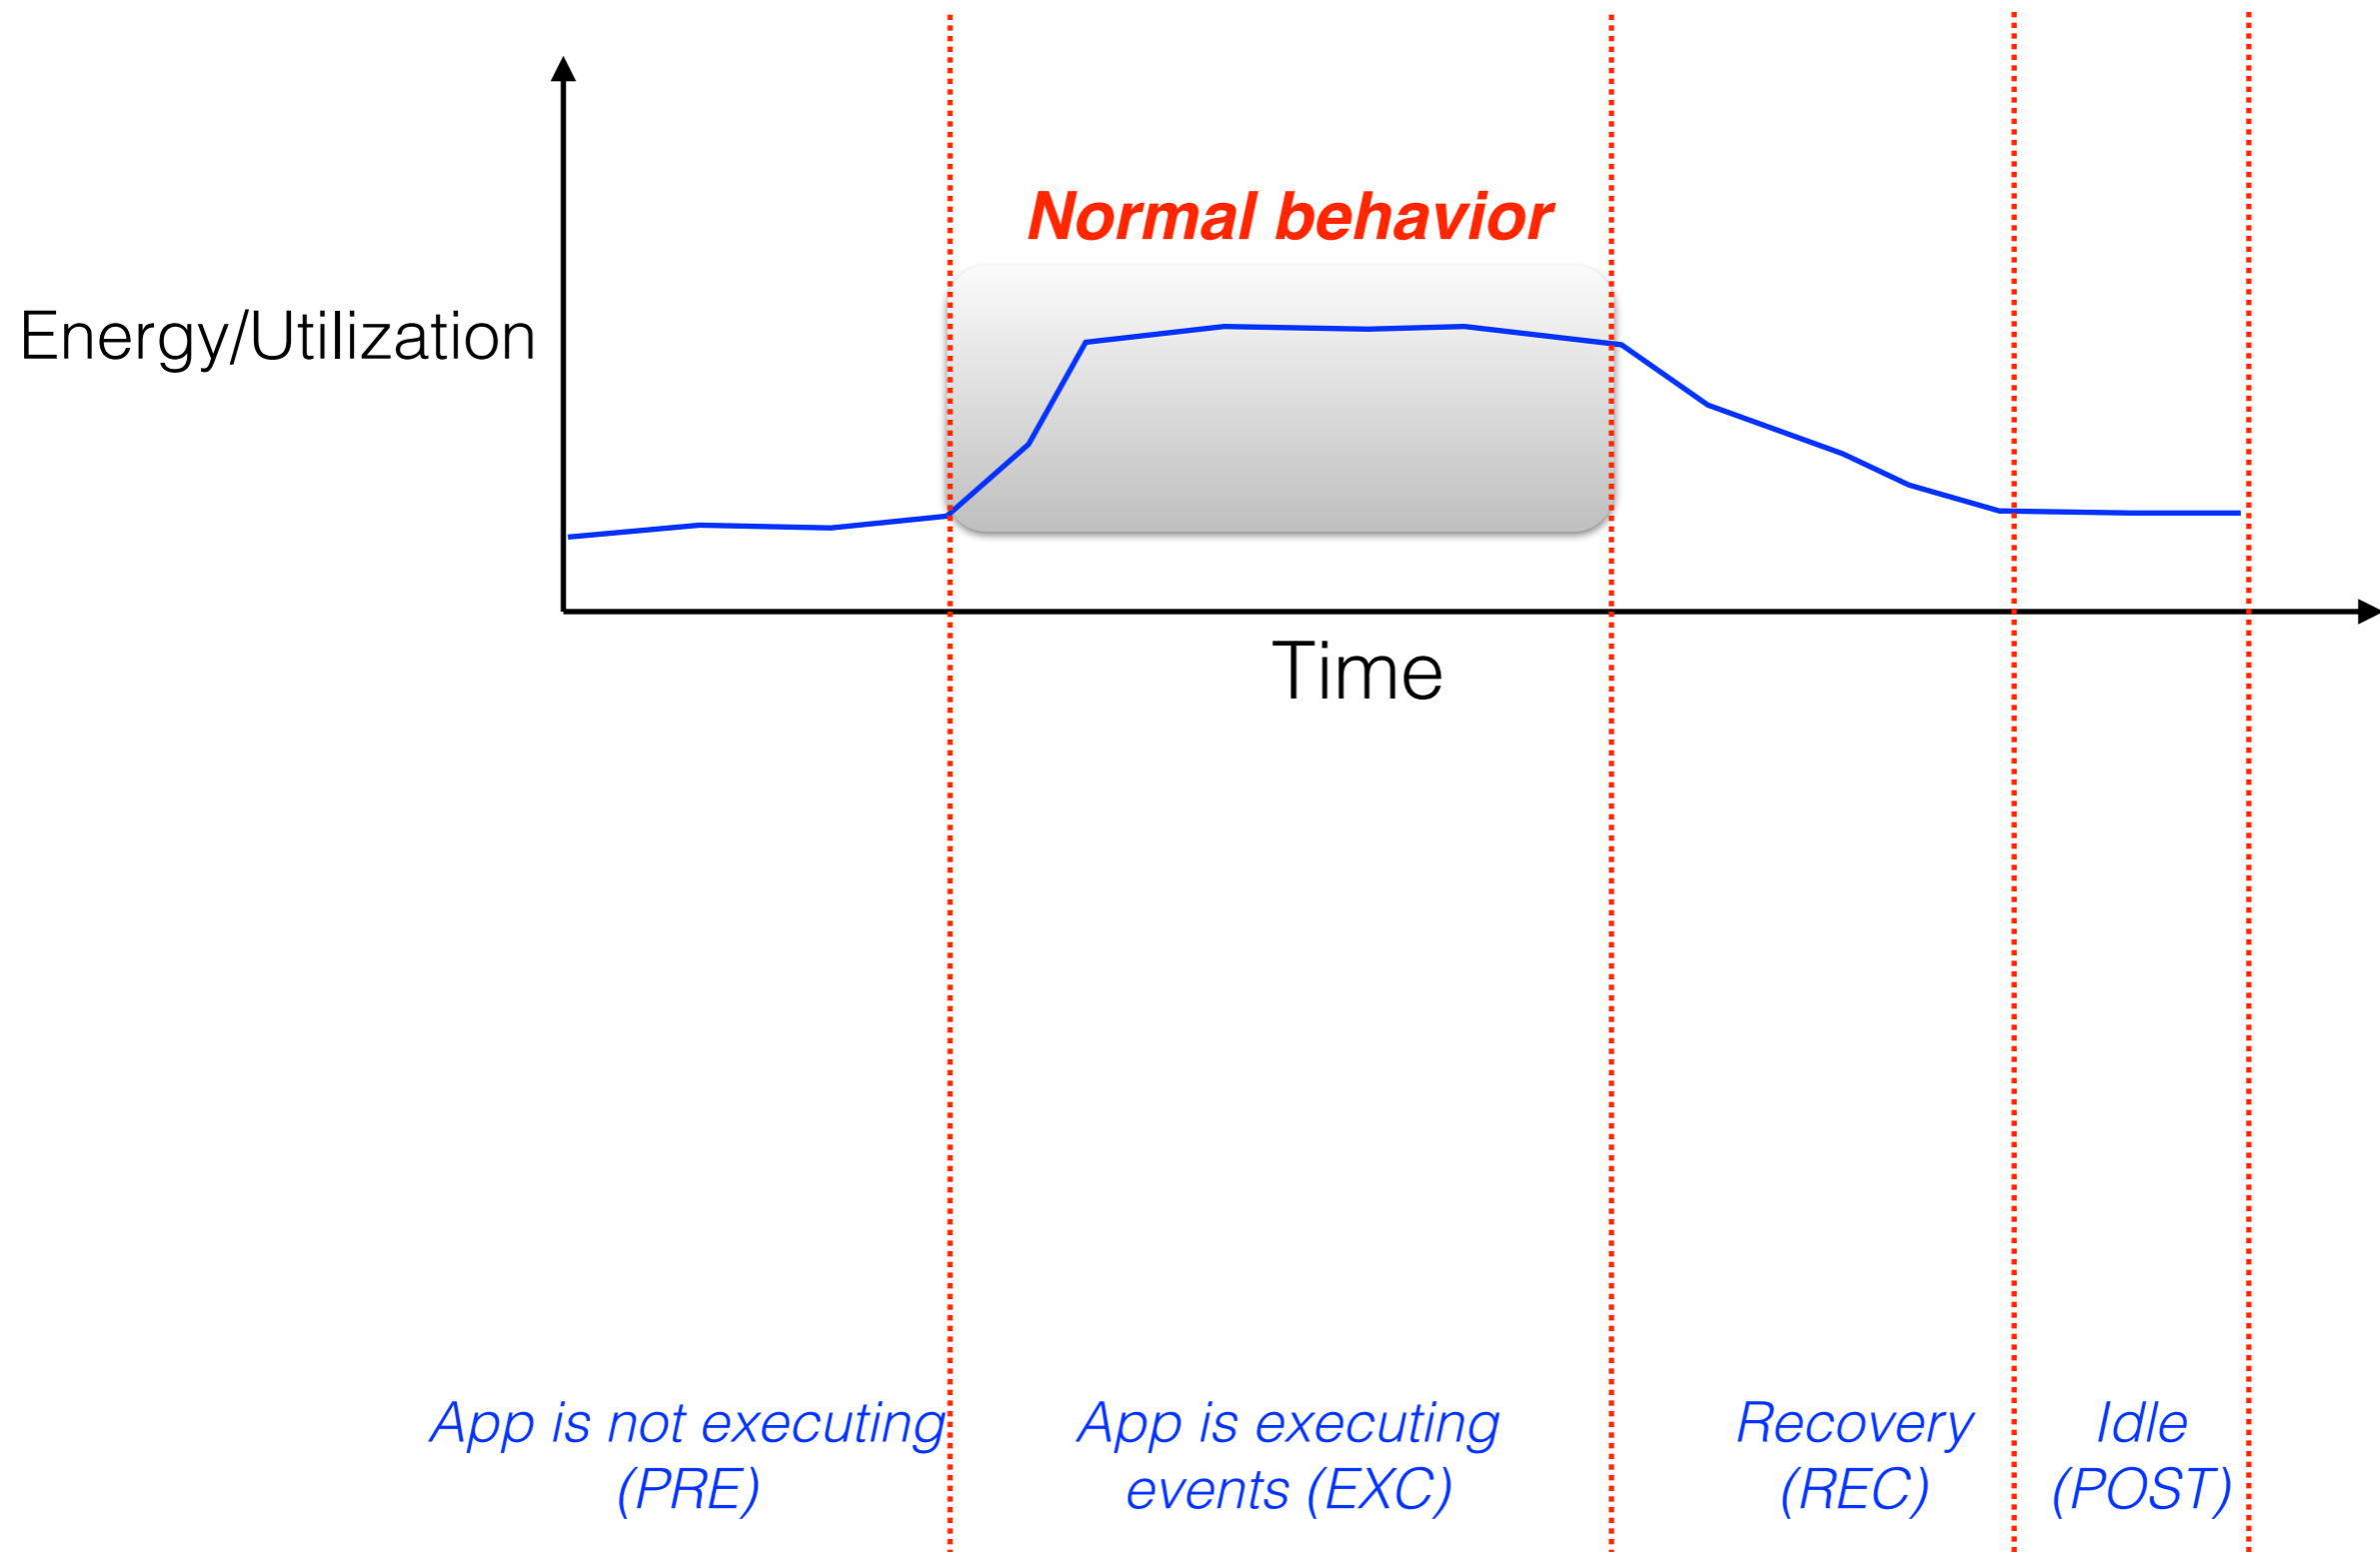
Energy Inefficiency




Energy Inefficiency



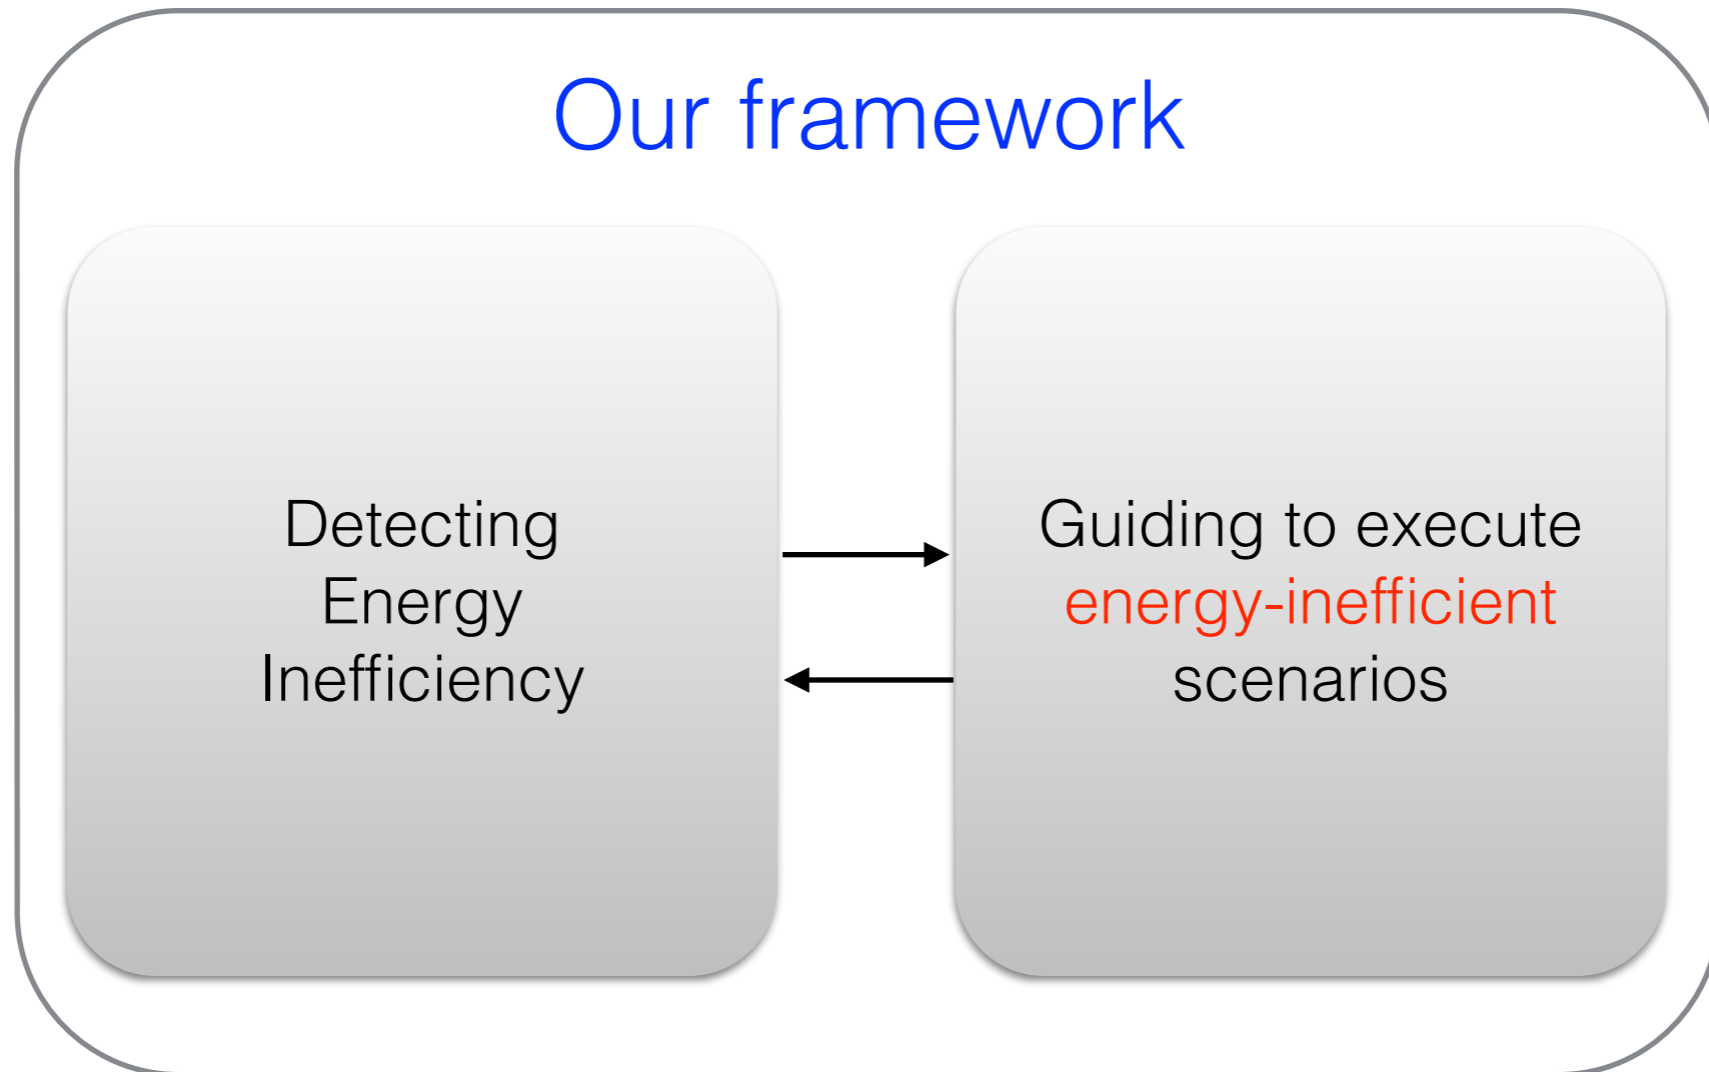
Energy Inefficiency



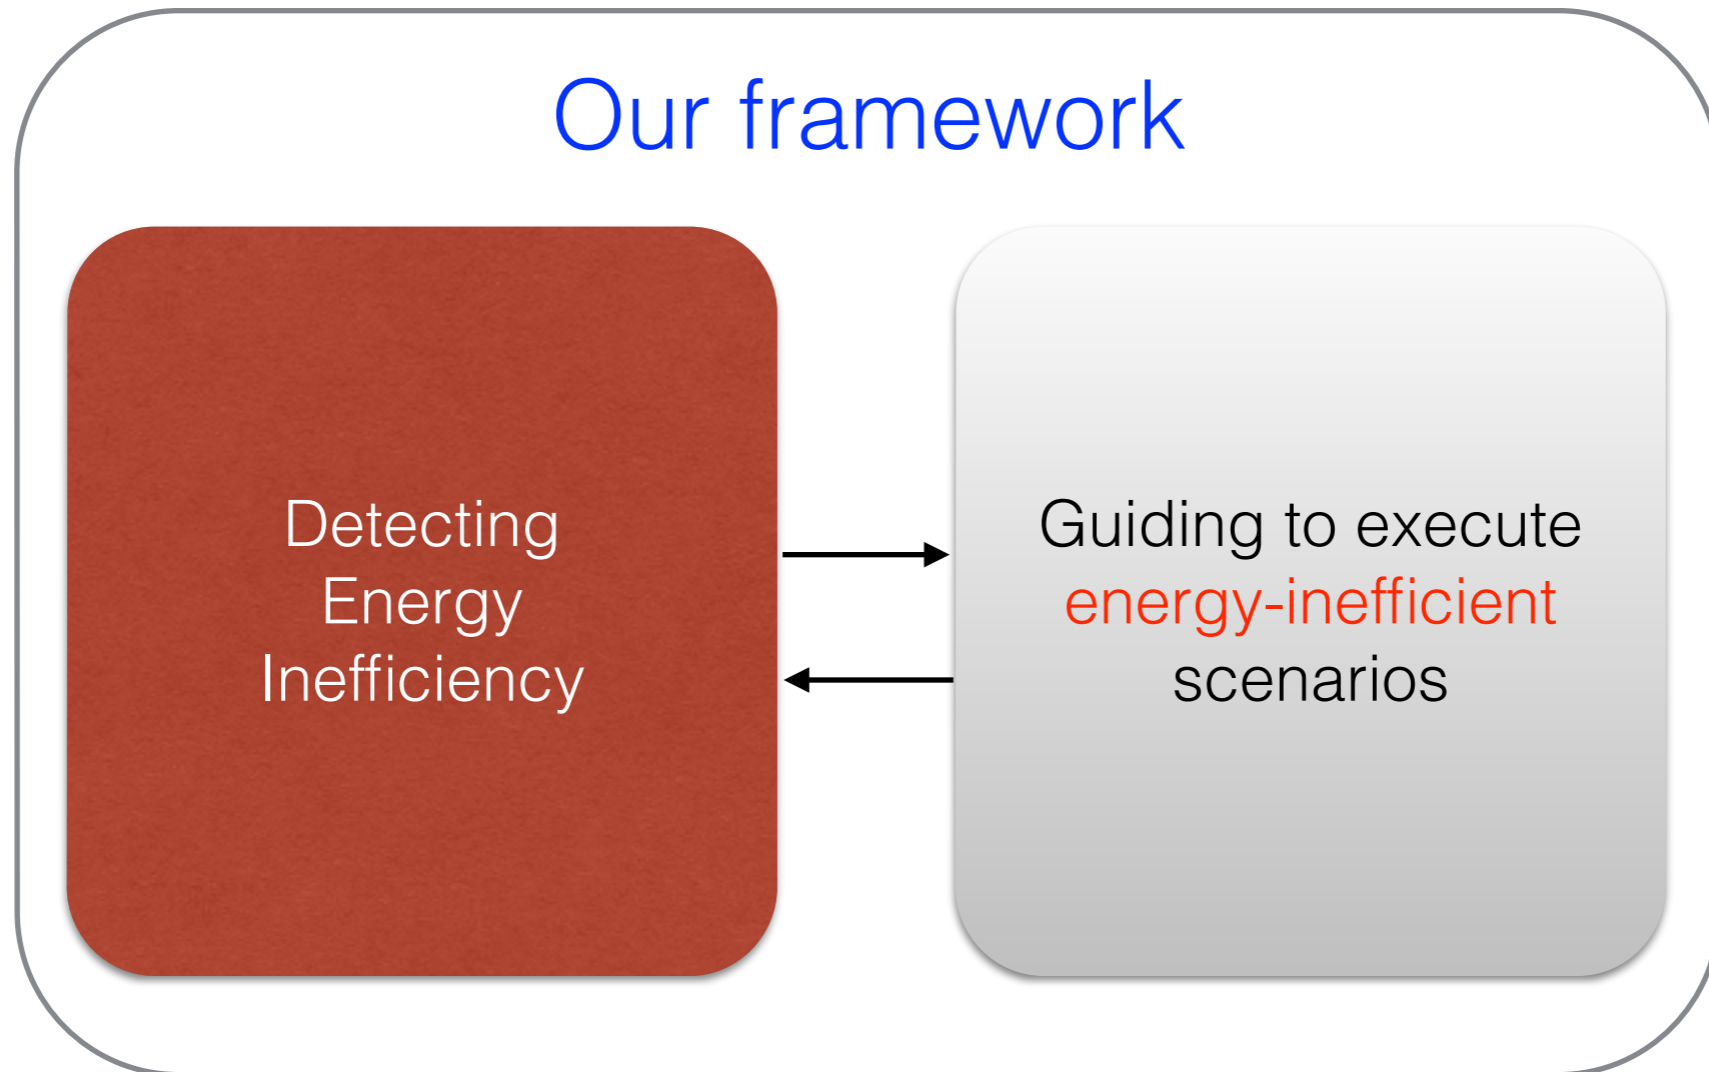
Energy Inefficiency



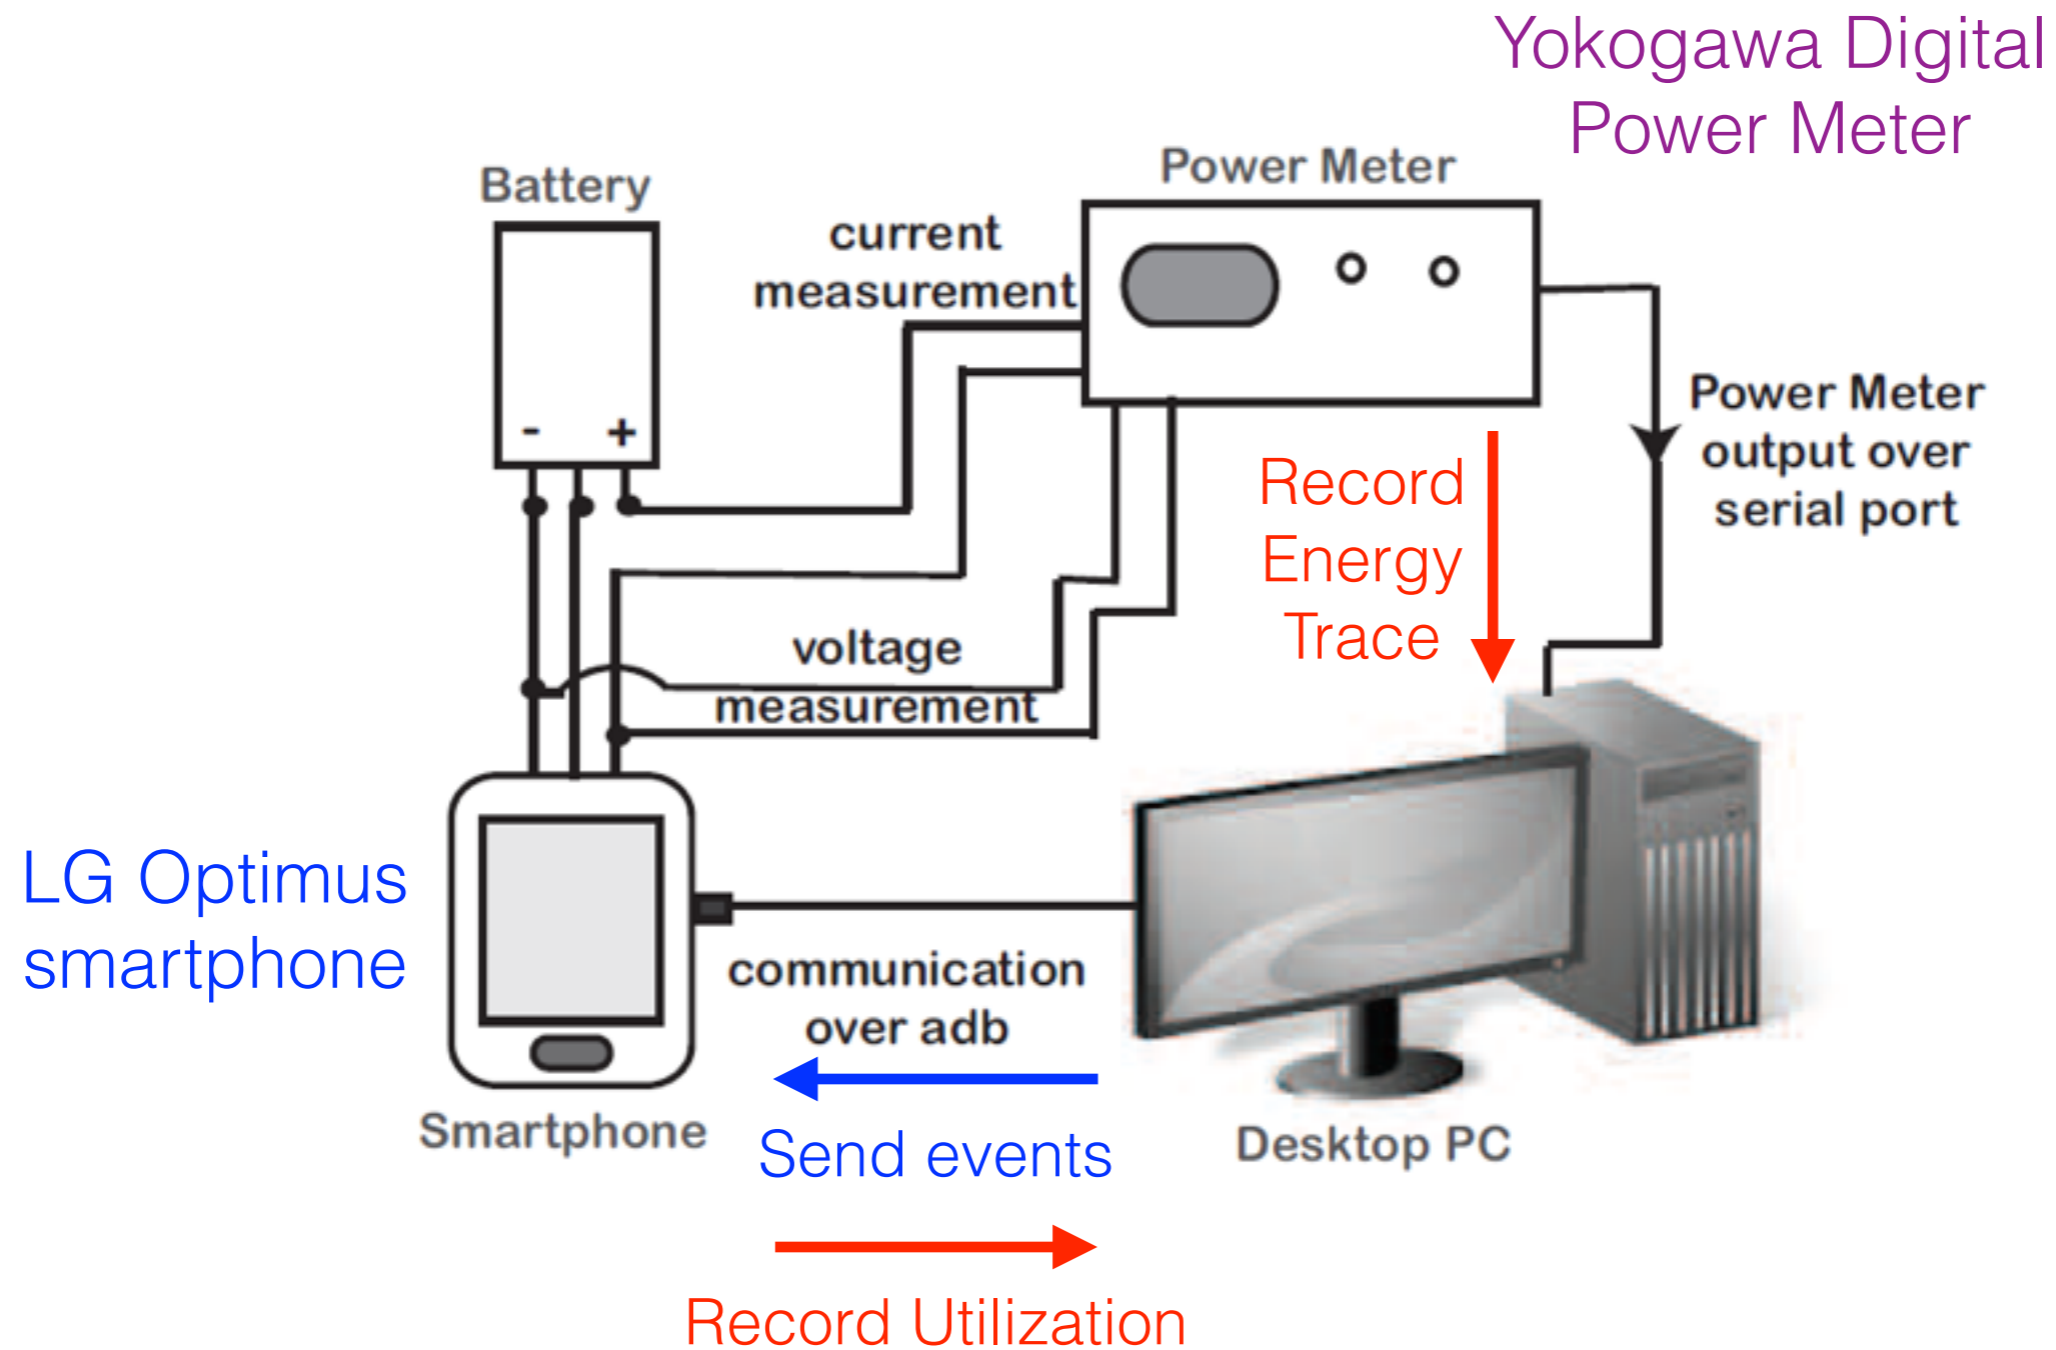
Test Generation



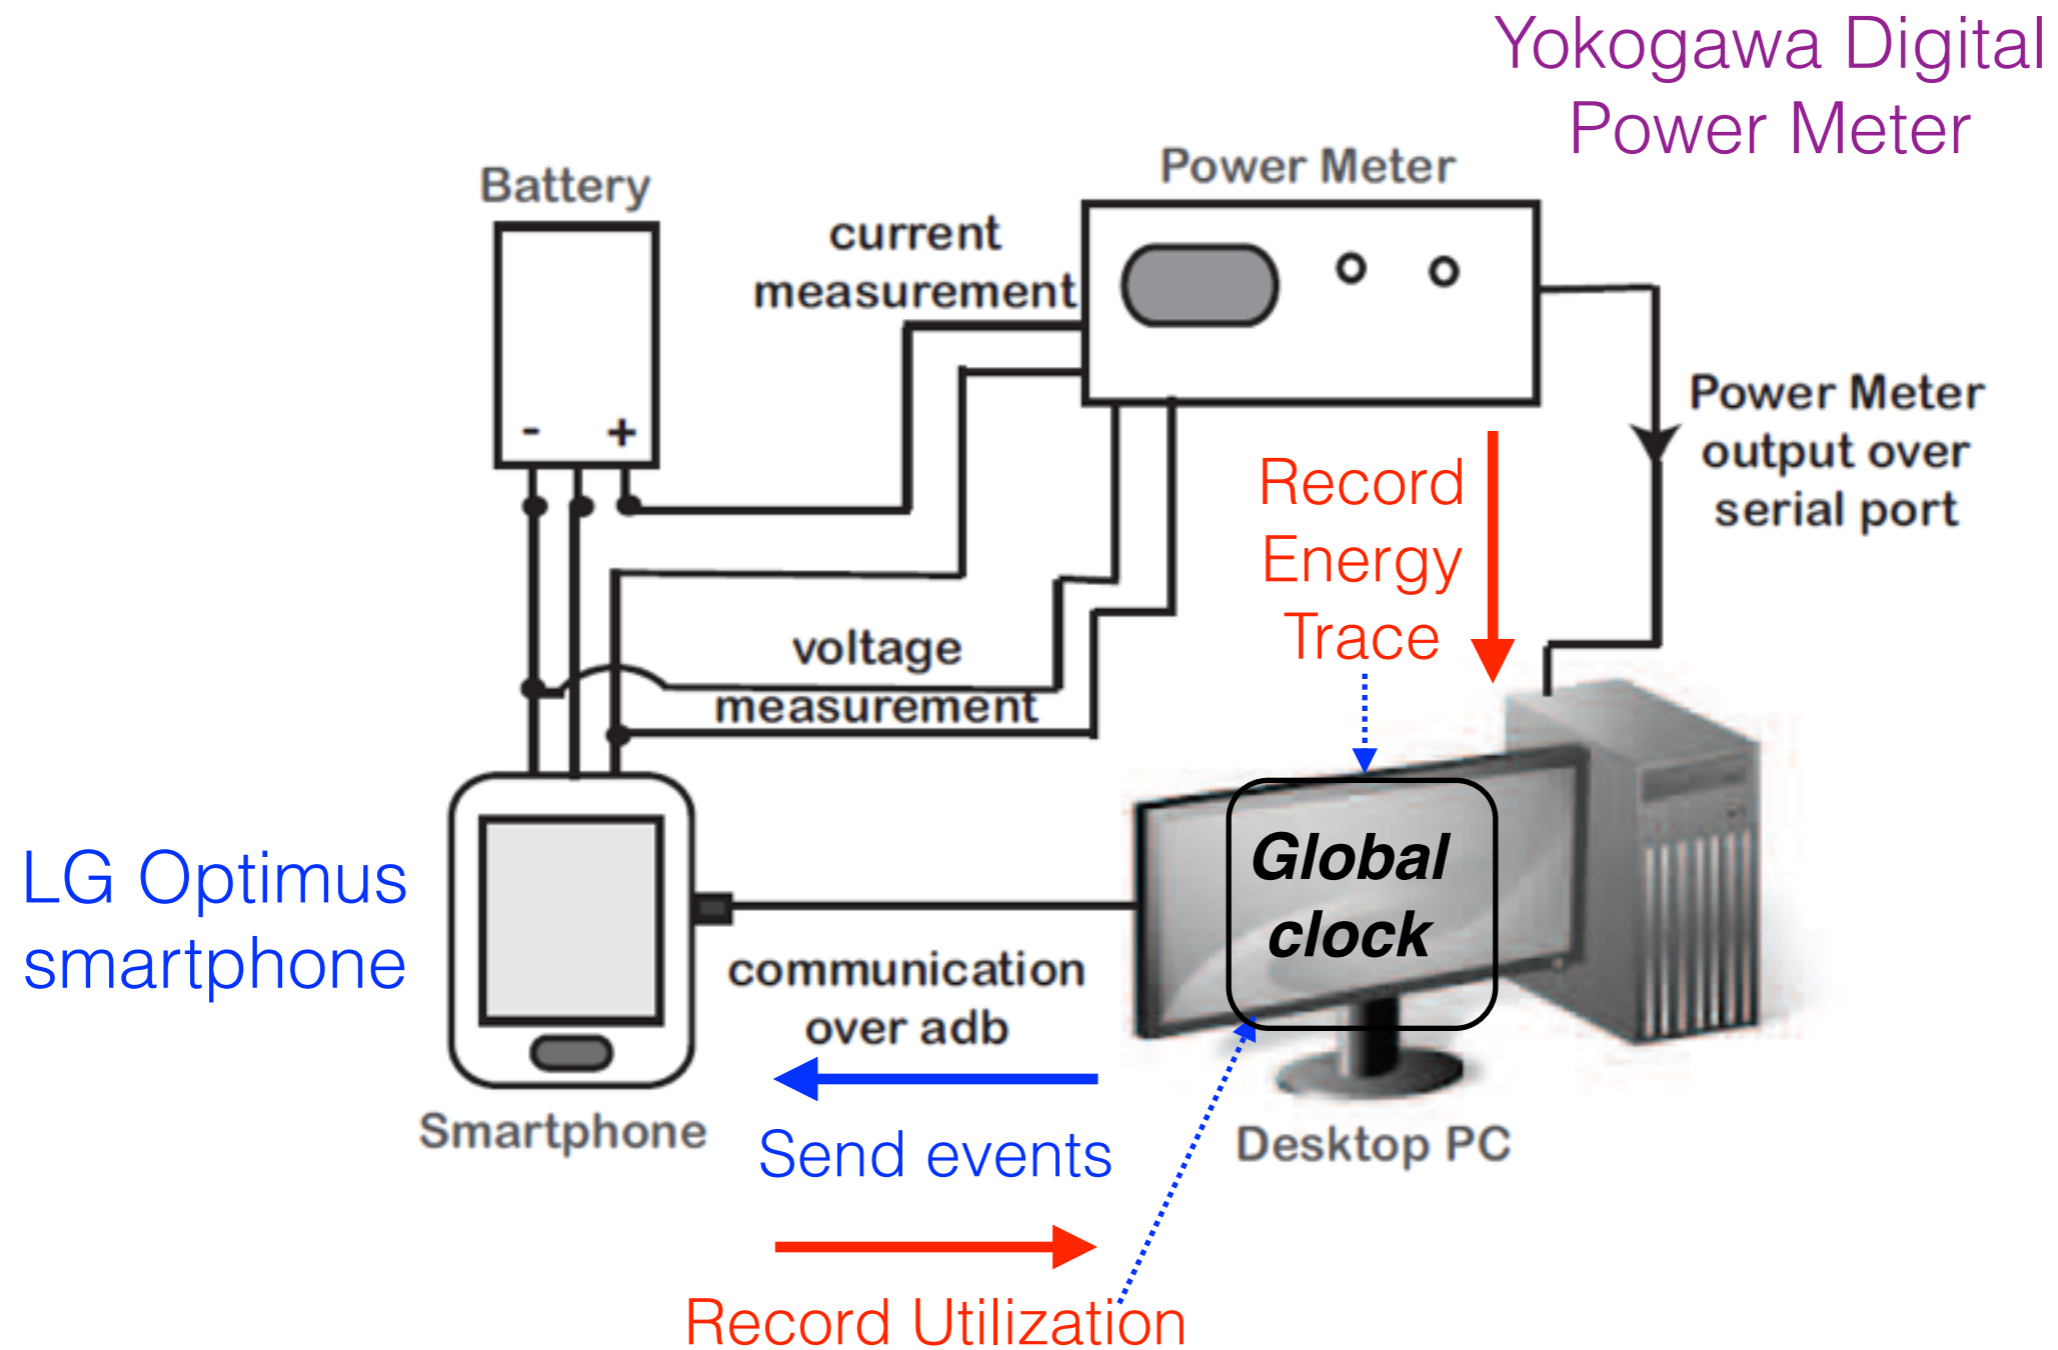
Test Generation



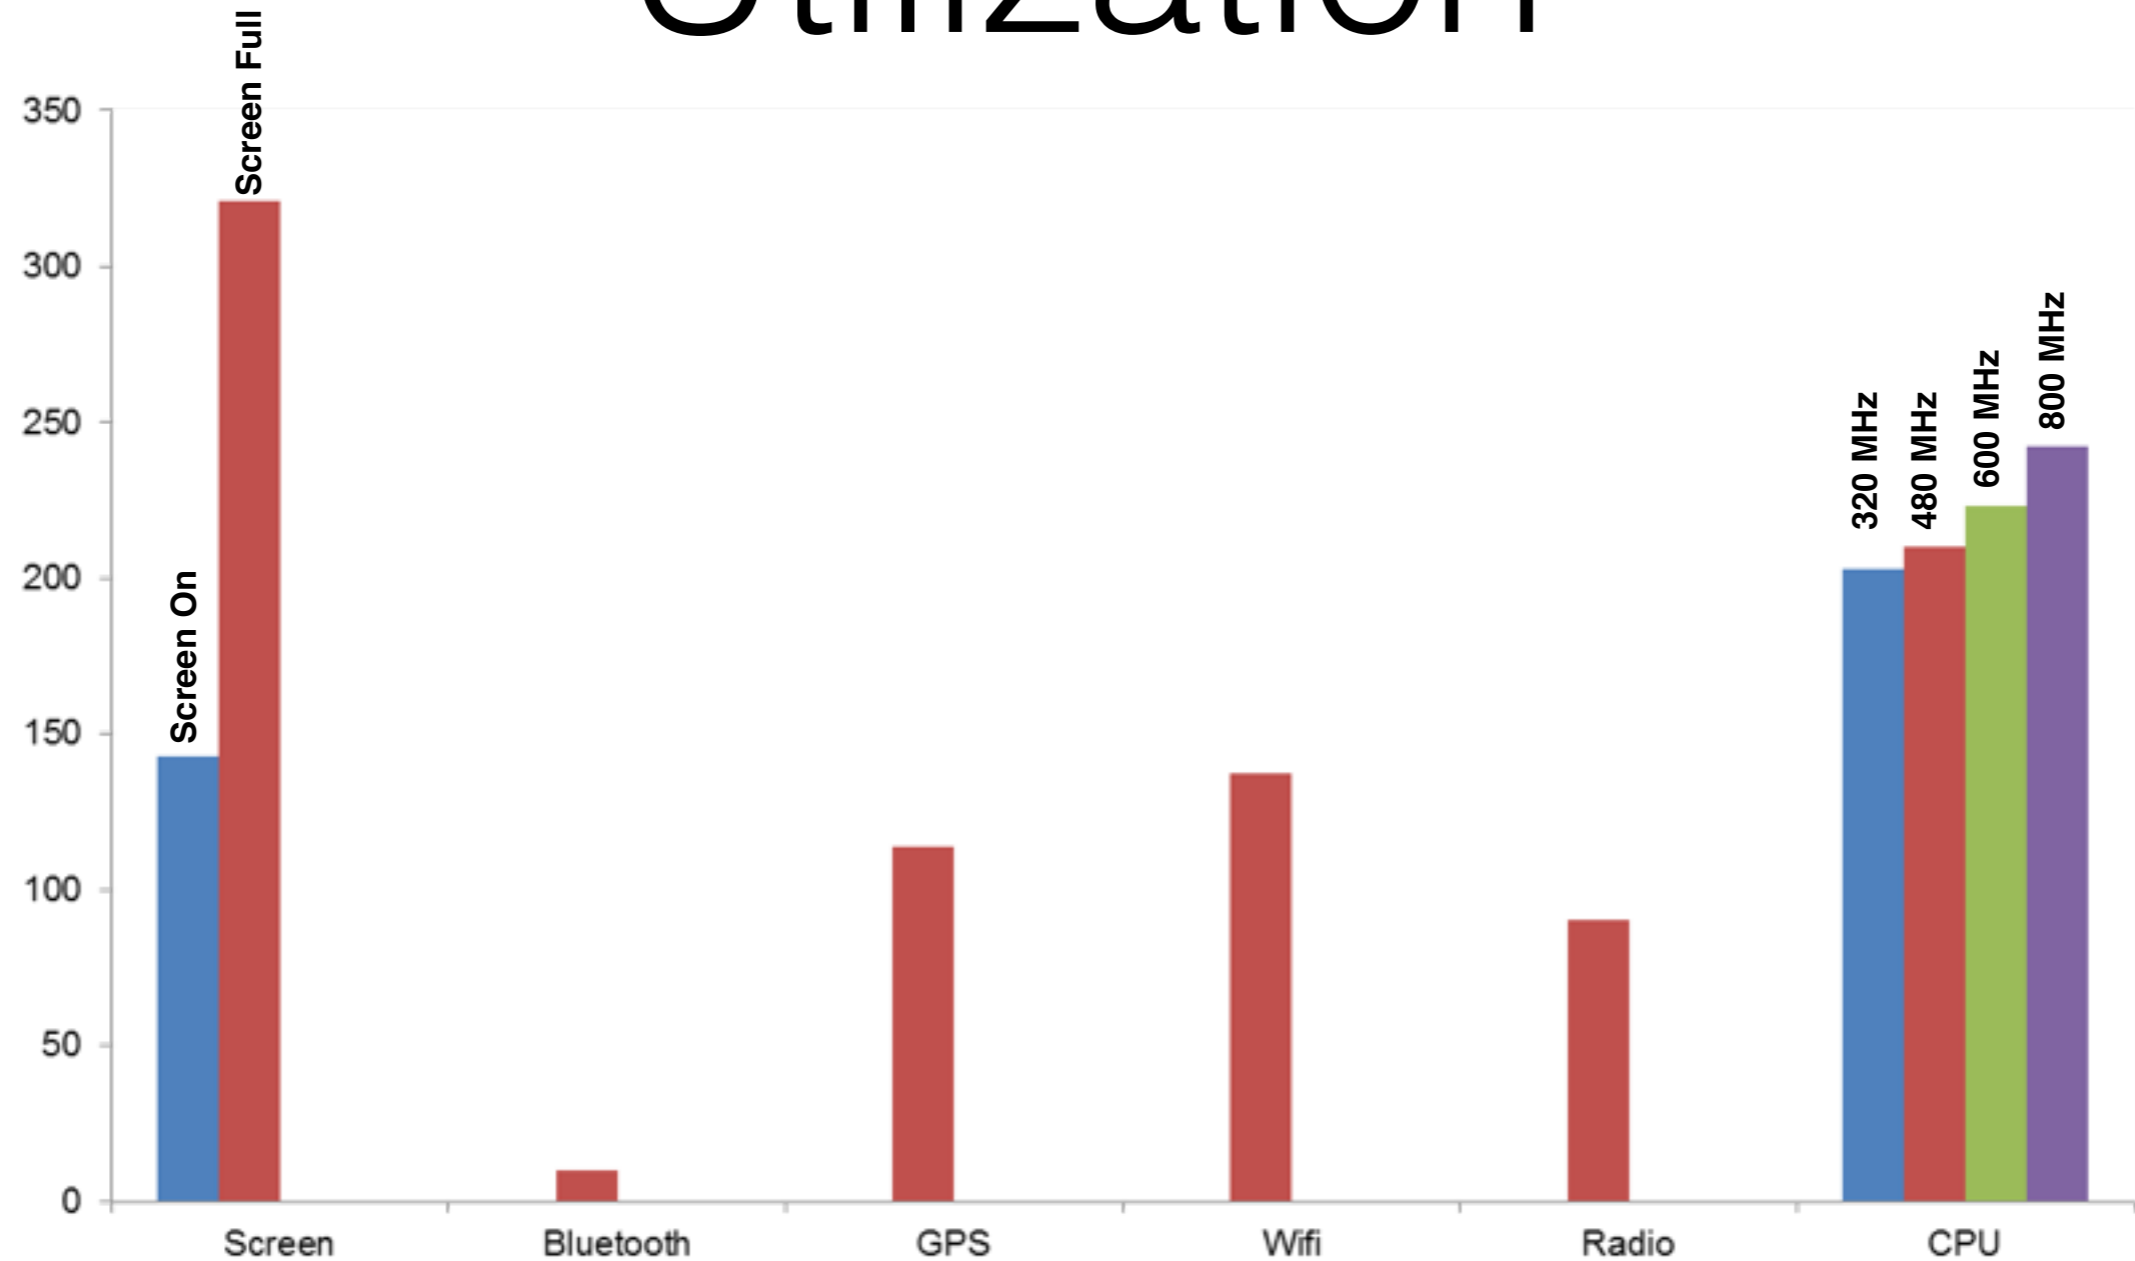
Measurement



Measurement

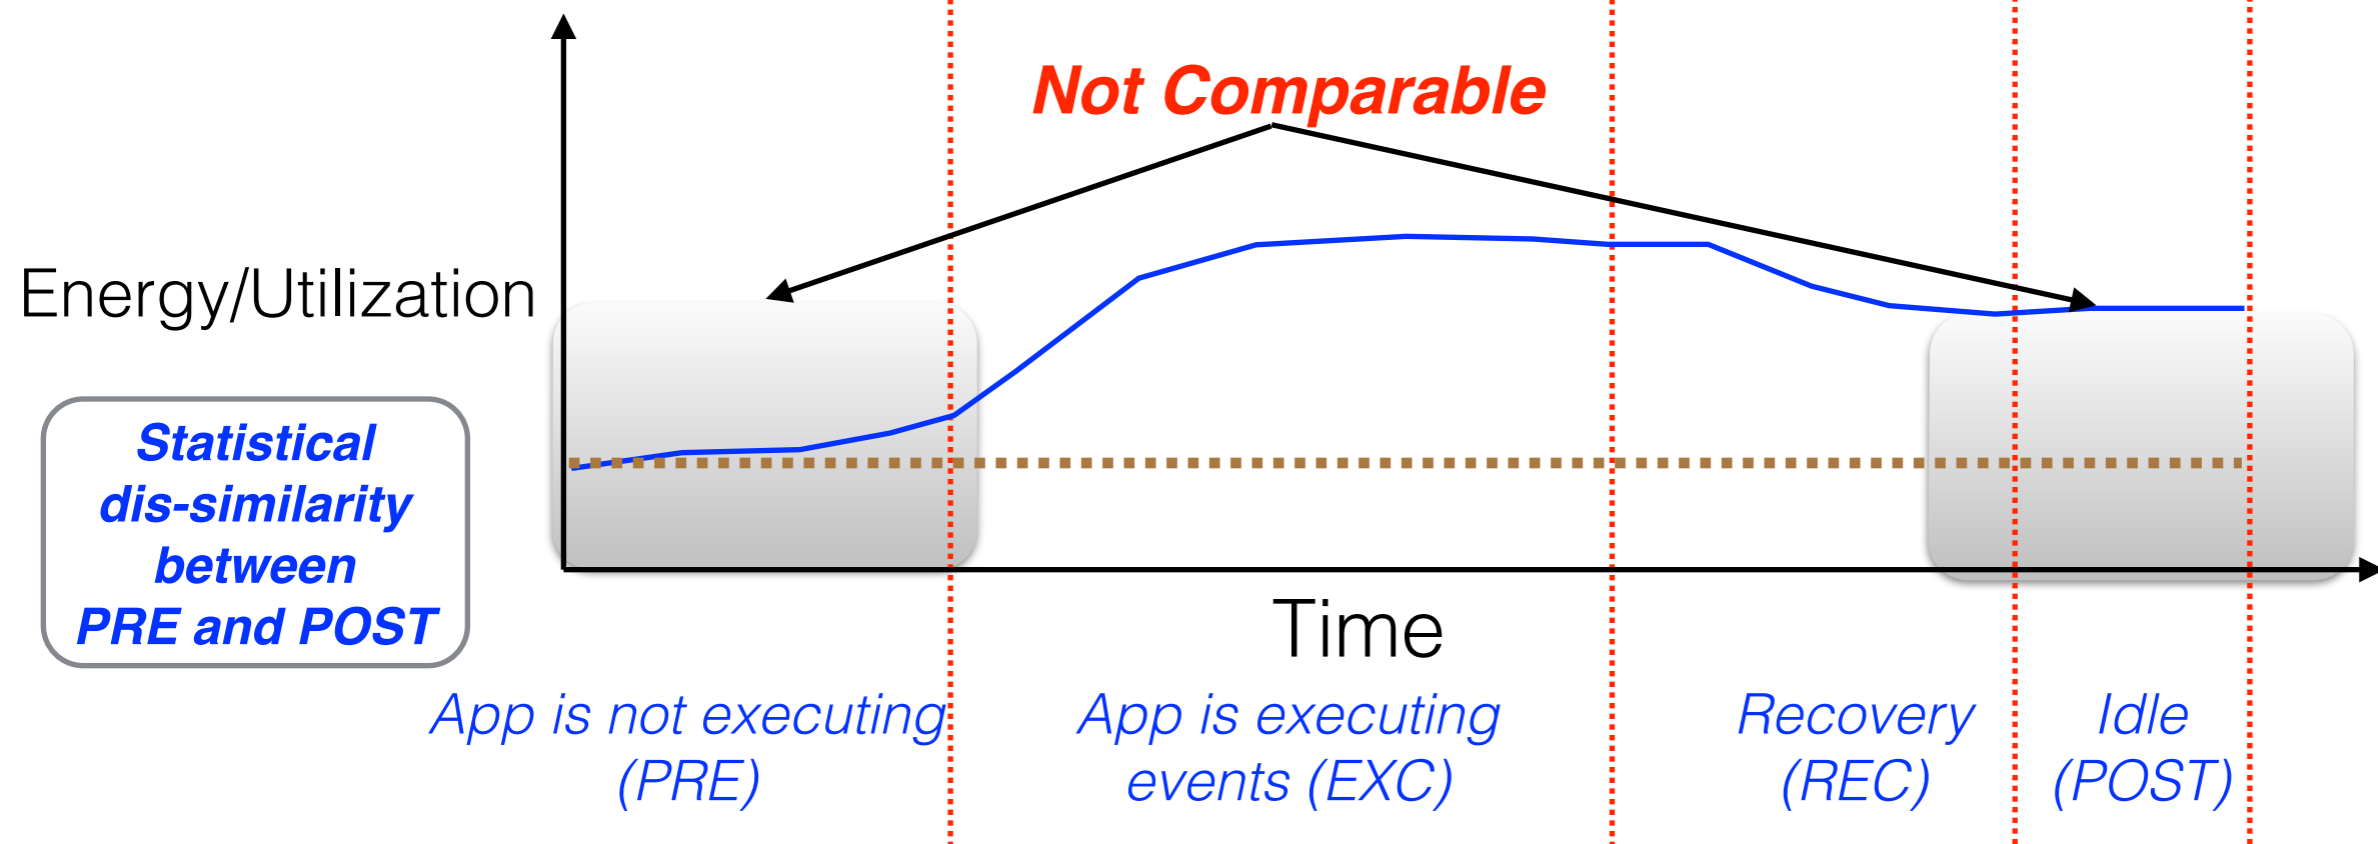



Utilization




Energy consumption of different components is not even (GPS < CPU)
100% CPU does not consume same energy as GPS being on

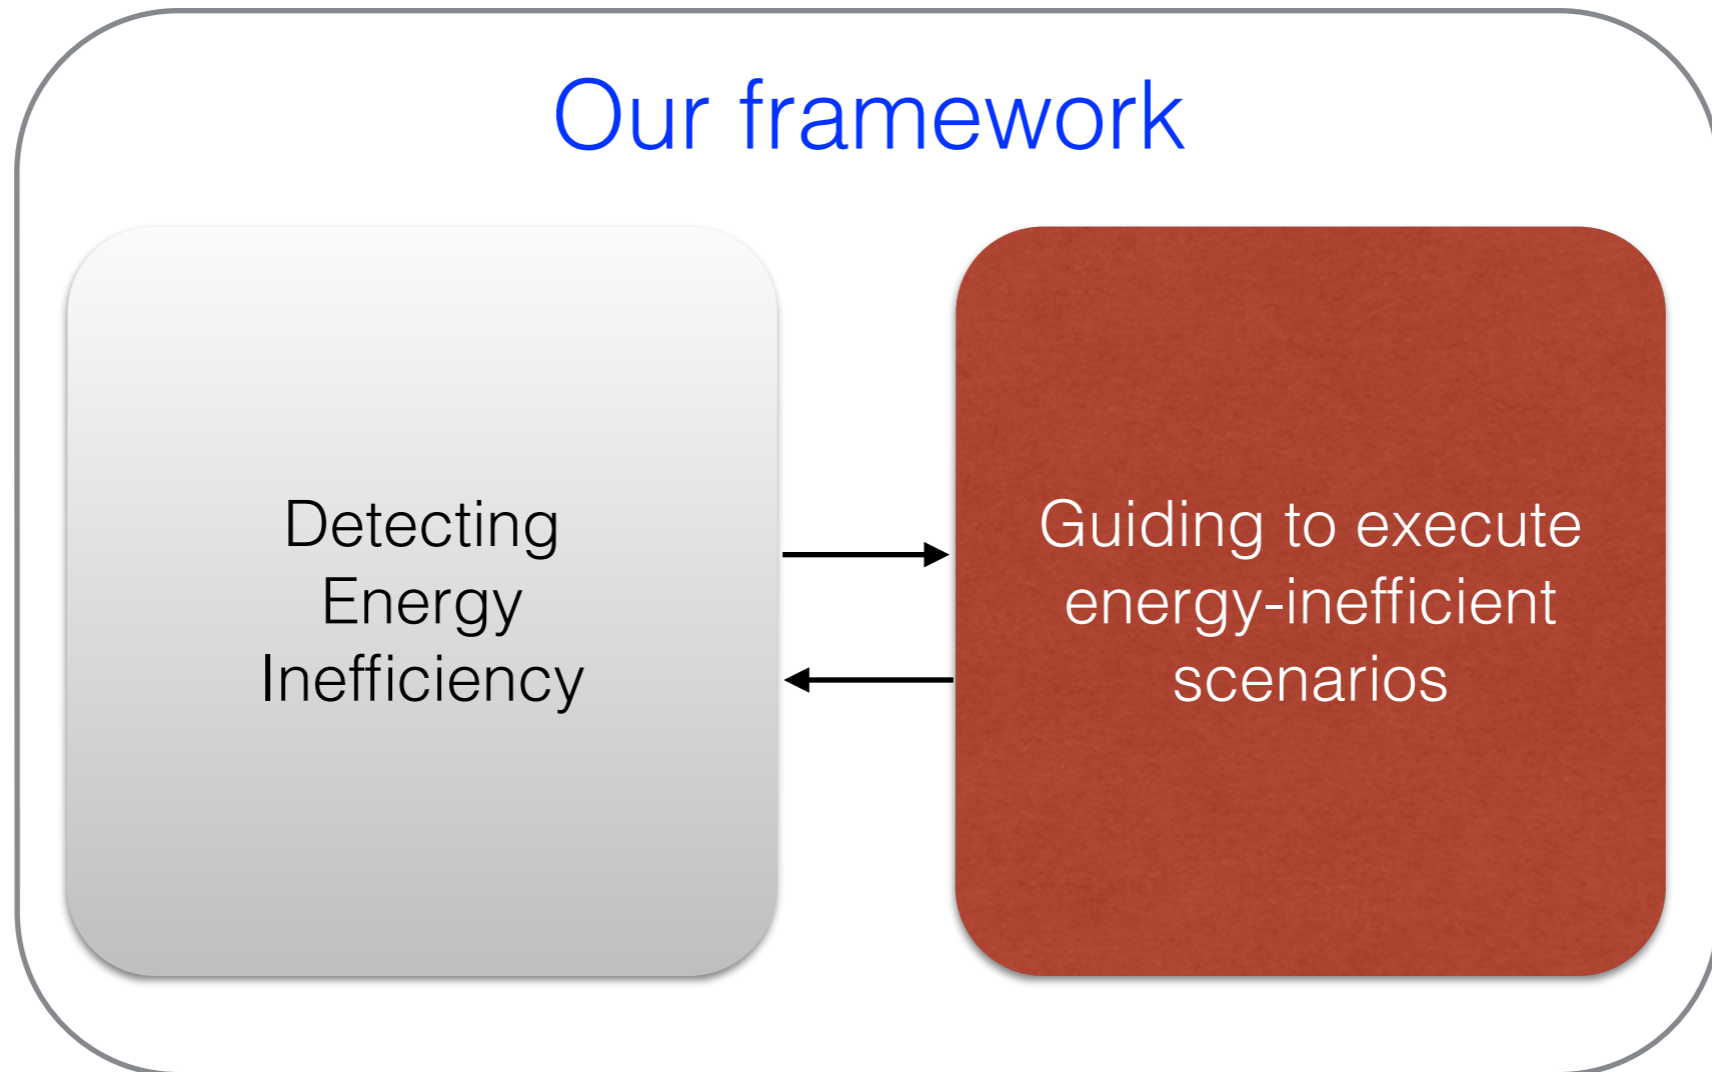
Detecting Energy Bugs



Detecting Energy Hotspots



Test Generation



Guided Exploration

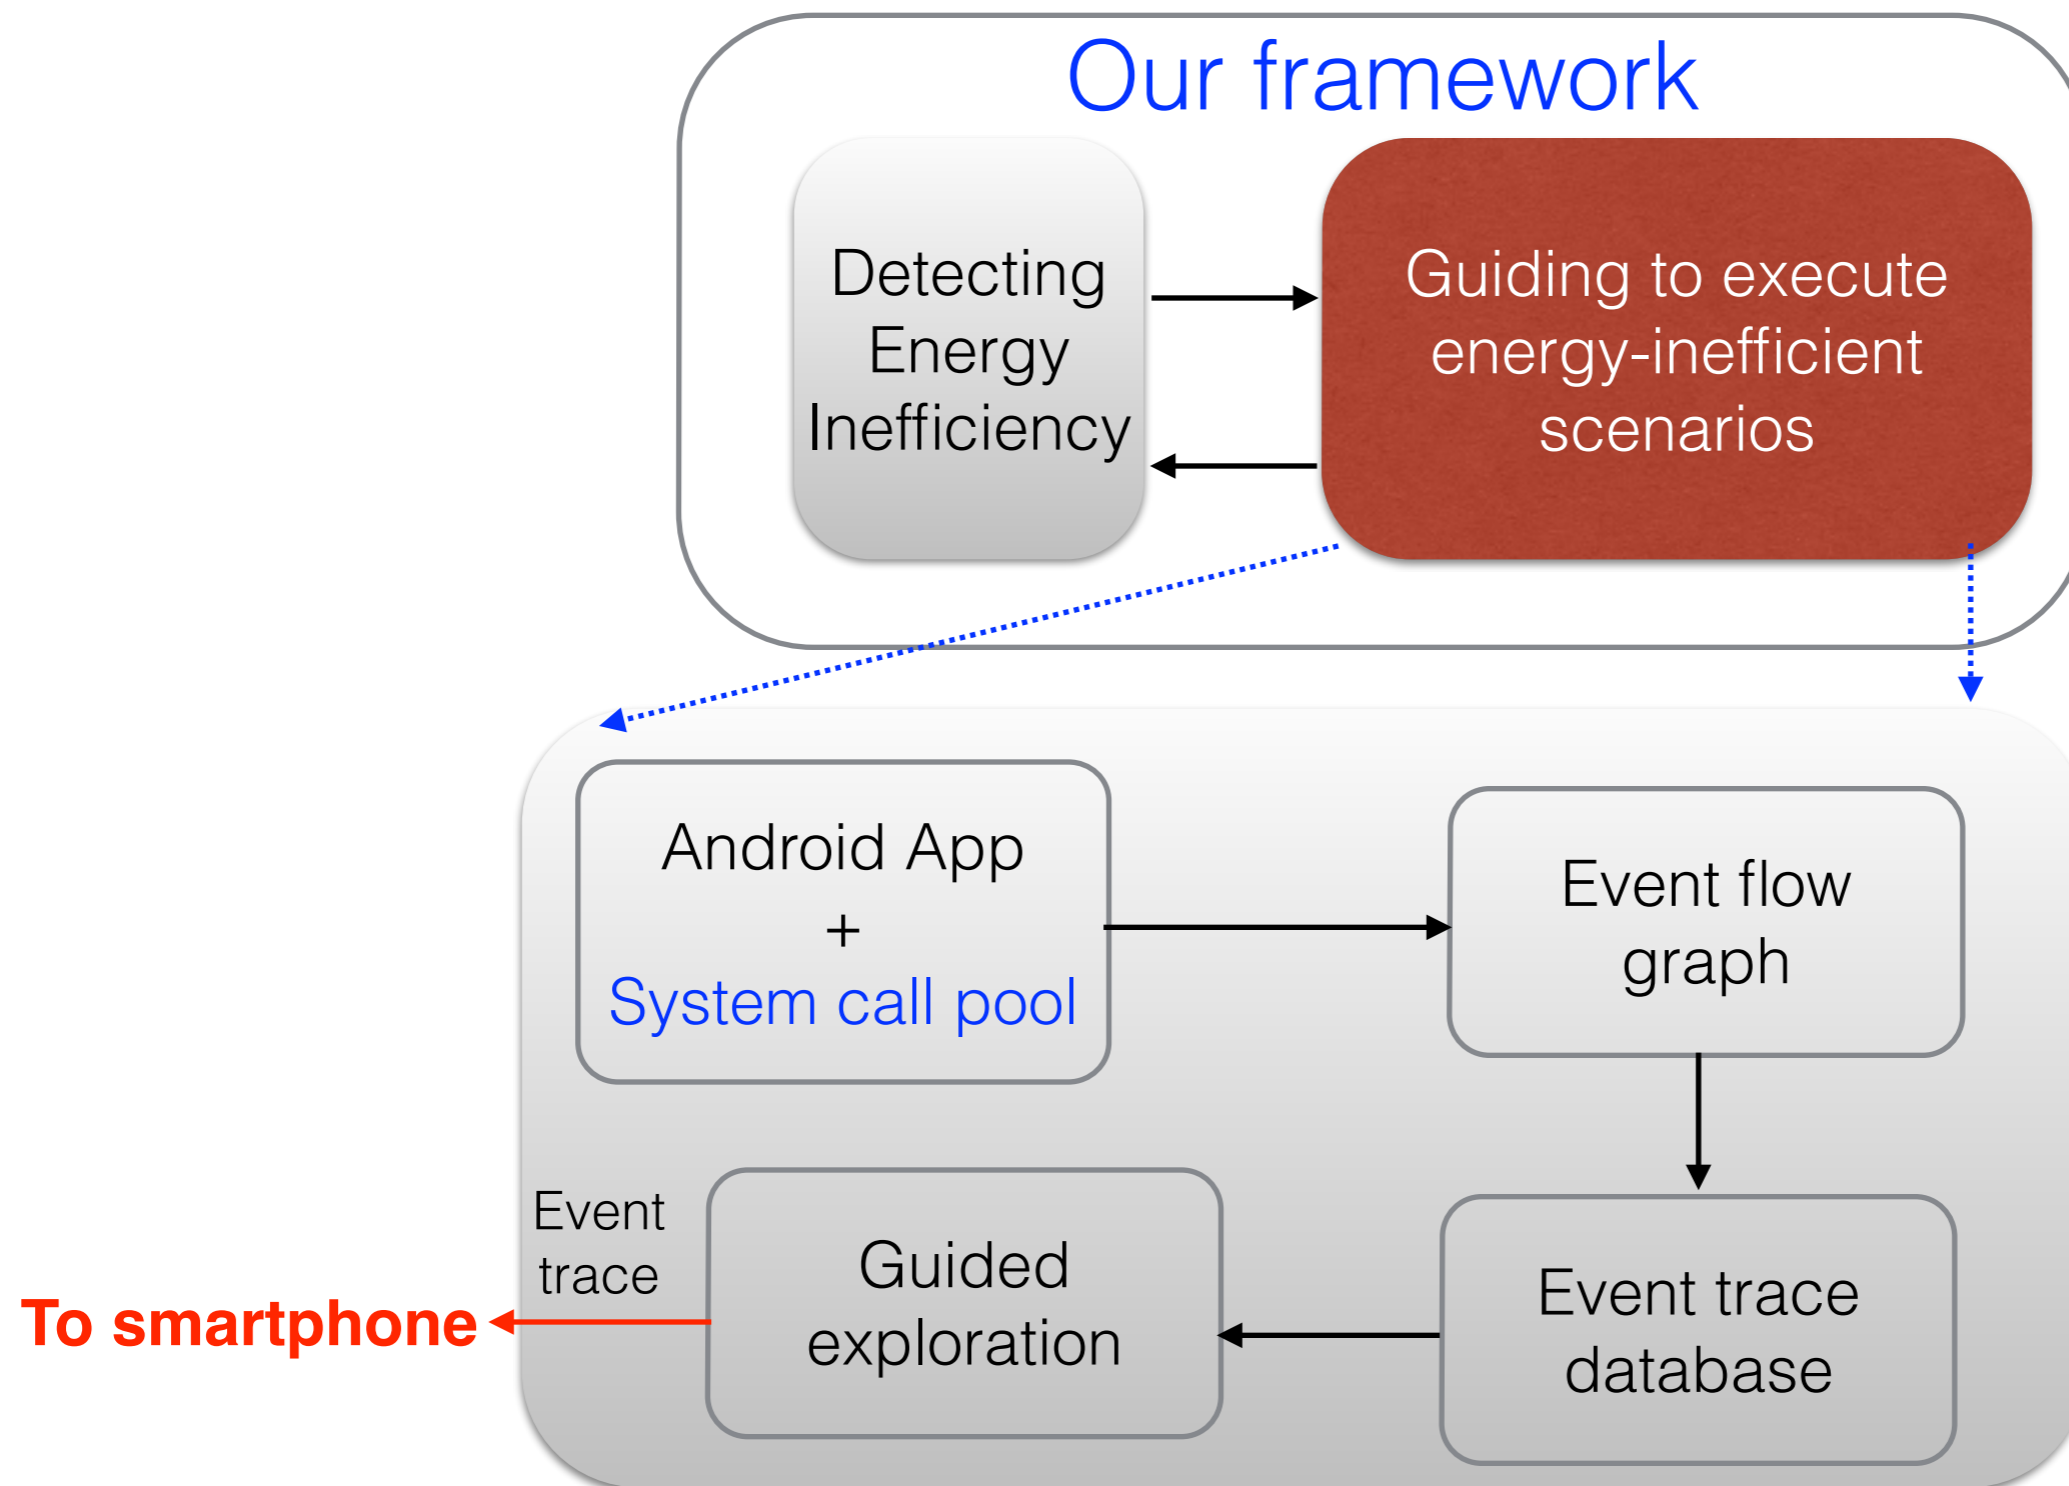
- Energy-inefficient execution
 - Which fragments are energy-inefficient?
 - What is an appropriate coverage metric?

A Broader Categorization

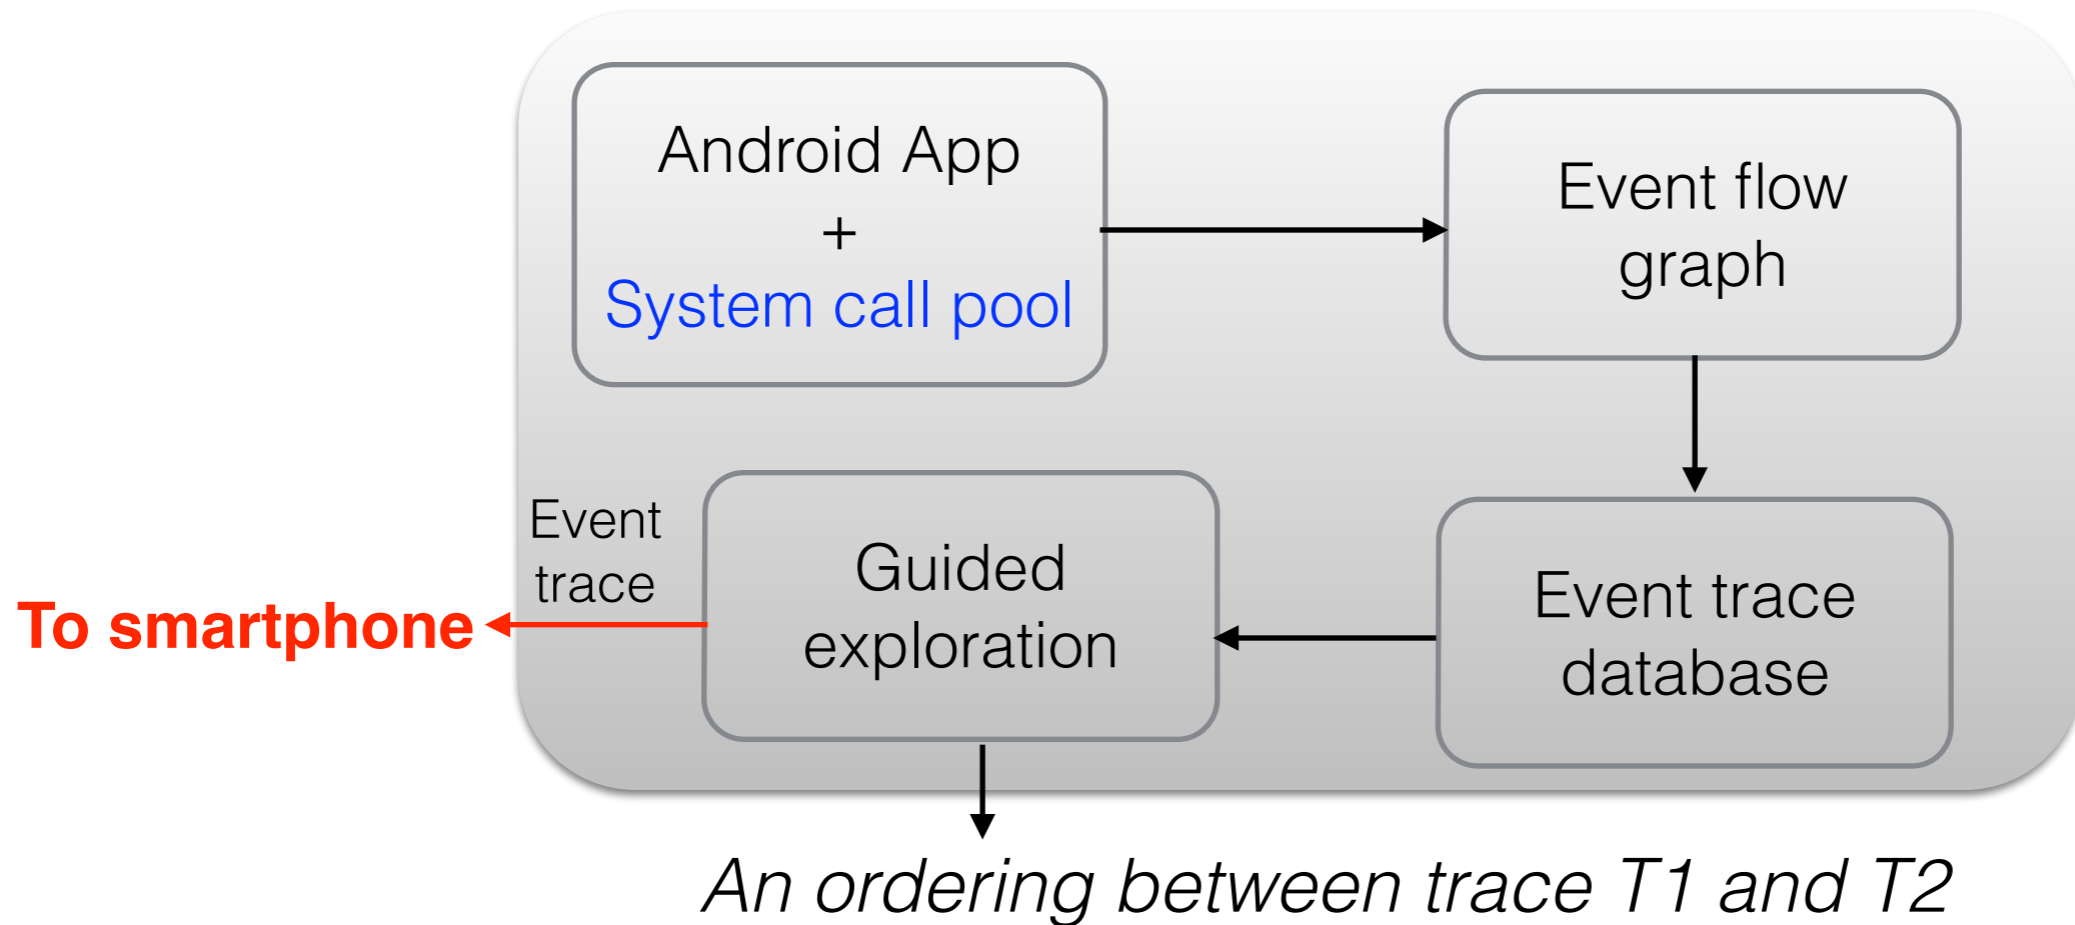
Cause/Source	Energy Bugs	Energy Hotspots
Hardware components	Resource leak	Suboptimal resource binding
Sleep state transition	Wakelock bug	Tail Energy hotspot
Background Service	Vacuous background service	Expensive background service
Defective Functionality	Immortality bug	Loop energy hotspot

Invoked via System Calls


Test Generation



Test Generation




$T1 > T2$




More system calls

Buggy trace $T1 > T2$



Similarity with buggy trace


Executed trace $T1 < T2$



Covering more system calls

Evaluation

- Travel/Transportation
- Lifestyle/Health
- Books/News
- Productivity
- Photography/Media
- Entertainment
- Puzzle
- Tools



Category of Android Apps Evaluated

Summary of Evaluation

App	Feasible traces	Energy Bugs	Energy Hotspots	Type	Reported before
Aripuca	502	Yes	No	Vacuous background service	No
Montreal Transit	64	No	Yes	Suboptimal resource binding and more	No
Sensor Test	2800	Yes	No	Immortality Bug	No
760 KFMB AM	26	Yes	Yes	Vacuous background service, suboptimal resource binding	No


*All Results are in the paper
(10 energy bugs and 3 energy hotspots found
out of 30 tested apps)*

Summary of Evaluation

App	System call	Statement	Lines of Code
Aagtl	100	21	11612
Android Battery Dog	100	17	463
Aripuca	100	15	4353
Kitchen Timer	100	30	1101
Montreal Transit	89	11	10925
NPR News	100	24	6513
OmniDroid	83	36	6130
Pedometer	100	56	849
Vanilla Music Player	86	20	4081

*To cover all system calls, exploring only a small part of the program suffices
A substantial portion of the code is used for provide user feedback, compatibility
over different OS*

Case Studies




Aripuca (Energy Bug)

Reason: *Vacuous Background Service*

Fix:


```
serviceConnection.getService().stopLocationUpdates();  
serviceConnection.getService().stopSensorUpdates();
```

Case Studies




Aripuca (Energy Bug)

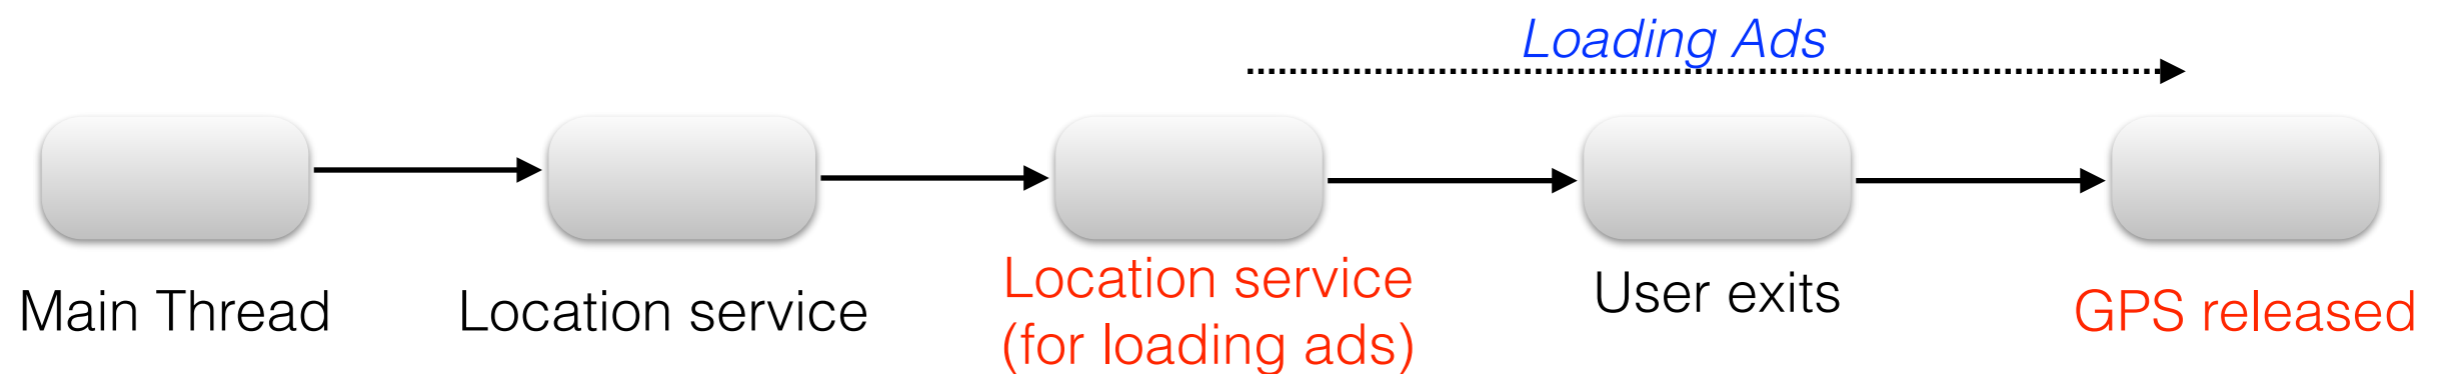
Reason: Vacuous Background Service




Case Studies



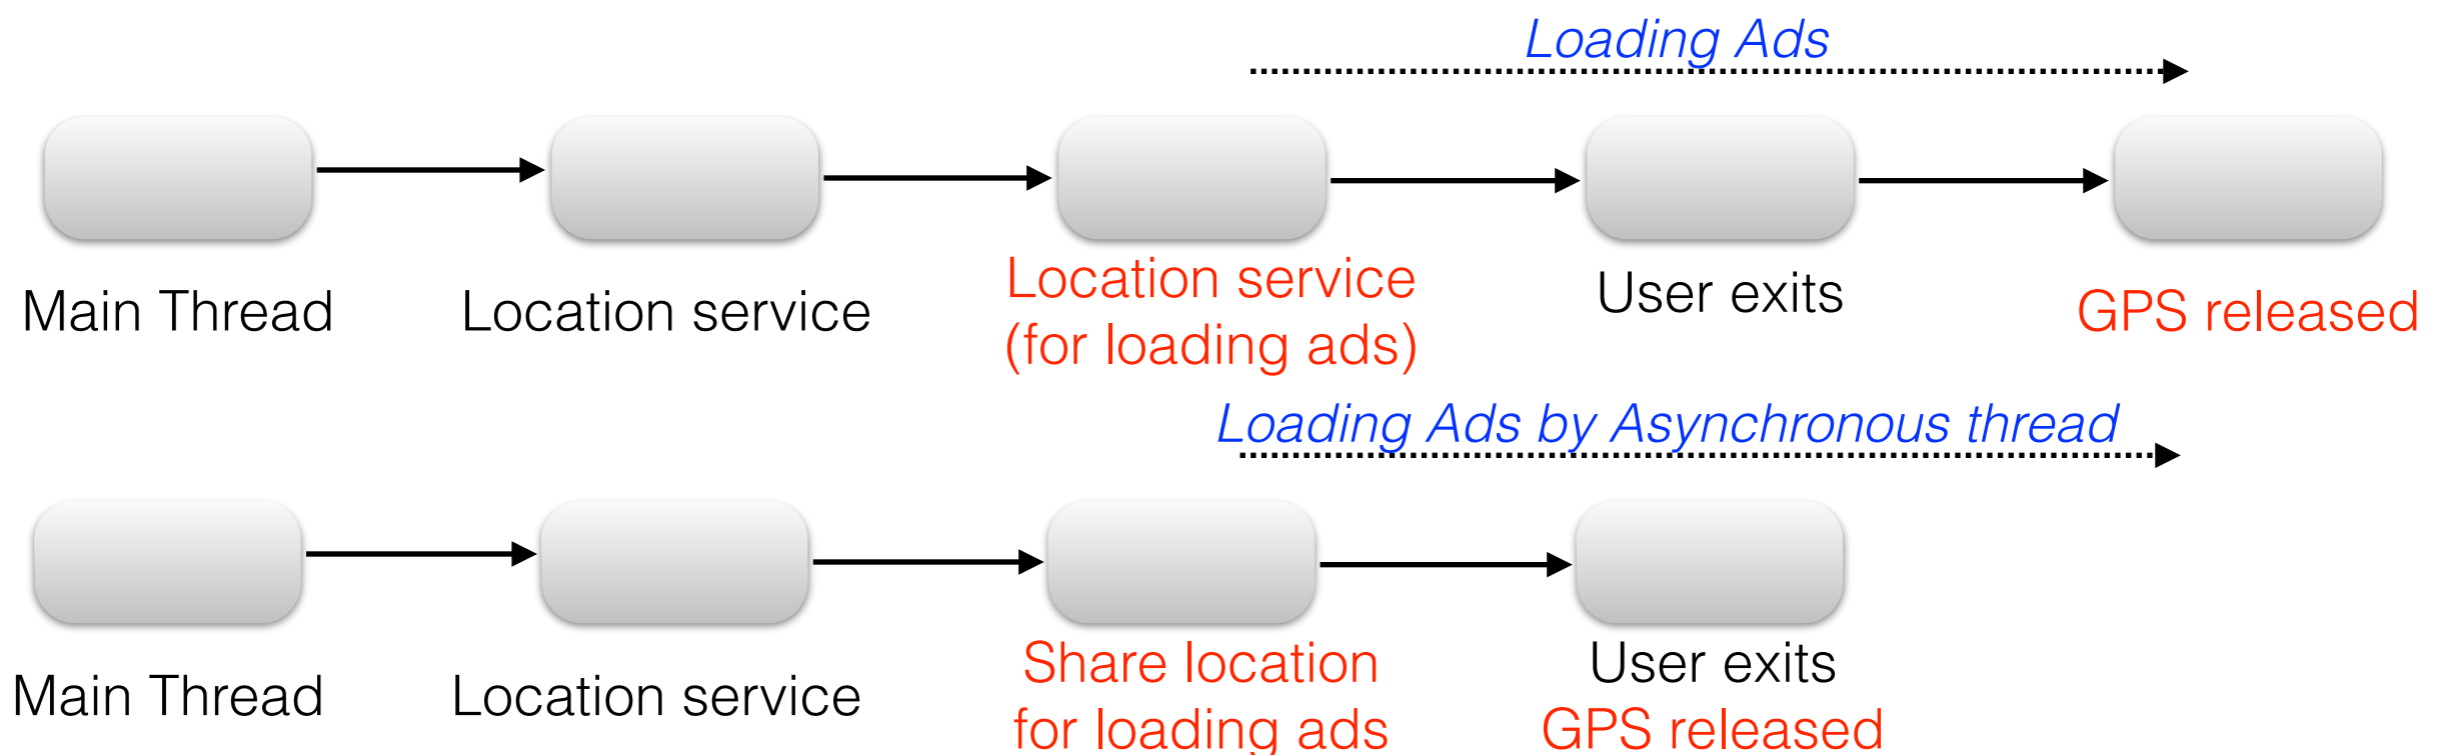
Montreal Transit (Energy hotspot for <5 sec)




Case Studies




Montreal Transit (Energy hotspot for <5 sec)



Case Studies



Montreal Transit (Energy hotspot for <5 sec)



Montreal Transit (After fixing)

Summary

- Categorization of energy inefficiency
 - Energy bugs
 - Energy hotspots
- A guided exploration of event traces
 - Targeting system call coverage
- Evaluation with Android apps
 - Energy bugs and hotspots exist in several Android apps


Is That All?

- Testing Non-functional Software Properties
 - Performance Testing
 - Energy Testing
- Far from being solved
 - Fresh look on the formal foundation of software testing
 - Automated debugging and fault localization
- *Information leaks via side channels (time, cache miss, power)*

Cache Side Channel

```
encrypt (message, key) {  
  :: :: :: :: :: :: :: :: :: :: :: :: :: :: :: ::  
  mid[0][0] = lookup[key[0]][0];  
  mid[0][0] ^= lookup[key[1]][1];  
  :: :: :: :: :: :: :: :: :: :: :: :: :: :: :: ::  
}
```

Intermediate state of the encrypted message




**lookup Table
(in memory)**

Cache Side Channel

```
encrypt (message, key) {  
  :: :: :: :: :: :: :: :: :: :: :: :: :: :: :: ::  
  mid[0][0] = lookup[key[0]][0];  
  mid[0][0] ^= lookup[key[1]][1];  
  :: :: :: :: :: :: :: :: :: :: :: :: :: :: :: ::  
}
```


Intermediate state of the encrypted message

*Depends on the value of **Key***



**lookup Table
(in memory)**

Cache Miss Distribution



Number of Occurrences of "N" Cache Misses (for different inputs)

- Abhijeet Banerjee, Sudipta Chattopadhyay and Abhik Roychoudhury. *Static Analysis Driven Cache Performance Testing*. IEEE Real-time Systems Symposium (**RTSS**), 2013 **Best Paper Candidate**
- Abhijeet Banerjee, Lee Kee Chong, Sudipta Chattopadhyay and Abhik Roychoudhury. *Detecting Energy Bugs and Hotspots in Mobile Apps*. Foundation in Software Engineering (**FSE**), 2014
- Sudipta Chattopadhyay, Petru Eles and Zebo Peng. *Automated Software Testing of Memory Performance in Embedded GPUs*. International Conference on Embedded Software (**EMSOFT**), 2014