


# Towards a Formal Foundation of Intermittent Computing

**Milijana Surbatovich**

Brandon Lucia, Limin Jia


# Batteryless Energy-harvesting Devices (EHDs) enable computing in inaccessible environments




Maintenance expensive or impossible



```
x := in()
y := x
z := y + 5
```



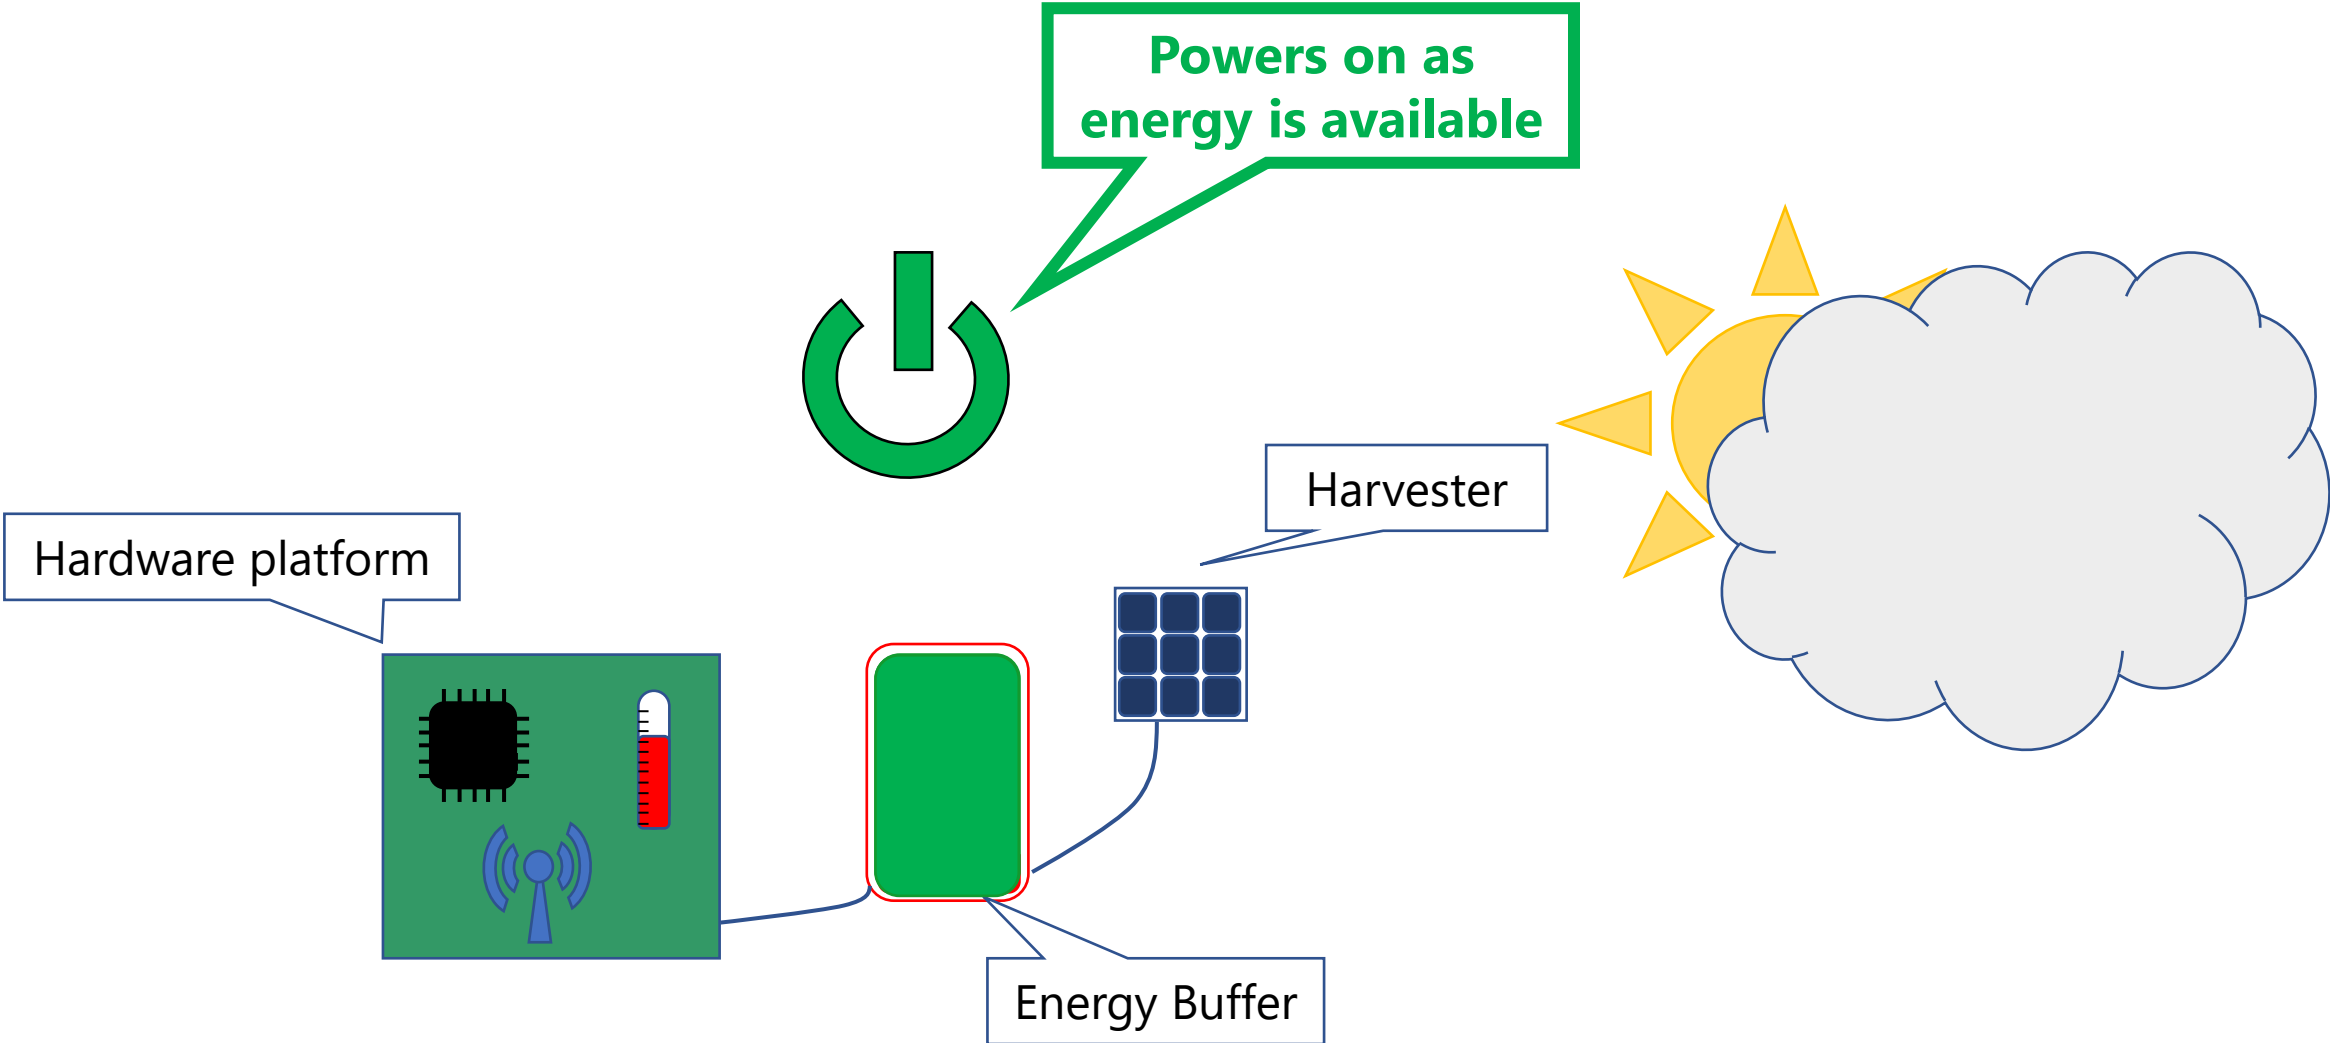
Batteryless EHDs




```
x := in()
y := x
z := y + 5
```



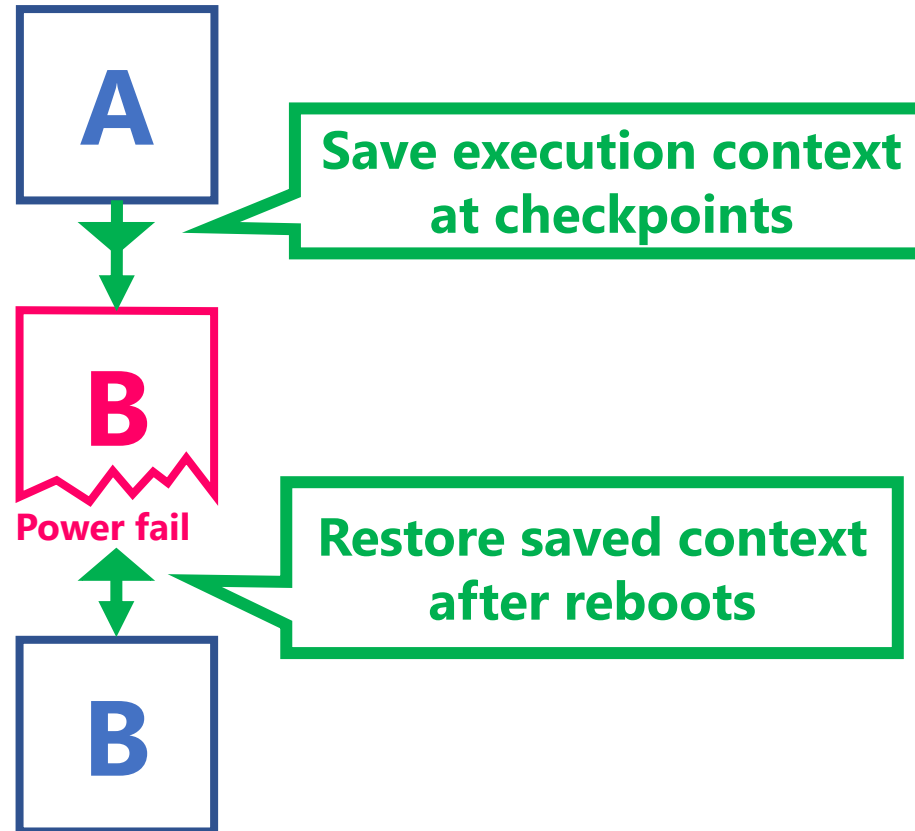
# Intermittent execution in energy harvesting devices



# Intermittent execution in energy harvesting devices

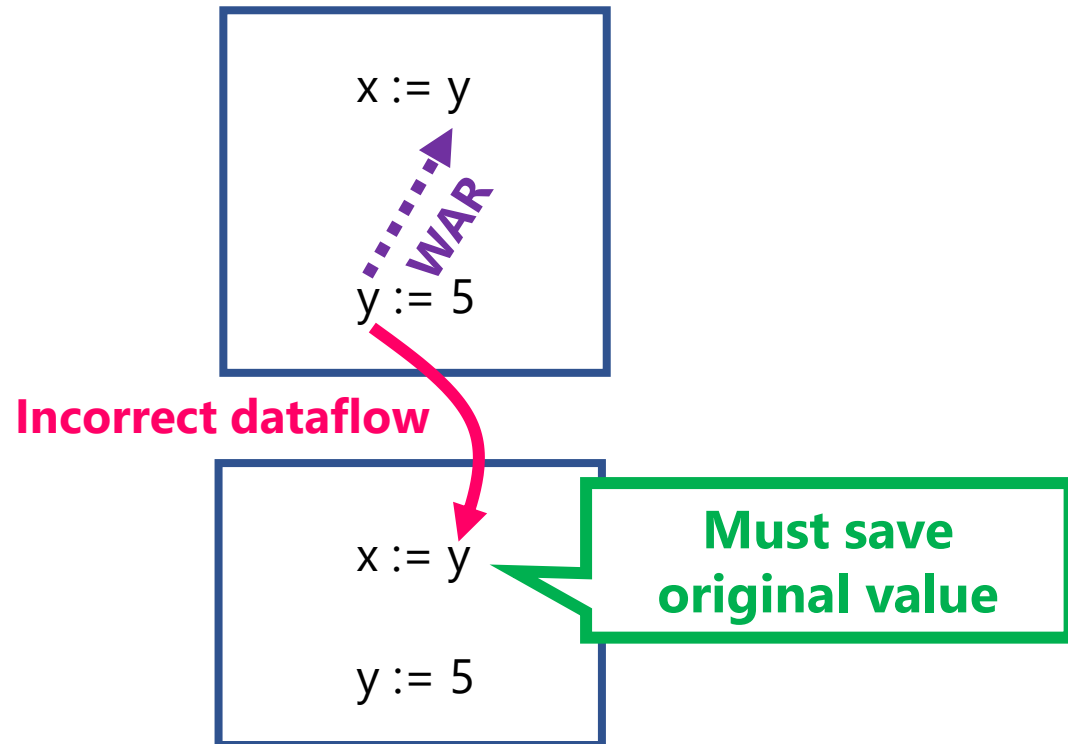


# Preserving progress by saving state



# Systems must re-execute regions correctly

## Write-After-Read (WAR)



Alpaca



Adds value of non-volatile variables with a WAR dependence to the saved execution context

*K. Maeng, A. Colin, B. Lucia. Alpaca: Intermittent Execution without Checkpoints. OOPSLA '17*

Others: DINO, Ratchet,  
Chinchilla

# Input re-executions are not handled correctly

Repeated-Input-Operation (RIO)

```
x := input()
If x > 5:
  y := 1
Else z := 1
```

**Different on re-execution**

```
x := input()
If x > 5:
  y := 1
Else z := 1
```

**Incorrect behaviour!**

IBIS




Detects and reports input-dependent branches that write to different sets of variables

*M. Surbatovich, L. Jia, B. Lucia. I/O Dependent Idempotence Bugs in Intermittent Systems. OOPSLA '19*

# The need to formalize intermittent execution

**No formal spec in existing works → systems subtly incorrect**

Our correctness definitions address both WAR and RIO problems, which no existing work has done



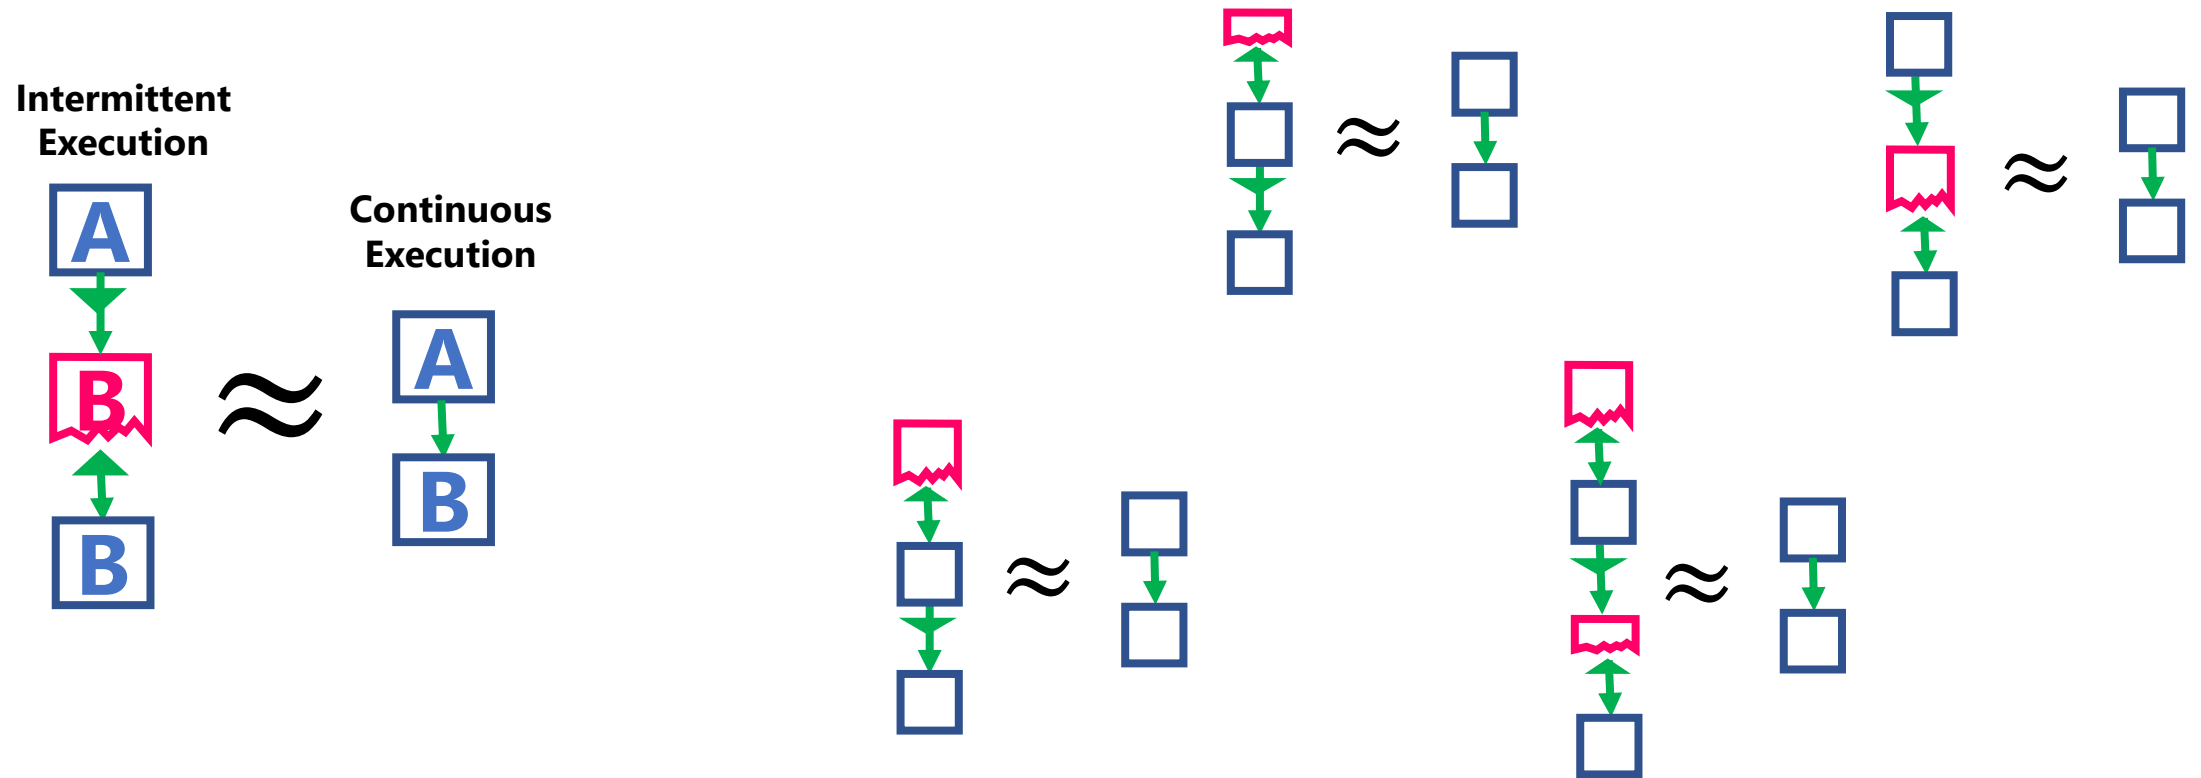


# Outline

- Challenge of intermittence
- **Memory consistency correctness definition**
- Memory relations
- Correct checkpoint set
- Evaluation and conclusion

# Correct intermittent execution

Continuous execution specifies correct program behaviour



# Difficulty of reasoning about equivalence


Equivalence: memory reads and memory state at checkpoints

**Intermittent Execution**




$\approx$

**Continuous Execution**



**Reboots don't restore to the exact same state**



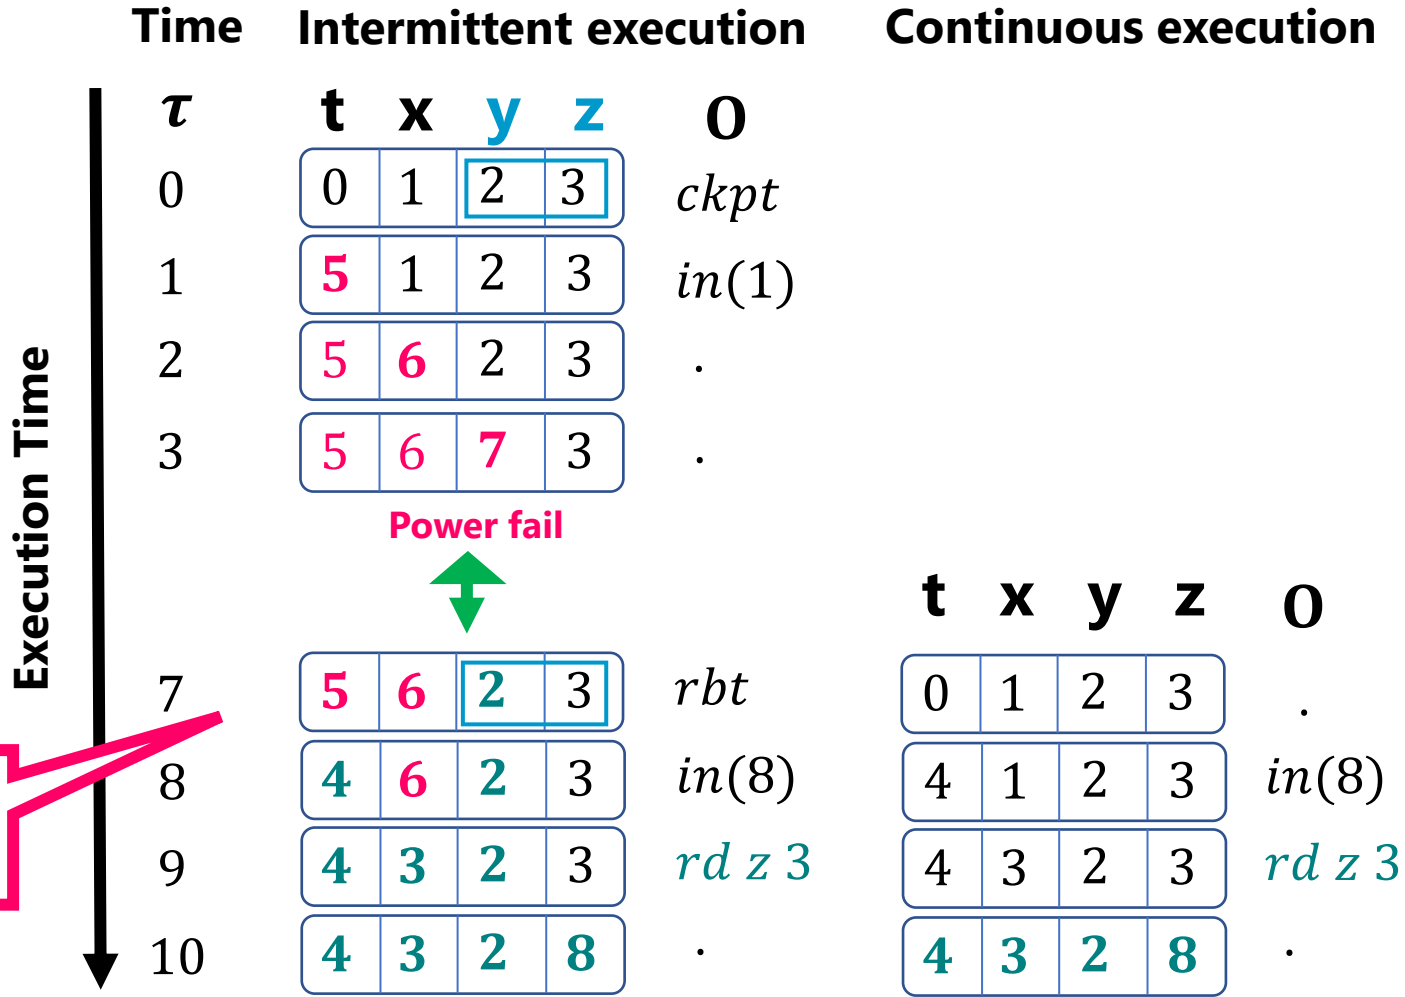
**Inputs cause different paths to be taken**



# Memory can be different at many points

```

0 checkpoint(y,z)
1 t = temp();
2 if t >= 5
3 then x := 6;
4     y := 7;
5 else x := z;
6     z := 8;
    
```



Not same state as at checkpoint

How different can memory get that the differences still resolve?

# Outline

- Challenge of intermittence
- Memory consistency correctness definition
- **Memory relations**
- Correct checkpoint set
- Evaluation and conclusion

# Any differences must resolve on re-execution

Intermittent

t	x	y	z
0	1	2	3
5	1	2	3
5	6	2	3
5	6	7	3

Power fail



5	6	2	3
4	6	2	3
4	3	2	3
4	3	2	8

Continuous

t	x	y	z
0	1	2	3
4	1	2	3
4	3	2	3
4	3	2	8

Differing locations must be written to before being read

Written values must be the same

# Any differences must resolve on re-execution

Intermittent

t	x	y	z
0	1	2	3
5	1	2	3
5	6	2	3
5	6	7	3

Power fail



5	6	2	3
4	6	2	3
4	3	2	3
4	3	2	8

Differing locations must be written to before being read

Continuous

t	x	y	z
0	1	2	3
4	1	2	3
4	3	2	3
4	3	2	8

What set of variables is safe?

All Variables

Safe	Checkpointed
t x	y z

Written values must be the same

# Must-first-write set

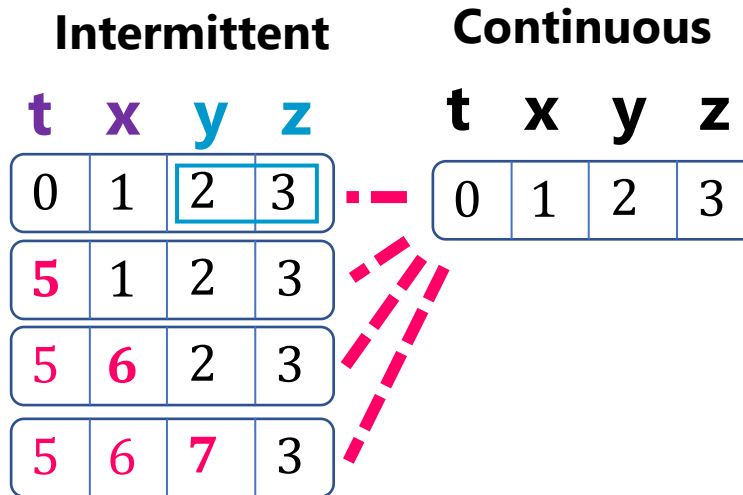
The *must-first-write* set – must-write variables with no preceding read

Any execution writes to these variables before reading them


```
0  checkpoint(y,z)
1  t = temp();
2  if t >= 5
3  then x := 6;
4       y := 7;
5  else x := z;
6       z := 8;
```



# Defining allowable differences



Power fail



Final state matches

Always writes to MstFstWt + checkpoint

Relation 1: Differing locations must be in the MstFstWt set or the checkpoint set

Relation 2: Differing locations must be in the MstFstWt set and cannot have been written to yet.

Each write resolves differences

# Outline

- Challenge of intermittence
- Memory consistency correctness definition
- Memory relations
- **Correct checkpoint set**
- Evaluation and conclusion

# Only checkpointing WAR variables is incorrect

Must-first-write

Checkpoint Set

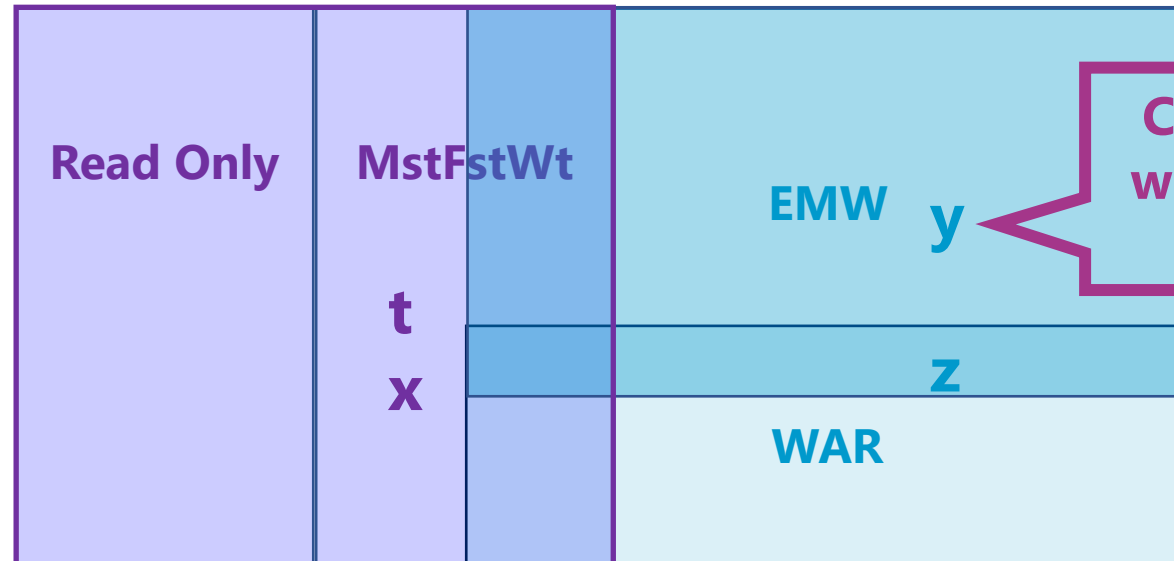
<b>t</b>	<b>x</b>	<b>y</b>	<b>z</b>
0	1	2	3
⋮			
4	3	2	8

```

0  checkpoint(y,z)
1  t = temp();
2  if t >= 5
3  then x := 6;
4     y := 7;
5  else x := z;
6     z := 8;
    
```

Safe

Must be Checkpointed



Conditionally written due to inputs

**Exclusive May-Write** set:  
may-writes minus must-write

# Collecting Exclusive May-Writes


Only inputs can cause a different path to execute after reboot

Use static taint analysis to identify input-dependent branches

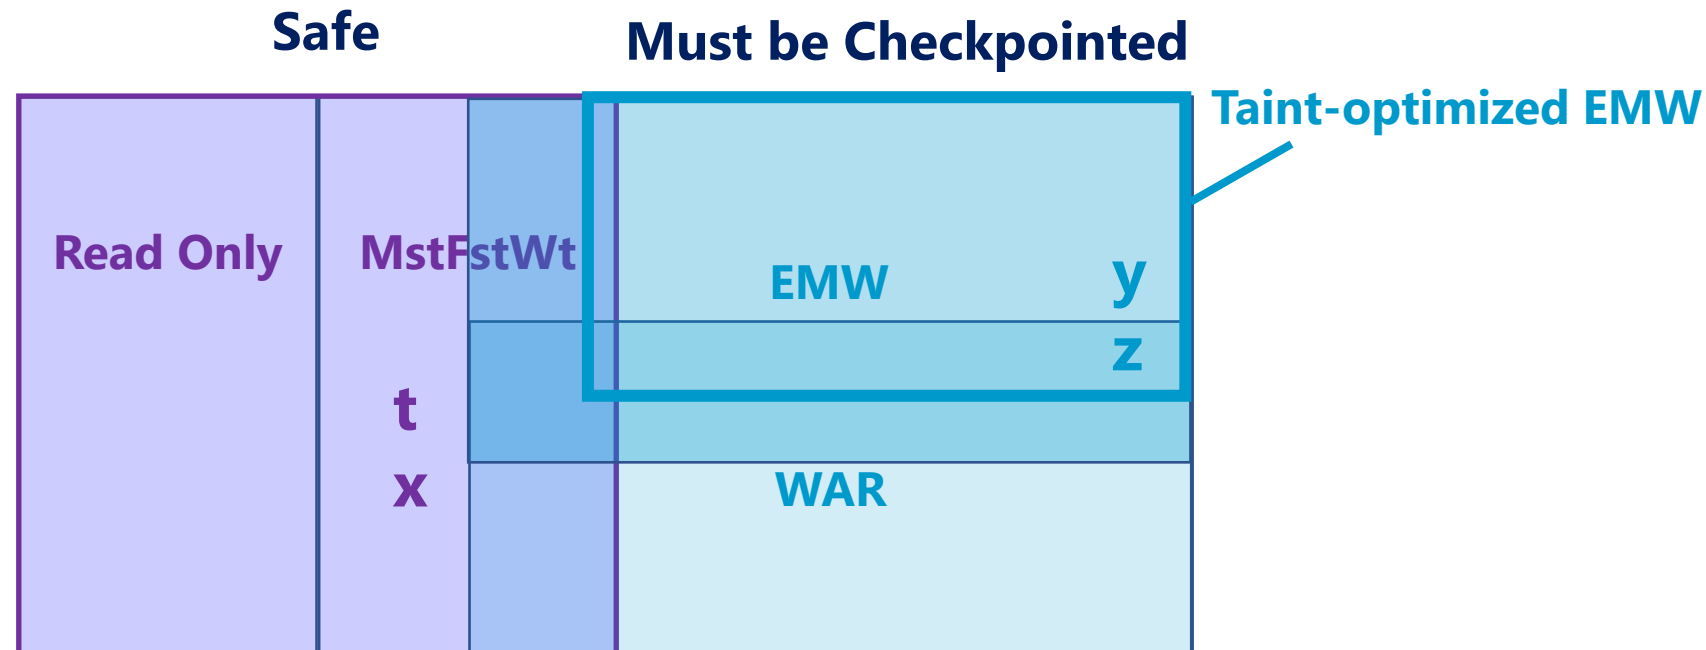
```

0  checkpoint(y,z)
1  t = temp();
2  if t >= 5
3  then x := 6;
4       y := 7;
5  else x := z;
6       z := 8;

```

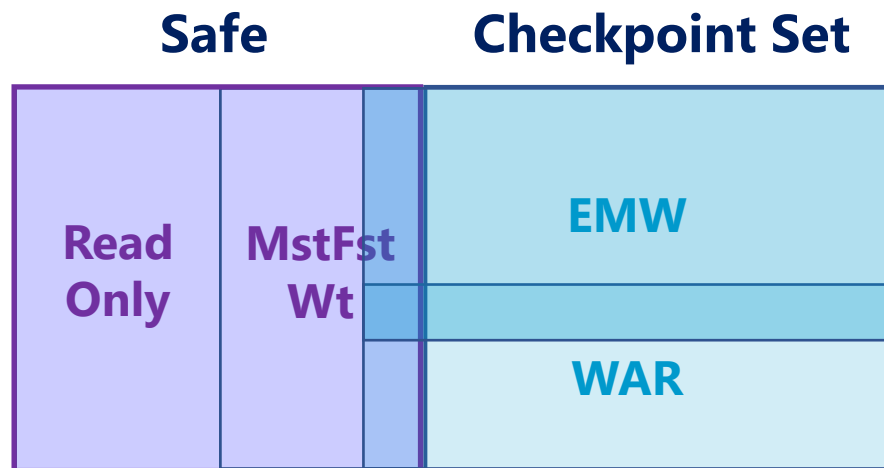


Mst-Wt EMW



# Correctness Theorem

If all unsafe WAR and EMW variables are in the checkpointed set, then an intermittent program will execute correctly



Intermittent Execution



≈

Continuous Execution



# Implementation

Compiler pass implemented in LLVM

Two versions: taint-optimized EMW and basic EMW

Analysis added to Alpaca, which tracks WAR

More in paper...


# Outline

- Challenge of intermittence
- Memory consistency correctness definition
- Memory relations
- Collecting the correct checkpoint set
- **Evaluation and conclusion**

# Goal of evaluation

Show that Modifying Alpaca with EMW is practically efficient

1) Low runtime overhead



2) Low programmer burden

```
x := in()
y := x
z := y + 5
```



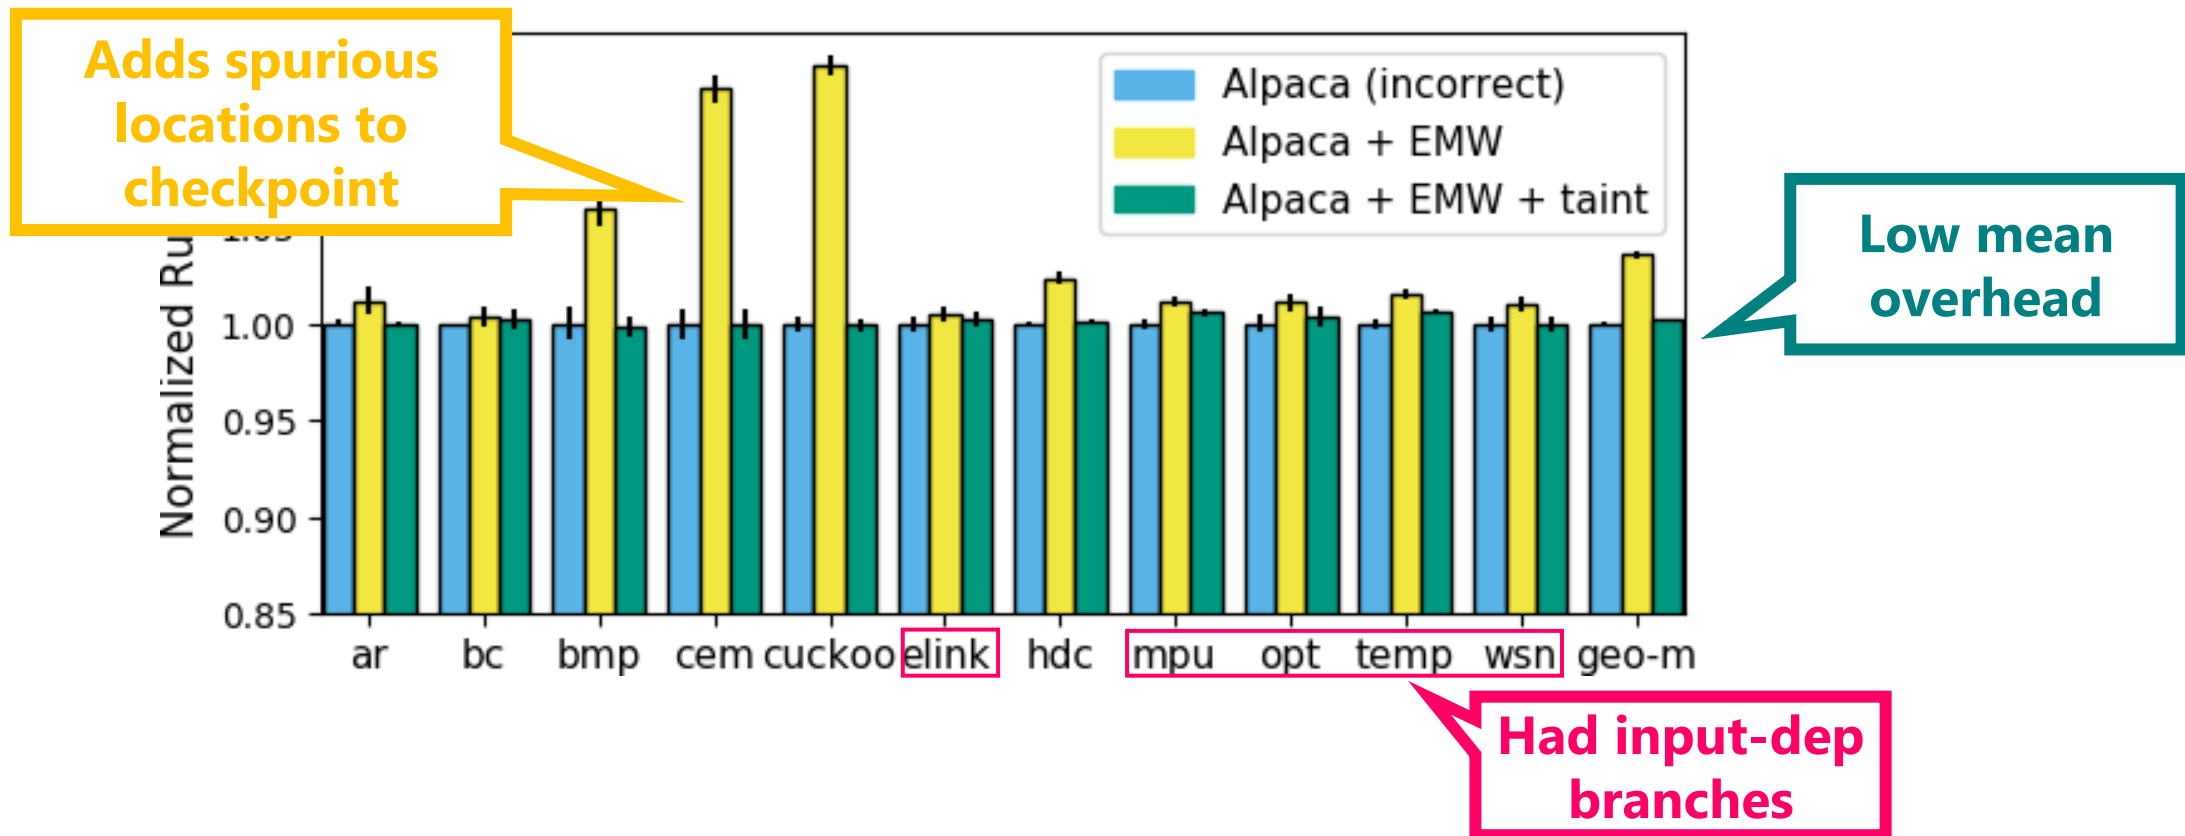
```
x := in()
y := x
z := y + 5
```





# EMW has little performance penalty

Experiments run on benchmarks from prior work on real hardware



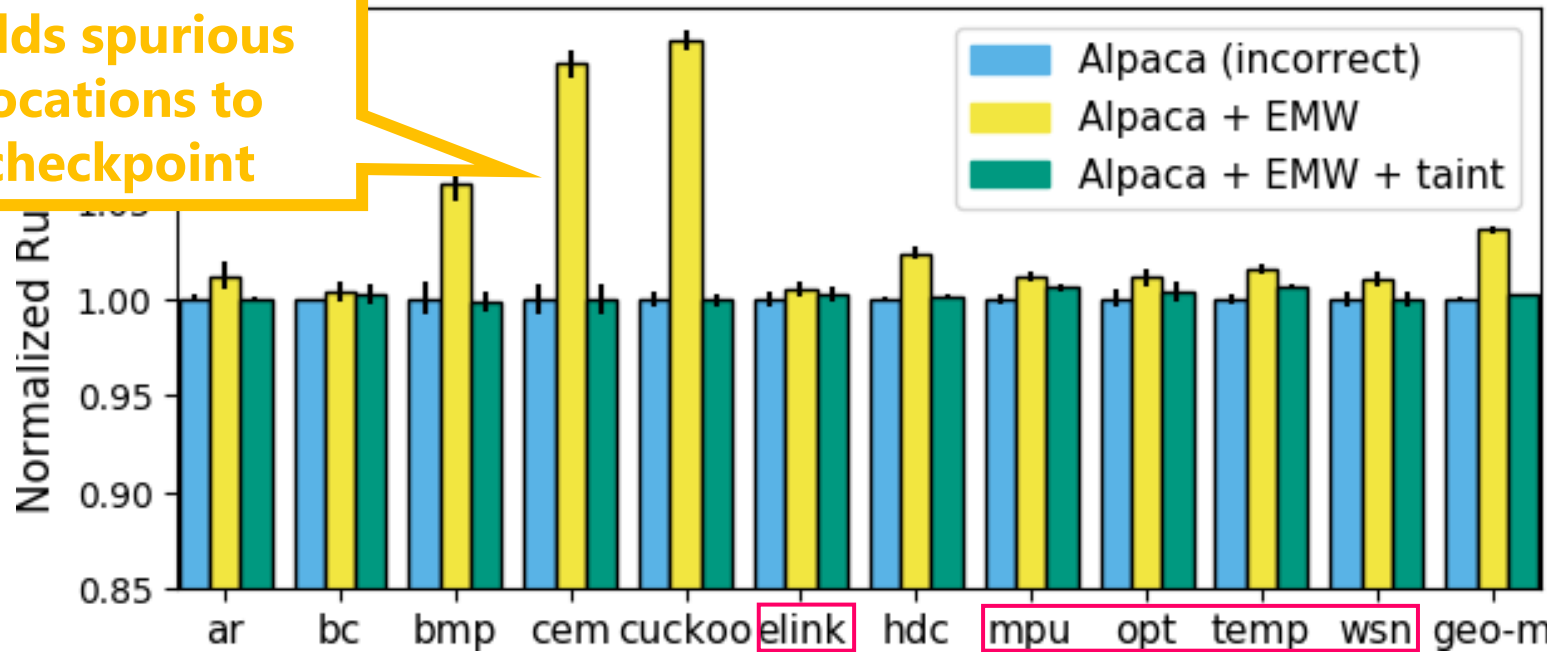
# EMW needs little to no programmer effort

Manual Fix, could still be incorrect

No effort, higher overhead

Specify input functions, little overhead

Adds spurious locations to checkpoint



Low mean overhead

Had input-dep branches

M. Surbatovich, L. Jia, B. Lucia. I/O Dependent Idempotence Bugs in Intermittent Systems. OOPSLA '19

# More in paper

Proving equivalence between execution models

Collection and checking algorithms

Implementation and experiment details

Application Discussion

# Connection to related work

## Persistent Memory Models

Persist vs execution order

Multi-threaded executions

ISA persistency semantics

[Raad et al., Israelevitz et al., Pelley et al.]

## Crash Consistency

Equivalence of crashy execution to non-crashy

Automated proof tools:

Yggdrasil, CHL

Fault Tolerant Resource Reasoning

Crash Consistency through Reachability

[Bornholt et al., Chen et al., Ntzik et al., Koskinen and Yang]

## This work

Explicitly considers **non-deterministic inputs**

Defines **correctness conditions for intermittent executions**

# Summary

Intermittent computing systems need to be correct and reliable

We develop a framework and give a formal definition of correctness

We apply the framework to reason about equivalence and develop a compiler analysis to make existing systems correct

# Towards a Formal Foundation of Intermittent Computing

**Milijana Surbatovich**

Brandon Lucia, Limin Jia