


# Generating Robust Parallel Programs via Model Driven Prediction of Compiler Optimizations for Non-determinism

**Girish Mururu, Kaushik Ravichandran, Ada Gavrilovska, Santosh Pande**

"Insanity is doing the same thing over and over again and expecting different results"

*Albert Einstein*



Computers can do the same thing over and over again  
emitting different results

Are computers **sane** or **Insane**?



# Are Computers Sane or Insane ?

Regardless of using them to do sane or insane things, computers are **Non-deterministic**

# Non-determinism

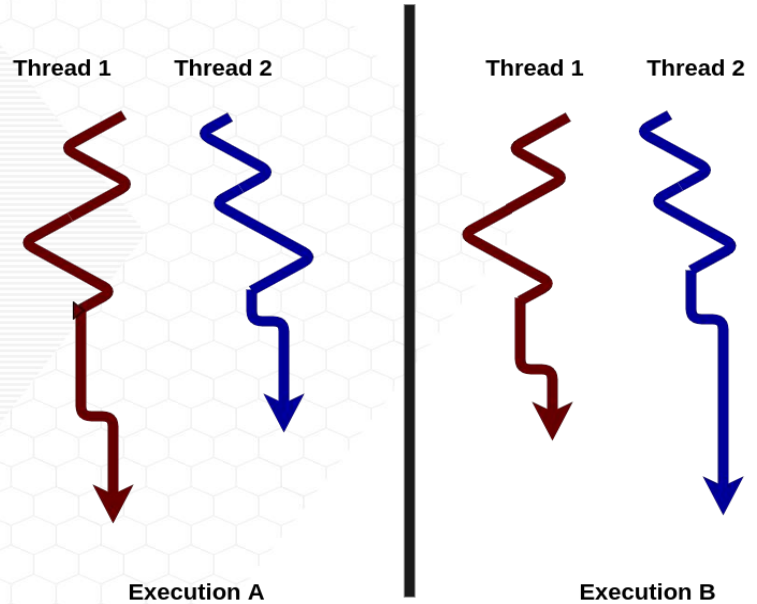
- Variant behavior exhibited during repeated execution with same input
- Different sources
  - ◆ Architecture, OS, runtime

## Non-determinism in Sequential Programs

- Co-executing programs
- Context switches
- Architectural causes
  - ◆ Branch mispredictions , cache and TLB, misses etc.

# Non-determinism in Parallel Programs

- In addition to sequential causes
- interference
- resource sharing
- scheduling decisions



## Parallel Programs: Implications of Non-determinism

- Can suffer from extreme to debug concurrency behaviors
- Uncommon thread interleavings (Corner cases)
  - ◆ Can carry bug into production
  - ◆ Difficult for developers to reproduce
- Occurrence of the bug is related to Non-determinism
  - ◆ Thread interleavings and more (such as a pointer state)



## Non-determinism and Bugs

- Non-determinism in parallel programs causes
  - ◆ Variability in thread execution interleavings
- Influencing Non-determinism affects rare interleavings
  - ◆ More Non-determinism, more interleavings possible
  - ◆ Less non-determinism, less interleavings possible
- Can Non-determinism be a control-knob for **Robustness?**

# Non-determinism and Robust Programs

- Decreasing Non-determinism
  - ◆ Makes exhibition of a rare thread interleaving *less likely* in production
  - ◆ Testing covers most *likely* behaviors (interleavings)
- Increasing Non-determinism
  - ◆ Help developers reproduce rare bugs exercising corner cases of interleavings
- Completely Removing non-determinism slows down programs (prior works)

## Compiler flags - debug vs optimize

- Developers select compiler optimization levels
- For speed or code-size
  - ◆ Release - e.g. O3
  - ◆ Debugging convention - e.g. O0
- Can compiler optimization effect Non-determinism?
  - ◆ No known non-determinism models

## Goals

- We quantify and characterize architectural non-determinism
  - ◆ Develop non-determinism model
    - Use Hardware Performance Counters
  - ◆ Minimize OS effects
- Use model to predict the compiler optimization level that corresponds to the least and most non-determinism
  - ◆ We empirically show this correlation through regression testing
- We use the predicted flag to generate the corresponding debuggable or robust parallel program

## Definitions

- Commit Order - order in which threads complete
  - ◆ E.g. 2,1,3,4 or 4,3,2,1 , different permutations
  - ◆ Commit orders possible =  $n!$
- Non-determinism - Unique commit orders
  - ◆ All  $n!$  orders need not be exhibited
  - ◆ A new permutation exposes new synchronization behavior

## Definitions

- Most Common Permutation (MCP)
  - ◆ The most frequent commit order
- Commit Distance of a thread ( $d_t$ )
  - ◆ Distance from the usual position in MCP
    - If MCP  $\rightarrow 2,3,1,4$  ; then  $1,2,3,4 \rightarrow d = 2$  for thread 1
 

Pos  $\rightarrow$  0 1 2 3

1

Pos  $\rightarrow$  0 1 2 3

1

## Definitions

→ Architectural Factors:

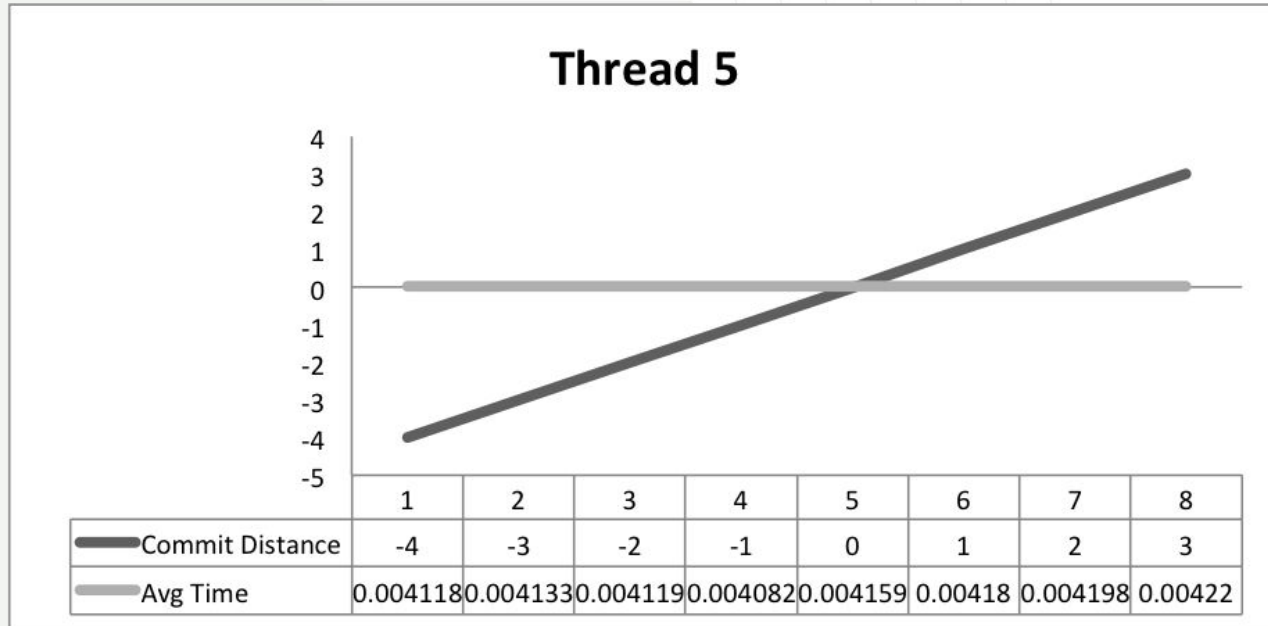
- ◆ Architectural artifacts such as cache misses, branch mispredicts etc. which influence thread execution time

→ We model commit distance in terms of change in artifacts

- ◆ 
$$d_t = f(\Delta h_{1avg}, \Delta h_{2avg}, \dots, \Delta h_{zavg})$$

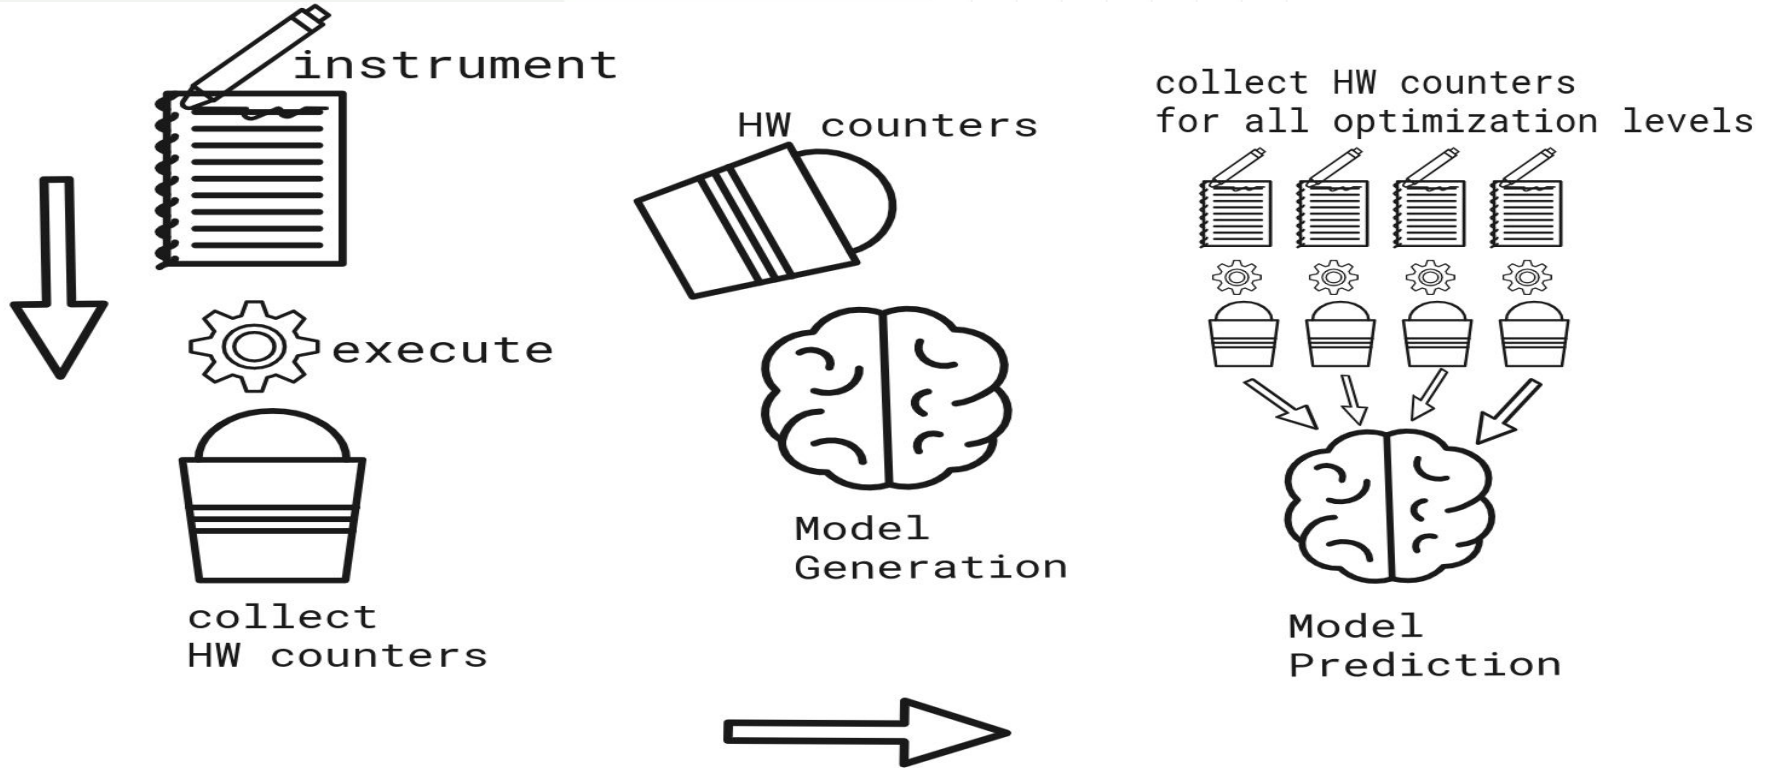
- ◆  $h_1, h_2, \dots, h_z$  are different architectural artifacts

# Model Generation





# Overall Framework



# Model Generation

- Profile hardware counters per thread
  - ◆ Average the counters for a thread in a commit order

$$h_{javg} = \frac{\sum_{k=0}^m h_{jk}}{m}, \forall 1 \leq j \leq z$$

- Deviation from the hardware counter of the thread in MCP

$$\Delta h_{javg} = h_{javg} - hg_{javg}, \forall 1 \leq j \leq z.$$

## Model Generation

- We learn non-deterministic behavior of each thread
  - ◆ Correlate commit distance to deviation in architectural artifacts

$$d_t = f\left(\sum_{j=1}^z \Delta h_{j_{avg}}\right)$$

- We learn a linear regression because the relation between commit distance and execution time is linear

## Model Generation Details

→ We use R to find the coefficients in the linear model:

$$d_t = f\left(\sum_{j=1}^z \Delta h_{j_{avg}}\right)$$

→ Small and Medium inputs for training

→ We normalize the hardware counter values of all commit orders

over MCP

$$\Delta h_{j_{avg}} = \frac{h_{j_{avg}} - hg_{j_{avg}}}{hg_{j_{avg}}}, \forall 1 \leq j \leq z$$

## Model Generation Details

- We generate a model for every step size ( 500 runs) using O0
- Check for convergence
  - ◆ Adjusted rsquared (**ar**): statistical measure of correctness
  - ◆ New thread model is adopted if **ar** is greater than previous
- Stopping condition: once model stabilizes for at least 4 iterations
- We use drop-term from MASS library to remove factors

## Model-based Prediction

- Run O0, O1, O2, O3 at step size (500 runs)
- Plug average hardware counter values in the model
- All thread absolute values are added
- Must exercise the model adequate number of times
  - ◆ Too little will not capture correlation
  - ◆ Very large number of runs take too much time (infeasible)
- Stopping condition:
  - ◆ Based on Bollinger Bands

## Model-based Prediction

- Bollinger Bands are used to measure stability (convergence)
- Each iteration generates **six** Bollinger band values
  - ◆ Last value(LV), middle(M), upper(U), lower(LB), bandwidth(BW), percent(P)
- Stopping condition:
  - ◆ Majority of values (4/6) are min or max
  - ◆ Over four successive iterations
  - ◆ Trend stabilization based on first four (LV, M, U, LB) or derivatives (BW, P)
- Model predicts the Least or Most non-deterministic opt level

## Experiments

- Eight core Intel Xeon with Linux
- Threads pinned!
- GCC- 4.7
- PARSEC pthread benchmarks
- MCP observed in all benchmarks

Specifications	Values
CPU	Intel Xeon E5530
Core count	8
CPU Base frequency	2.40 GHz
Sockets	2
Core/socket	4
HyperThreading	disabled
LLC Cache	8 MB
Page Size	2MB
Main Memory	48 GB
OS	Linux



# Model Accuracy & Non-determinism Reduced

→ Verified by running each version 80,000 times

Benchmark	Most	Prediction Error
Blackscholes	O3	7.9%
Swaptions	O3	3.8%
Canneal	O2	0%
Bodytrack	O1	0%
Fluidanimate	O1	38%
Streamcluster	O2	13%
Raytrace	O0	0%

Least	Non-Determinism Reduced(%)
O2	25.94
O0	12.58
O0	32.35
O0	80.23
O2 (0.0%)	30.5
O3	15.34
O2	7.41

→ Least Non-determinism prediction is ~100% accurate

→ Non-determinism Reduced up to 80%

# Model Characteristics

Benchmark	Model #runs	Gen Time (mins)	Prediction #runs	Pred Time(mins)
Blackscholes	9500	58	5500	183.6
Swaptions	2500	17.16	1500	65.35
Canneal	3500	17.501	18500	490.12
Bodytrack	7000	209.53	1500	174.25
Fluidanimate	1500	6.68	3500	61.21
Streamcluster	5000	3.30	3500	6.62
Raytrace	2000	6.19	1500	23.98

# Bodytrack Model


Thread	Linear Model
Thread0	$0.047170\Delta h_{14} + 0.000902\Delta h_{5} + 0.003052\Delta h_{6} + 0.005743\Delta h_{1} + -0.000608\Delta h_{4} + 0.000129\Delta h_{7}$ $+ -0.036770\Delta h_{13} + -0.007761\Delta h_{12} + 0.000193\Delta h_{15}$
Thread1	$-0.004823\Delta h_{14} + -0.000140\Delta h_{5} + 0.004911\Delta h_{6} + -0.002535\Delta h_{1} + 0.000493\Delta h_{4} + 0.000586\Delta h_{7}$ $+ 0.005430\Delta h_{13} + 0.215039\Delta h_{12} + 0.001700\Delta h_{15}$
Thread2	$-0.052642\Delta h_{14} + 0.000185\Delta h_{5} + 0.004871\Delta h_{6} + 0.006364\Delta h_{1} + -0.001203\Delta h_{4} + 0.000815\Delta h_{7}$ $+ 0.095173\Delta h_{13} + 0.286653\Delta h_{12} + 0.001542\Delta h_{15}$
Thread3	$-0.001099\Delta h_{14} + -0.000692\Delta h_{5} + 0.010845\Delta h_{6} + 0.001082\Delta h_{1} + -0.001843\Delta h_{4} + 0.000357\Delta h_{7}$ $+ 0.075876\Delta h_{13} + 0.073614\Delta h_{12} + 0.001148\Delta h_{15}$
Thread4	$0.216100\Delta h_{14} + -0.000411\Delta h_{5} + 0.010000\Delta h_{6} + -0.002152\Delta h_{1} + -0.001766\Delta h_{4} + 0.000003\Delta h_{7}$ $+ 0.056330\Delta h_{13} + -0.508800\Delta h_{12} + 0.000469\Delta h_{15}$
Thread5	$0.040220\Delta h_{14} + 0.001287\Delta h_{5} + 0.004525\Delta h_{6} + 0.000302\Delta h_{1} + -0.002025\Delta h_{4} + 0.000351\Delta h_{7}$ $+ -0.014940\Delta h_{13} + -0.251100\Delta h_{12} + 0.000449\Delta h_{15}$
Thread6	$-0.034110\Delta h_{14} + 0.000163\Delta h_{5} + 0.001886\Delta h_{6} + 0.000022\Delta h_{1} + 0.000548\Delta h_{4} + 0.000757\Delta h_{7}$ $+ -0.074510\Delta h_{13} + 0.044290\Delta h_{12} + 0.000449\Delta h_{15}$
Thread7	$-0.014660\Delta h_{14} + 0.000183\Delta h_{5} + 0.002347\Delta h_{6} + 0.003954\Delta h_{1} + 0.000489\Delta h_{4} + 0.000794\Delta h_{7}$ $+ -0.054610\Delta h_{13} + 0.008929\Delta h_{12} + 0.000298\Delta h_{15}$

# Correlating Factors


Benchmark	Architectural Artifacts
Blackscholes	B, C, F, L, S, T
Swaptions	B, C, F, L, R, S, T
Canneal	B, C, F, L, R, S, T
Bodytrack	B, C, L, S, T
Fluidanimate	C
Streamcluster	C
Raytrace	B, C, F, L, R, S, T

(B - branches) (C - cache) (F - floating point operations) (L - loads) (S - stores)  
 (R - resource stall cycles) (T - TLB)


# Most vs Least Non-determinism



(d) Bodytrack



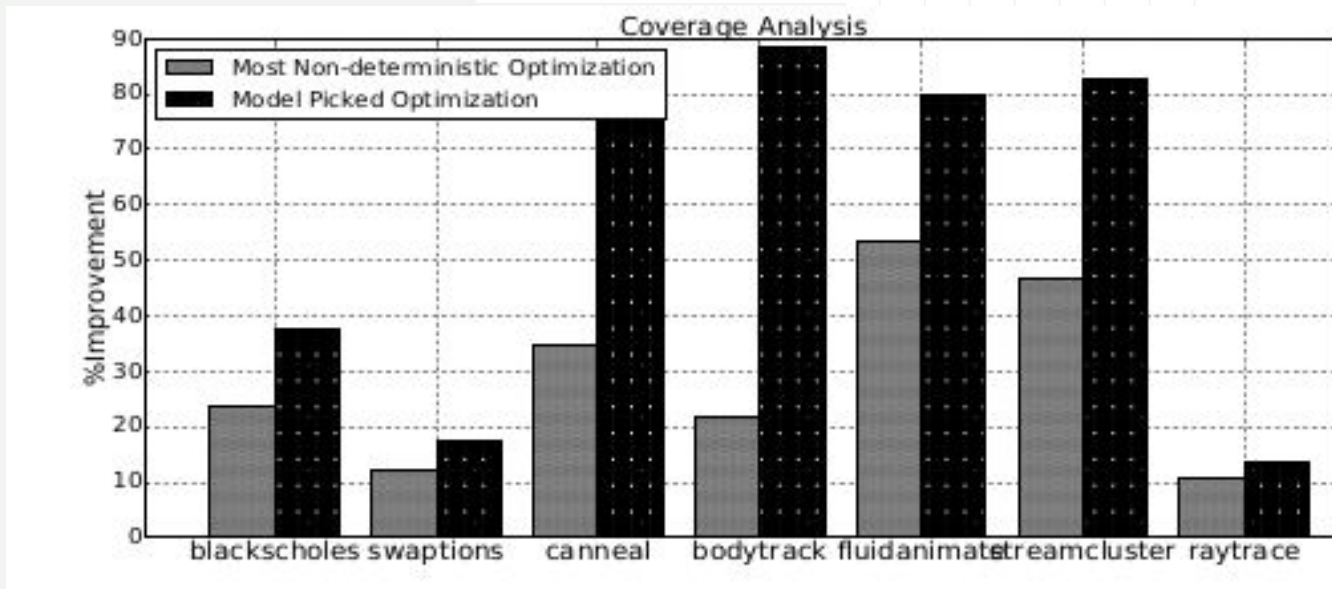
(e) Fluidanimate



(f) Streamcluster

# Robustness Improvement via Better Coverage

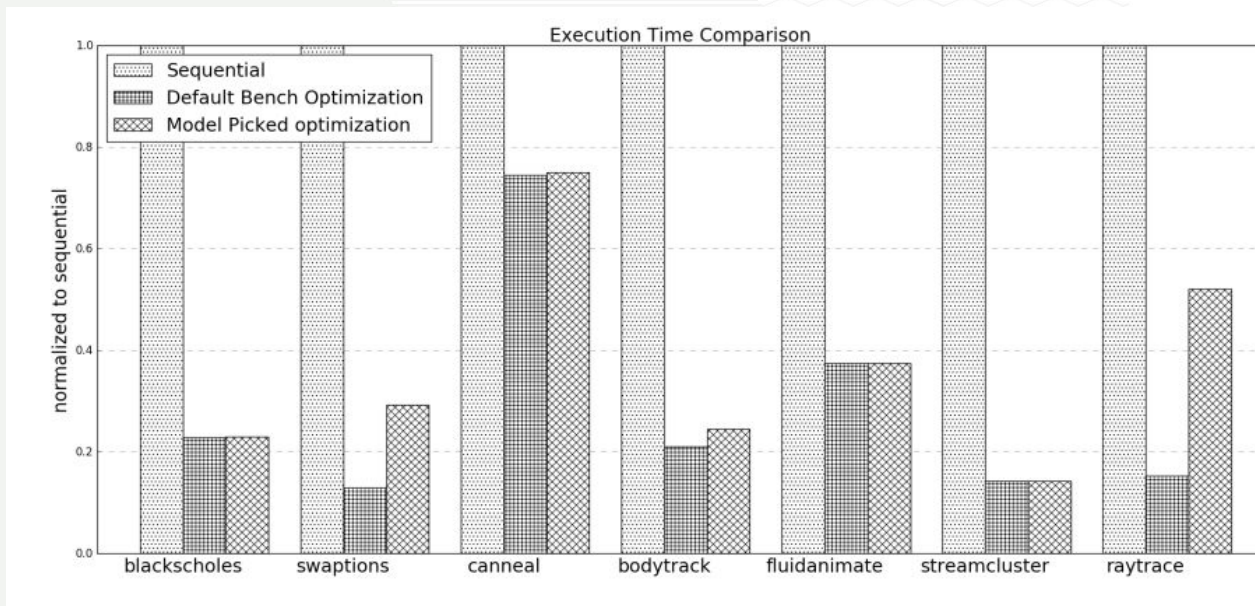
→ Percentage of 80,000 runs covered by top 1000 permutations





# Performance Comparison with Default Optimization

→ Section specific compilation achieves close to default opt



## Race Conditions Made Benign

- Known race conditions
  - ◆ Parsec 2.1 Streamcluster known race condition
    - Caused by a missing barrier
- Random race conditions ( 10 race conditions inserted)
  - ◆ Blackscholes - Debug 6 (11 failures), Release failed once
  - ◆ Swaptions - Debug 7 ( one 7 times) , Release 5 (once)



## Related Works

- [DeSTM: harnessing determinism in STMs for application development](#)
  - ◆ Kaushik Ravichandran, Ada Gavrilovska, Santosh Pande (Pact '14)
  - ◆ Improves debugging through deterministic replay of permutations
- [Quantifying and Reducing Execution Variance in STM via Model Driven Commit Optimization](#)
  - ◆ Girish Mururu, Ada Gavrilovska, Santosh Pande (CGO 19, PPOPP 18)'
  - ◆ Model Non-determinism to reduce timing variance in STM

## Future work

- Different definitions of non-determinism
- Modelling other sources of non-determinism
  - ◆ OS, runtime
- Compiler optimizations geared towards non-determinism
- Debugging methodologies by actively controlling non-determinism

## Conclusion

- Systematically models non-determinism
- Model predicts most and least non-deterministic versions
- Least non-deterministic version reduces non-determinism
  - ◆ Up to 80.23%, Coverage up to 66.48%
  - ◆ Make bugs benign
- Most non-deterministic version easy to reproduce bugs