

EPVF: AN ENHANCED PROGRAM VULNERABILITY FACTOR METHODOLOGY FOR CROSS-LAYER RESILIENCE ANALYSIS

Bo Fang †, Qining Lu †, Karthik Pattabiraman †, Matei Ripeanu †

and Sudhanva Gurumurthi *

† The University of British Columbia, Canada

*Cloud Innovation Lab, IBM, USA







What are we facing?



C is increasing





Why Software-based Fault Tolerance



Software-based techniques: more cost-effective

a place of mind

Mitigating Silent Data Corruption (SDC): Key to Error Resilience





Error Resilience Estimation: Accuracy vs Cost





Identifying SDC-causing Bits

 AVF/PVF: Identify Architecturally Correct Execution (ACE) Bits [MICRO03, HPCA10]



e(nhanced)PVF: a methodology that distinguishes crash-causing bits from ACE bits

C a place of mind





Our Approach: ePVF R2 ADDR1 LD **R8** LD R1 R3 ADD ADD R4 R6 ST ADD ADDR2 R7 R5 ADD

- Source of crashes
 - Segmentation faults (99% of crashes are due to segfaults)
- Direct crash-causing bits
 - Crash model
- Indirect crash-causing bits Source of crashes
 - Propagation model









- Determining the bits that cause an out-of-bound memory access
- Applied on every memory instruction

Crash model





Identifying all possible bits that can affect the bits identified by the crash model R1 = LD R2 R4 = ADD R1, R3 R5 = ADD R6*4 + R7 ST R4, R5 R8 = LD R2

 max(R6) = (max(R5) - R7)/4 min(R6) = (min(R5) - R7)/4
 max(R7) = max(R5) - R6*4 min(R7) = min(R5) - R6*4



Overall ePVF methodology





Experimental setup

- Scientific benchmarks
 - 8 from Rodinia [IISWC 09]
 - Matrix Multiplication
 - LULESH: DOE proxy app [IPDPS 2013]
- Fault Model
- LLFI [DSN 14]
 - 3,000 runs per benchmark



Evaluation

- RQ1: Accuracy of the models
- RQ2: Effectiveness of the ePVF methodology
- RQ3: Performance





90% of the Mod Recall 80% Crash trials vnorimont 70% Randomly pick Pick the flipped bit for a crash trail 60% a bit from the 100% models Check that bit for the model **Recall of the Model** 90% 80% 100% Crash trials **Precision of the Model** Pre 70% 90% Flip the exact bit during the **60%** 80% Our models achieve average execution 50% Pick 70% 89% recall and 92% bit for a crash 60% trail precision 50% Check if a 30 crash occurs Chè. for the model

RQ1: Accuracy of the models

15

RQ1. Accuracy of the Models



On average, 90% of the time the ePVF methodology is accurate to identify crash-causing bits



RQ2: Effectiveness of the ePVF

SDC estimate using PVF analysis, ePVF analysis and Fault Injection





ePVF-informed Duplication

- Rank instructions based on their ePVF value
 - ePVF value per instruction = $\frac{ACE \ bits Crash cuasing \ bits}{ACE \ bits}$
 - Higher the ePVF value, Higher chance to lead to SDCs
- Duplication highly-ranked ePVF instructions
- 30% more SDC coverage than hot-path duplication for the same performance overhead



RQ3: Performance

- Modeling time ranges from 30s (lavaMD) to ~ 4 hours (pathfinder).
 - Depending on the size of the DDG, hence the number of dynamic instructions
- Optimization (Sampling and Extrapolation)
 - Intuition scientific applications usually have repetitive behaviors.



Extrapolated ePVF values based on 10% of the graph, and showing less than 1% difference on average

19



Conclusion

- ePVF removes the crash-causing bits from PVF to get a more accurate estimate of SDC rate.
 - A crash model that predicts direct crash-causing bits
 - A propagation model that identifies bit that lead to direct crash-causing bits
 - Implementation with LLVM compiler
 - Drive selective protection of SDC-causing instructions

Email: bof@ece.ubc.ca Code: https://github.com/flyree/enhancedPVF