Dynamic Programming Based FPGA Hosted Control-Flow Integrity (CFI)

Design Document

sdmay21-10
Professor Akhilesh Tyagi
Advisers: Ananda Biswas and Zelong Li
Greg Wendt/Meeting Scribe
Cole Schumacher/Meeting Facilitator
Nickolas Mitchell/Chief Engineer
Samuel Henley/Chief Engineer
Maxwell Wangler/Test Engineer
Tristan Duyvejonck/Report Manager
sdmay21-10@iastate.edu
http://sdmay21-10.sd.ece.iastate.edu

Revised: 1/15/2020 Final

Executive Summary

Development Standards & Practices Used

We use the following standards for interfacing with the FPGA:

- FMC (FPGA Mezzanine Card): Most popular and versatile interface standard
- Digital PMOD: A nice, simple alternative to FMC
- SYZYGY: newer standard that provides a middle ground between these two
- Standard design practices:
    - Well-documented
    - Efficient use of resources
    - Using proper IP's and memory cores

We use the following standards in regards to Control-Flow Integrity:

- Current CFI strategies
    - Stack Canaries
        - Non-executable memory
    - Write XOR eXecute
    - Code randomization
- Coarse-grained CFI
- Fine-grained CFI
- We want the best possible security with the lowest possible performance impact

The following IEEE Standards were considered in this design:

- IEEE 1008-1987: IEEE Standard for Software Unit Testing
- IEEE 1500-2005: IEEE Standard Testability Method for Embedded Core-based Integrated Circuits
- IEEE 15288-2004: Systems and Software Engineering System Life Cycle Processes
- IEEE 1220-2005: IEEE Standard for Application and Management of the Systems Engineering Process

Summary of Requirements

A high level overview of the function requirements is given below:

- We will create an algorithm to match two sets of strings
- One string is generated based on program control flow, and is compared to a defined string
- It will use a two dimensional grid with its sized based on the number of basic blocks in the program

A high level overview of the nonfunctional requirements is given below:

- The algorithm must perform fast. The goal is O(nlog(n)).

Applicable Courses from Iowa State University Curriculum

The following courses at Iowa State University were applicable to our project:

CPRE 281 - Digital Design

COM S 228 - Introduction to Data Structures

COM S 309 - Software Development Practices

COM S 311 - Design and Analysis of Algorithms

COM S 321 - Introduction to Computer Architecture and Machine-Level Programming

COM S 327 - Advanced Programming Techniques

ENGL 314 - Technical Communication

S E 329 - Software Project Management

New Skills/Knowledge acquired that was not taught in courses

List all new skills/knowledge that your team acquired which was not part of your Iowa State curriculum in order to complete this project.

LLVM   Parallelism with FPGAs

PMU   CFI

Table of Contents

List of figures/tables/symbols/definitions

1 Introduction

1.1 ACKNOWLEDGEMENT

Work on this project is taking place along with Dr. Akhilesh Tyagi, and graduate assistants Ananda Biswas and Zelong Li. All of which have a significant expertise in the computer engineering and software engineering field. Our team is working closely with these assistants to familiarize ourselves with the project and the problems we need to solve.

1.2 PROBLEM AND PROJECT STATEMENT

General Problem Statement:

      We are trying to determine if a break exists within the control flow of a program. When a program is compiled, it is split into sets of instructions called basic blocks. The computer jumps between basic blocks during execution. A vulnerability exists where a malicious user could inject some code which causes the program to jump to an unexpected block. In this case, the Control Flow Integrity has been breached.

General Solution Approach:

      To do this, we will develop an algorithm for comparing the expected control flow of a program to the given control flow. The algorithm will determine if the difference between control flows constitutes a breach in Control Flow Integrity. The algorithm will implement concepts of race logic and will run on an FPGA board.

1.3 OPERATIONAL ENVIRONMENT

The Operational Environment for this program will be lab rooms. No extreme conditions are expected.

1.4 REQUIREMENTS

A high level overview of the function requirements is given below:

- We will create an algorithm to compare two sets of strings
- One string is generated based on program control flow, and is compared to a defined string

- It will use a two dimensional grid with its sized based on the number of basic blocks in the program

- Standardized solution to DE2-115 FPGA board.

A high level overview of the nonfunctional requirements is given below:

- The algorithm must perform fast. The goal is $O(n\log(n))$.

1.5 Intended Users and Uses

As the program is experimental, we don't have a set user base in mind at the moment beyond the Professor and Graduate Students. As far as intended uses are concerned, we intend our project to be used in applications where security with low overall system resource impact is important.

One of the core use cases is a scenario where a malicious user attacks a computer server by attempting to alter the control flow of a program. Our program will notify the system of a control flow breach in an effort to mitigate the attack. In this scenario, the user would be the personnel managing the computer server.

1.6 Assumptions and Limitations

Assumptions:

➢ The solution for DNA sequencing can be applied to this problem.

Limitations:

➢ We must use dynamic programming to determine if we lost control flow integrity.
➢ We must use a FPGA board to test our solution.
➢ The FPGA board shouldn't cost more than $100.

1.7 Expected End Product and Deliverables

The deliverable for this project is an FPGA board with a working program. The program shall compare two strings: one generated based on program control flow and a defined one. The expected delivery date for this is April 30, 2021.

The expected delivery dates are as follows:
➢ Ideate solutions: October 26, 2020
➢ Test solution in higher-level programming language: November 20, 2020
➢ Implement race conditions logic: February 12, 2021
➢ Implement solution in FPGA: April 30, 2021

2   Project Plan

2.1 Task Decomposition

1. Develop solution (Meant for understanding race logic/cfg)
   a. Write pseudocode for solution
   b. Develop solution in higher-level programming language
   c. Test solution

      d.   Look at binary/cfg (LLVM)
2. Implement solution in FPGA board
      a.   Implement using FPGA emulator
      b.   Implement on actual FPGA
3. Features (TBD at a later date with the clients)

2.2 Risks And Risk Management/Mitigation

➢ Possibly override code of other team members
    ○   The use of git prevents this to a high degree.

2.3 Project Proposed Milestones, Metrics, and Evaluation Criteria

Our first key milestone includes, completing development on the software solution which will detect whether there is a significant breach in Control Flow Integrity with above 50% accuracy. Note, the target accuracy is subject to change based on testing. Our team's last milestone will be completing implementation on hardware with accuracy similar to our software solution.

2.4 Project Timeline/Schedule

Tasks	9/4/2020	9/11/2020	9/18/2020	9/25/2020	10/2/2020	10/9/2020	10/16/2020	10/23/2020	10/30/2020	11/6/2020	11/13/2020	11/20/2020	11/27/2020	12/4/2020	12/11/2020	12/18/2020	12/25/2020
1a	■	■	■	■	■												
1b					■	■	■	■	■								
1c										■	■						
1d					■	■	■	■	■								
2a												■	■	■	■	■	
2b																	
3																	

Table 2.1: Gantt Chart for tasks (2020)

| Tasks	1/1/2021	1/8/2021	1/15/2021	1/22/2021	1/29/2021	2/5/2021	2/12/2021	2/19/2021	2/26/2021	3/5/2021	3/12/2021	3/19/2021	3/26/2021	4/2/2021	4/9/2021	4/16/2021	4/23/2021	4/30/2021	5/7/2021
1a																			
1b																			
1c																			
1d																			
2a	■	■	■	■	■														
2b						■	■	■	■	■	■	■							
3													■	■	■	■	■		

Table 2.2: Gantt Chart for tasks (2021)

Tables 2.1 and 2.2 show the proposed schedule for completing tasks laid out in Section 2.1.

2.5 PROJECT TRACKING PROCEDURES

- ➢ Our group is using Slack to communicate with the team
- ➢ Our group is using a Google Drive to store team documents.
- ➢ Our group is using Github to store files relating to the project.
- ➢ Our group is a combination of GMAIL, Webex, and ZOOM, to communicate with our clients.
- ➢ Our group is using Trello to keep track of assigned tasks and due dates.

2.6 PERSONNEL EFFORT REQUIREMENTS

Task	Hours Per Person	Team members assigned	Total person-hours
1a	4	1	4
1b	5	2	10
1c	5	3	15
1d	15	3	45
2a	10	3	30
2b	15	2	30
3	TBD	TBD	TBD

Table 2.3: Effort Estimate by task

Our team will expect all team members to complete the tasks assigned to them by a predetermined deadline determined by the team. Our team expects total weekly hours per member to not exceed

around 40 hours in week given several tasks to be completed in said week. In general, our team expects all members to give an appropriate number of hours to provide quality while still time managed work.

2.7 OTHER RESOURCE REQUIREMENTS

Our team plans to use a DE2-115 FPGA board due to its accessibility from Iowa State's ETG department. The DE2-115 FPGA board also offers features important to our project such as JTAG and Active Serial programming use over USB, various LEDs including red and green LEDs to receive real time visual feedback, and more than enough memory to support our solution.

2.8 FINANCIAL REQUIREMENTS

Our team plans on no financial requirements being necessary.

3 Design

3.1 PREVIOUS WORK AND LITERATURE

A. Madhavan, T. Sherwood and D. Strukov, "Race Logic: Abusing Hardware Race Conditions to Perform Useful Computation," in *IEEE Micro*, vol. 35, no. 3, pp. 48-57, May-June 2015, doi: 10.1109/MM.2015.43.

This paper was useful in understanding the algorithm we use in the system. The authors applied race conditions to solve a string-matching problem related to protein sequencing. We want to adapt the algorithm to analyze control flow integrity.

3.2 DESIGN THINKING

To define the problem that shaped our design, we worked closely with the professor to understand the needs and core issues of monitoring Control Flow Integrity. Our final goal is to implement the solution on an FPGA board, but first we wanted to focus on developing a solution that could be implemented and tested in a high-level environment.

When we ideated solutions for the problem, we discussed ideas with our graduate assistants. We focused on one solution and built on it until we could determine if the solution was feasible.

3.3 PROPOSED DESIGN

We are not using any particular design standards at this time. Our first step was to design pseudo code and come up with ideas that theoretically solve the problem. After we discuss the pseudocode with the client, we want to try to implement the pseudo code in a higher-level programming language and test its functionality. Our final proposed design includes implementing our software solution on an FPGA board.

One of our project goals is to implement race logic in the design of our string comparison algorithm. We believe that abusing race conditions and delay operations will result in more efficient execution of our program.

Along with developing the string comparison algorithm, we have begun experimenting with LLVM, a set of tools for analyzing and manipulating compiled programs. LLVM will be useful for injecting functions into another program. LLVM can also insert static labels into a program.

3.4 TECHNOLOGY CONSIDERATIONS

One design alternative available is to use an ASIC system instead of an FPGA. We think the project would be worse off with an ASIC as the processing cycles would be interrupted by the primary compute unit. The ASIC unit would also be more difficult to scale.

Another tradeoff our team is considering alongside our clients is whether to use PMU or static labels to determine if control flow integrity has been breached. Static labels would simply be labels we assign to different blocks of code which would be simpler then PMU but potentially comes at a performance cost. PMU utilizes the libperf library which is a performance analyzing tool which can be used to keep track of the number of times a block of code has been run. PMU will offer a less resource intensive solution when compared to static labels but the reliability of using PMU is uncertain. With more testing of both of these possible solutions, we hope to determine the best for our project.

3.5 DESIGN ANALYSIS

Thus far, our team has only performed minor testing but we plan to expand these efforts in the near future. Moving forward, we hope to test the three main modules described in section 3.7. Testing the readLabels module will include checking whether the labels are correctly assigned in binary. Testing the readExpected module will include checking whether the labels our solution generated match the labels we found in the readLabels module. Lastly, we plan on testing the classify module using LLVM and PMU to determine if control flow integrity is upheld or not.

3.6 DEVELOPMENT PROCESS

Our team is using a waterfall model development process. We decided the nature of our project was not suited for an agile development process with small, incremental iterations. Our project will move through each stage of the development process once and in order.

3.7 DESIGN PLAN

The system is composed of three modules:

1. Reading labels from binary
2. Read expected labels
3. Classifier

The first module reads control labels from the executed binary. The labels could be the PMU data from the processor, or another label to be determined. The second module reads the expected control flow model for the executable. The data from both modules is passed to the classifier, which determines whether there is a breach in control flow integrity.
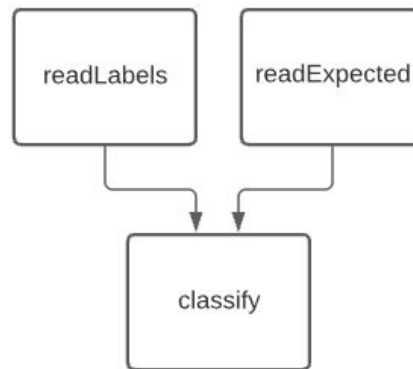

Figure 3.1: Program execution flowchart

This system fulfills the use cases specified for our solution. Our user needs to be notified if a control flow breach is detected. The 'classify' module will determine if there is a control flow breach, which receives data from the 'readLabels' and 'readExpected' modules.

4 Testing

4.1 UNIT TESTING

Software we plan to test in an isolated environment includes our dynamically programmed race logic algorithm, and any CFG programs we design. Hardware we plan to test in an isolated environment includes an FPGA board running our solution.

4.2 Interface Testing

The communication between the reading modules and the classification module will be tested by recording the data output by the reading modules and analyzing the behavior of the classification module.

4.3 Acceptance Testing

To verify the functionality of the system, develop input sets for best-case and worst-case scenarios. Apply the input sets to the system. For example, create a scenario when the model and the input data match exactly. Upon completion of the project, we will demonstrate the system to the client during one of our meetings. We will use LLVM to aid in acceptance testing.

4.4 Results

Thus far, our team has yet to produce any meaningful results during our testing phase. Please note, our testing phase is ongoing.

5 Implementation

Preliminary implementation for our project includes finalizing our software solution before translating it to an FPGA board.

Next semester, we will start by ensuring that our solution is functional on an FPGA emulator. Next we will work on implementing the solution on a physical FPGA board. Once the solution is working and tested, we will discuss additional features with the client.

6 Closing Material

6.1 Conclusion

Thus far, our team has worked with our professor and graduate assistants to understand control flow integrity and the basics of program analysis. Based on our meetings with the advisers, we developed an algorithm in pseudocode that we could discuss and revise. Then, we implemented the algorithm in a higher-level language. We are currently in the process of testing the algorithm.

Our primary goal is to implement the solution in a physical FPGA board. Secondary to this, we would like the algorithm to run with time complexity $O(n\log(n))$. Our best plan of action is to implement the algorithm in a high-level language first, then an FPGA emulator, then a physical

FPGA board. This way, we are able to transition from one medium to the next, testing and detecting problems along the way.

6.2 REFERENCES

A. Madhavan, T. Sherwood and D. Strukov, "Race Logic: Abusing Hardware Race Conditions to Perform Useful Computation," in *IEEE Micro*, vol. 35, no. 3, pp. 48-57, May-June 2015, doi: 10.1109/MM.2015.43.

6.3 APPENDICES

https://www.intel.com/content/dam/altera-www/global/en_US/portal/dsn/42/doc-us-dsnbk-42-14 04062209-de2-115-user-manual.pdf

The above documentation is the manual to the DE2-115 FPGA board our team plans to use in our solution.

Figure 6.1: Conceptual Sketch