An Approach to Reducing the Cost of Fault Injection

Raul Barbosa, Jonny Vinter, Peter Folkesson and Johan Karlsson

Department of Computer Science and Engineering
Chalmers University of Technology
412 96 Göteborg, Sweden
{rbarbosa, vinter, peterf, johan}@ce.chalmers.se

Abstract. This paper describes a fully automated pre-injection analysis technique aimed at reducing the cost of fault injection campaigns. The technique reduces the fault-space by utilizing assembly-level knowledge of the target system in order to place single bit-flips in registers and memory locations only immediately before these are read by the executed instructions. This way, faults (time-location pairs) that are overwritten or have identical impact on program execution are removed. Experimental results obtained by random sampling of the optimized fault-space and the complete (non-optimized) fault-space are compared for two different workloads running on a MPC565 microcontroller. The pre-injection analysis yields an increase of one order of magnitude in the effectiveness of faults, a reduction of the fault-space of two orders of magnitude in the case of CPU-registers and four to five orders of magnitude in the case of memory locations, while preserving a similar estimation of the error detection coverage.

1. Introduction

Computer systems are increasingly being used in safety-critical applications such as aerospace or vehicular systems. To achieve the high safety integrity levels required by these applications, systems are designed with fault tolerance mechanisms in order to deliver correct service even in the presence of faults. Faults may, for instance, occur when processors are disturbed by high energy particles such as neutrons or heavy-ions. Such particles may sometimes interfere with the processor and cause a single event upset (SEU) – an error that typically changes the state of a single bit in the system.

In order to validate the correctness and efficiency of their fault tolerance features, safety-critical systems must be thoroughly tested. Fault injection has become an effective technique for the experimental dependability validation of computer systems. The objective of fault injection is to test fault tolerance mechanisms and measure system dependability by introducing artificial faults and errors.

A problem commonly observed during fault injection campaigns is that not all faults fulfil the purpose of disturbing the system [1]. Often 80-90% of randomly injected faults are not activated [1, 2]. A fault placed in a register just before the register is written or faults that are injected into unused memory locations are examples of faults with no possibility of activation. In most tools the location and the time for fault injection are chosen randomly from the complete fault-space, which is typically extremely large. The statistical implication of this is that the cost of obtaining appropriate confidence levels for dependability measures becomes unnecessarily high.

This paper presents a pre-injection analysis technique that is applicable to injection of transient bit-flips into CPU user registers and memory locations. The bit-flip fault model is often used in fault injection experiments to emulate the effects of single event upsets and other transient disturbances.

A bit-flip in any resource will only manifest itself once this resource is read to perform an operation. The pre-injection analysis optimizes[1] the fault-space by *placing bit-flips only in resources[2] immediately before these are read by instructions*. It should be noted that all faults in a given class are collapsed into a single fault in the optimized fault-space. This may cause a bias in the estimated dependability measures (e.g. error detection coverage). One of the objectives of this research is therefore to investigate the magnitude of this bias.

The pre-injection analysis technique was implemented in the GOOFI [3] tool, for Nexus-based fault injection [2, 4, 5], and is also suitable for implementation in other platforms. The effectiveness of the technique was assessed by comparing fault injection results with results obtained by non-optimized fault injection on the same target system. The system is based on the Motorola MPC565 [6] – a microcontroller aimed at the automotive and other control-intensive applications based on the PowerPC architecture. By applying assembly-level knowledge of this architecture it is possible to identify which resources are read by each executed instruction. This information, along with the time of the fault injections, is used to define the optimized fault-space. The fault injection experiments are then conducted by random sampling of faults from the optimized fault-space.

[1] The word *optimize* should not suggest that the optimal fault-space is found but rather an improvement on the usual random approach. Further optimization is therefore achievable.

[2] In this paper we use the word *resource* as a common term for CPU-register, main memory locations and other state-elements where bit-flips may occur.

2. Fault-space Optimization Method

For single bit-flip fault injection, we define a fault-space to be a set of time-location pairs that determines where and when the bit-flip is injected. The time is selected from an interval during the execution of the workload selected for the experiment. The time granularity is based on the execution of machine instructions, i.e. bit-flips can only be injected between the execution of two machine instructions. The complete (non-optimized) fault-space consists of all possible time-location pairs.

Required Input

In order to determine the optimized fault-space it is necessary to gather information about the code of the application and the computer system executing it. This information consists of the assembly code of the application, the Program Counter (PC) trace over time, the definition of which resources are read by each assembly instruction and the effective address of each memory read access.

In our experimental setup, the assembly code is textual information obtained by disassembling the executable binaries of the application, processed automatically by the optimization program. The Program Counter trace and the values of the General Purpose Registers are stored during the execution of the reference run. The effective address of each memory read access is calculated with these values. The definitions of which resources are read by each assembly instruction are built into the optimization program. These were obtained from Motorola's RISC CPU Manual [7] and are available in [8].

Output Produced

The resulting output (the optimized fault-space) consists of a list of possible locations and times for fault injection. The optimization procedure has been adapted to both one-shot applications and control applications executing in loops. Each element on the optimized fault-space contains the following information: control loop index, breakpoint address, number of breakpoint invocations within the control loop, the fault injection location.

The *control loop index* is specific for control applications which execute in cycles. It defines the cycle during which a fault should be injected. For applications that do not execute in loops, the control loop index is always set to one. The *breakpoint address* specifies the breakpoint position inside the control loop and the *number of breakpoint invocations* specifies the number of times this breakpoint should be reached before fault injection.

Performing the Optimization

Using the Program Counter trace over time, the disassembled code of the application is parsed to obtain the sequence of assembly instructions executed. Each of the instructions is then analyzed in order to determine which resources the instruction reads (where faults should be placed).

The most important stage is therefore the identification of the resources read by each instruction. To accomplish this, the first step is to find the definition on the list matching the given instruction. This is done by matching the *opcode* and the *operands*. Then, by examining the possible assembly constructs, the symbols available in the *read list* of the definition are replaced by the resources actually read by the given instruction. Figure 1 illustrates this process with an example of the optimization procedure applied to two instructions.

The instruction at address *39DE8* adds R10 to R11 and stores the result in R5. The definition for this instruction is found in the table and the read list contains rA and rB, respectively, R10 and R11. Since these are the two resources read by this instruction, two new lines are inserted into the fault locations for code address *39DE8* (the control loop index and the breakpoint invocation are assumed to hold the specified values).

The second instruction, at address *39DEC*, fetches the memory word addressed by the effective address (R6) + 24 and stores it in R7. Its definition in the table specifies rA and MEM32(d+rA), respectively, R6 and the 32-bit word at 1000+24, as being read. The value 1000 of R6 is obtained during the reference run. The two resources along with the timings are then inserted into the fault-space.

3. Experimental Setup

Figure 2 describes the evaluation platform used to evaluate the effectiveness of the optimization technique for experiments performed on the jet engine control software, which is one of two workloads investigated in this paper. The GOOFI fault injection tool controls the experiments by using the winIDEA debugging environment [9] in conjunction with the iC3000 debugger. Faults are injected into the MPC565 microcontroller running the control software. In the case of the jet engine controller one computer board was used to run the jet engine control software and one board to execute the model of the jet engine. The experimental setup used for the other workload (an implementation of the quicksort algorithm) used only one computer board.

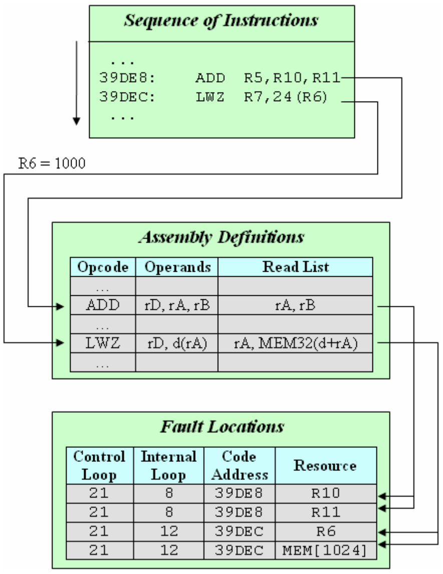

Opcode	Operands	Read List
...		
ADD	rD, rA, rB	rA, rB
...		
LWZ	rD, d(rA)	rA, MEM32(d+rA)
...		

Control Loop	Internal Loop	Code Address	Resource
21	8	39DE8	R10
21	8	39DE8	R11
21	12	39DEC	R6
21	12	39DEC	MEM[1024]

Fig. 1. Example of the optimization procedure

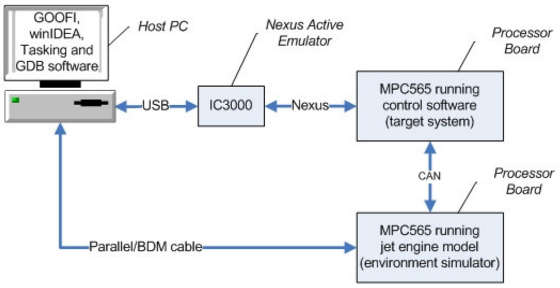
GOOFI, winIDEA, Tasking and GDB software
Host PC
Nexus Active Emulator
Processor Board
USB
IC3000
Nexus
MPC565 running control software (target system)
CAN
Processor Board
Parallel/BDM cable
MPC565 running jet engine model (environment simulator)

Fig. 2. Evaluation platform for the jet engine application

Fault Injection Tool

GOOFI is a fault injection tool developed at the Department of Computer Engineering, Chalmers University of Technology. A plug-in [5] has recently been developed in GOOFI which uses the Nexus [4] port to inject faults on Motorola's MPC565. Nexus is suitable to be used for fault injection [2] since it provides read/write access to the processor's resources and code execution trace capture.

The pre-injection analysis technique was implemented to enhance the existing Nexus fault injection plug-in. The target platform for the current implementation is therefore the MPC565 microcontroller. The technique may however be implemented for any microprocessor.

MPC565 Microcontroller

The MPC565 is a microcontroller developed by Motorola that implements the PowerPC instruction standard architecture. The complete computer system was based on the phyCORE-MPC565 [10] development board. The iSYSTEM iC3000 Active Emulator [11, 12] was used to access the Nexus working environment. The iC3000 emulator was, in its turn, controlled by GOOFI via winIDEA – an integrated development environment.

Workloads

Fault injection campaigns were conducted to evaluate the optimization technique using two different workloads: a sort program using the quicksort algorithm and a jet engine controller. Different campaigns targeting registers and data memory, using both optimized and non-optimized fault selection, were carried out. The technique is fully implemented in the sense that all the assembly instructions executed by the workloads are analysed and all registers and data memory locations where optimization is achievable with this method are considered. The outcome of each fault injection experiment was classified into one of the following categories:

Detected Error – All effective errors signalled by hardware error detection mechanisms in the processor.

Wrong Output – All effective errors not detected by the processor that lead to the production of wrong results.

Non-effective Error – Errors that do not affect the system execution during the chosen experiment time frame.

Quicksort Workload

The quicksort workload is a recursive implementation of the well-known sorting algorithm. It sorts an array containing seven double-precision floats.

The reference run execution takes two minutes during which the processor is being stepped and all the required data is obtained. The optimization procedure takes 20 seconds to complete. Each fault injection experiment takes less than half a minute to perform. During the execution of the reference run for this application, the MPC565 processor executed 34 distinct assembly instructions (opcodes) and a total of 815 instructions.

Jet Engine Controller Workload

This workload is a real-world control application that executes in loops in order to control a jet engine. At the end of each loop the controller has to produce results and exchange information with the engine (sensor values from the engine and actuator commands from the controller).

The execution of the reference run takes almost 12 hours. The optimization procedure takes 10 minutes to complete. Each fault injection experiment is then performed in less than two minutes for the selected configuration (number of control loops and memory locations to be logged).

Forty control loops of execution were logged during each experiment. From these, ten loops (21 to 30) were chosen as possible temporal locations for fault injection (corresponding to 50ms of real-time execution of the controller). During these ten control loops, in the reference run, the MPC565 processor executed 231.097 instructions. A total of 88 different assembly instructions (opcodes) were executed.

Fault Model and Fault Selection

The faults in the non-optimized campaigns were chosen using a uniform distribution. In the case of the optimized campaigns the faults are selected randomly from the optimized fault-space. This implies that the distribution of faults in resources is proportional to the representation of each resource in the optimized fault-space.

Microprocessor registers were selected as spatial locations for fault injection both in the quicksort and in the jet-engine controller campaigns. Memory locations were only targeted using the jet-engine controller. The registers targeted in the non-optimized campaigns are the ones considered by the optimization method: General Purpose Registers (32 registers of 32 bits), Floating Point Registers (32 registers of 64 bits), Link Register, Condition Register, Integer Exception Register and Count Register.

Two limitations of winIDEA (the debugging environment) are important to mention. The floating point registers are only allowed to be injected with faults in the least significant 32 bits. The Floating Point Status And Control Register (FPSCR), targeted by the optimization, is also not available for fault injection.

The fault injection campaigns in memory targeted the stack, heap and all other read/write and read-only data segments of the controller. A total of 100KB of memory were targeted as spatial locations.

The analysis of faults in the code segment was still not implemented and was therefore not studied. The optimization is easily extendable to support faults in the code segment by targeting, in each instruction, the 32-bit memory contents addressed by the Program Counter.

4. Experimental Results

Fault Injection in Registers

Table 1 shows the distribution of the outcomes of faults in the fault injection campaigns targeting microprocessor registers for both the quicksort and the jet engine controller workloads. The quicksort campaigns include approximately the same number of experiments. For the non-optimized jet engine controller campaign, a much higher number of experiments had to be performed in order to increase the confidence in the results.

Campaign		*# Exp.*	*Non-effective*	*Detected*	*Wrong Output*
Quicksort	*Non-optimized*	2739	2603 (95.0%)	83 (3.0%)	53 (2.0%)
	Optimized	2791	1461 (52.3%)	744 (26.7%)	586 (21.0%)
Jet Engine Controller	*Non-optimized*	5708	5457 (95.6%)	200 (3.5%)	51 (0.9%)
	Optimized	1559	964 (61.8%)	466 (29.9%)	129 (8.3%)

Table 1. Distribution of outcomes of fault injection in registers

The percentage of *effective* faults (*detected* or *wrong output*) increases from 5.0% using non-optimized fault selection to 47.7% choosing faults from the optimized fault-space when targeting the quicksort workload. In the jet engine controller this increase is from 4.4% to 38.2%. The improvement in the effectiveness of faults is, therefore, one order of magnitude.

Table 2 shows the estimated error detection coverage obtained in each campaign. We here define error detection coverage as the quotient between the number of detected and the number of effective faults.

Campaign		*Estimated error detection coverage (95% confidence)*
Quicksort	*Non-optimized*	61.0 ± 8.2%
	Optimized	55.9 ± 2.7%
Jet Engine Controller	*Non-optimized*	79.7 ± 5.0%
	Optimized	78.3 ± 3.3%

Table 2. Error detection coverage estimations (registers)

The values of the error detection coverage estimations are quite similar whether applying non-optimized or optimized fault selection. In the optimized campaigns the faults are only injected in the location that will activate them (at the time that the register is read). Since no weights are applied to reflect the length of the data lifecycle on the outcomes of faults, it could be expected that the error detection coverage would be skewed.

Fault Injection in Memory

Fault injection in memory locations was performed only for the jet engine controller. Table 3 shows the distribution of the outcomes of faults for both non-optimized and optimized fault selection.

Campaign		*# Exp.*	*Non-effective*	*Detected*	*Wrong Output*
Jet Engine Controller	*Non-optimized*	6666	6532 (98.0%)	40 (0.6%)	94 (1.4%)
	Optimized	2658	2150 (80.9%)	166 (6.3%)	342 (12.8%)

Table 3. Distribution of outcomes of fault injection in memory

The effectiveness of faults increases from 2.0% using non-optimized fault selection to 19.1% choosing faults from the optimized fault-space. The improvement in the effectiveness of faults is one order of magnitude, similar to one obtained for faults in microprocessor registers.

Table 4 shows the error detection coverage estimated with optimized and non-optimized fault selection.

Campaign		*Estimated error detection coverage (95% confidence)*
Jet Engine Controller	*Non-optimized*	29.9 ± 7.7%
	Optimized	32.7 ± 4.1%

Table 4. Error detection coverage estimations (memory)

We here observe a similar pattern to that observed for microprocessor registers, where the error detection coverage estimation using non-optimized or optimized fault selection is quite similar. In this case the estimation from the non-optimized campaign is not very accurate since the 95% confidence interval is still wide due to the small number of effective faults (2%).

Fault-space Considerations

Applying the optimization method to the fault-space of registers for the jet engine controller resulted in the determination of 7.7×10^6 distinct time-location pairs for bit-flips. All the targeted registers are 32 bit registers[3]. The complete non-optimized fault-space of these registers is obtained by flipping each bit of each register, for each instruction executed. This results in a set containing over 500 million bit-flips.

Campaign		*Size of the fault-space (time-location pairs for bit-flips)*
Jet Engine Controller	*Non-optimized*	5.0×10^8
	Optimized	7.7×10^6
Ratio		1.5%

Table 5. Comparison between fault-space sizes (registers)

In the case of the memory fault-space 3.3×10^6 possible time-location pairs for bit-flips were determined using optimized fault selection. The complete fault-space of memory is obtained by flipping each bit of each memory

[3] Floating Point Registers are 64-bits long limited by winIDEA to the least significant 32-bits.

location used by the program, for each instruction executed. Considering a memory usage of 100KB for data by the jet engine controller, the size of the complete fault-space is near 200 billion bit-flips.

Campaign		*Size of the fault-space (time-location pairs for bit-flips)*
Jet Engine Controller	*Non-optimized*	1.9×10^{11}
	Optimized	3.3×10^{6}
Ratio		0.0017%

Table 6. Comparison between fault-space sizes (memory)

5. Conclusions and Future Work

The study presented in this paper shows the efficiency of eliminating faults with no possibility of activation and determining equivalence classes among faults. A comparison with traditional non-optimized fault selection in the complete fault-space shows an order of magnitude increase in the effectiveness of faults. The fault-space itself is reduced two orders of magnitude for the registers and four to five orders of magnitude for the memory.

All faults targeting the same bit of a given resource, before this resource is read, are considered equivalent. This way, only one representative of these faults is injected. To obtain an accurate estimation of the error detection coverage (or any other dependability measure) it would be necessary to apply a weight corresponding to the number of faults in each equivalence class. However, the error detection coverage estimated by the optimized fault selection is found to be quite similar to the coverage estimated by non-optimized fault selection.

Even though activation of faults is ensured by the optimization technique (activation in the sense that the faulty resources are always utilized) not all faults result in effective errors. This occurs when the data is used in a non-sensitive way by the code (regarding the single bit-flip model). An interesting topic for further studies would be to investigate which activated faults are non-effective and why.

The outcome of a fault is highly dependent on the targeted resource. Faults in some registers were observed to have a greater tendency to cause wrong output while faults in other registers cause detected errors more frequently. This motivates a possible evolution in fault selection by using the results of previous fault injection experiments to select the faults that should be injected next (a combination of pre-injection and post-injection analysis). It would be possible to achieve a faster evaluation of specific error detection mechanisms by injecting faults in the resources that are more likely to activate them.

There is still room for further optimization by analyzing the error propagation. When a bit-flip is copied from one resource onto another and the first resource is overwritten, the fault in the new location is equivalent to the fault in the first location. The implementation of an analysis taking advantage of this has been started and preliminary results show additional improvement.

References

1. H. Madeira and J. Silva, “Experimental Evaluation of the Fail-Silent Behavior in Computers Without Error Masking”, Proc. FTCS-24, June 1994, pp. 350-359.
2. P. Yuste, J. Ruiz, L. Lemus, P. Gil, “Non-Intrusive Software-Implemented Fault Injection in Embedded Systems”, LADC 2003, LNCS 2847, 2003, pp. 23-38.
3. J. Aidemark, J. Vinter, P. Folkesson, J. Karlsson, “GOOFI: Generic Object-Oriented Fault Injection Tool”, Proc. DSN 2001, July 2001, pp. 83-88.
4. IEEE-ISTO, "The Nexus 5001 Forum™ Standard for a Global Embedded Processor Debug Interface", 1999.
5. D. Skarin, J. Vinter, P. Folkesson and J. Karlsson, "Implementation and Usage of the GOOFI MPC565 Nexus Fault Injection Plug-in", Tech. Report No. 04-08, Dept. of Comp. Eng., Chalmers University of Technology, Göteborg, Sweden, 2004.
6. Motorola Inc., “MPC565/MPC566 User’s Manual”, 2nd edition, 2003.
7. Motorola Inc., “RISC Central Processing Unit Reference Manual”, revision 1, 1999.
8. R. Barbosa, J. Vinter, P. Folkesson and J. Karlsson, "Fault Injection Optimization through Assembly-Level Pre-Injection Analysis", Tech. Report No. 04-07, Dept. of Comp. Eng., Chalmers University of Technology, Göteborg, Sweden, 2004.
9. winIDEA – iSystem’s Integrated Development Environment, http://www.isystem.se/products/ide.htm, June 3, 2005.
10. PHYTEC Technology Holding Company, http://www.phytec.com/sbc/32bit/pc565.htm, June 3, 2005.
11. iSYSTEM AB, http://www.isystem.se, June 3, 2005.
12. iC3000 Active Emulator, http://www.isystem.se/products/emulators.htm#three, June 3, 2005.