



**HAL**  
open science

# Moving code analysis from safety to security : attacker model

Soline Ducousso

► **To cite this version:**

Soline Ducousso. Moving code analysis from safety to security : attacker model. Modeling and Simulation. Université Grenoble Alpes [2020-..], 2023. English. NNT: 2023GRALM084. tel-04662172

**HAL Id: tel-04662172**

**<https://theses.hal.science/tel-04662172v1>**

Submitted on 25 Jul 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE

Pour obtenir le grade de

**DOCTEUR DE L'UNIVERSITÉ GRENOBLE ALPES**

École doctorale : MSTII - Mathématiques, Sciences et technologies de l'information, Informatique

Spécialité : Informatique

Unité de recherche : VERIMAG

**Aller de la sûreté à la sécurité en analyse de code : le modèle d'attaquant**

**Moving code analysis from safety to security: attacker model**

Présentée par :

**Soline DUCOUSSO**

Direction de thèse :

**Marie laure POTET**

PROFESSEURE DES UNIVERSITES, GRENOBLE INP

Directrice de thèse

**Sébastien BARDIN**

INGENIEUR DE RECHERCHE, CEA DE PARIS-SACLAY

Co-encadrant de thèse

Rapporteurs :

**KARINE HEYDEMANN**

INGENIEURE HDR, THALES RESEARCH & TECHNOLOGY

**GUILLAUME HIET**

PROFESSEUR, CENTRALESUPELEC RENNES

Thèse soutenue publiquement le **14 décembre 2023**, devant le jury composé de :

**MARIE-LAURE POTET**

PROFESSEURE DES UNIVERSITES, GRENOBLE INP

Directrice de thèse

**KARINE HEYDEMANN**

INGENIEURE HDR, THALES RESEARCH & TECHNOLOGY

Rapporteuse

**GUILLAUME HIET**

PROFESSEUR, CENTRALESUPELEC RENNES

Rapporteur

**LAURE GONNORD**

PROFESSEURE DES UNIVERSITES, GRENOBLE INP

Présidente

**GUILLAUME BOUFFARD**

INGENIEUR DE RECHERCHE, ANSSI

Examineur

Invités :

**SEBASTIEN BARDIN**

INGENIEUR DE RECHERCHE, CEA DE PARIS-SACLAY



---

---

# Résumé

## Aller de la Sûreté à la Sécurité en Analyse de Code: le Modèle d'Attaquant

Des travaux majeurs ont été réalisés, ces dernières décennies, dans le domaine de l'analyse de programme, tirant parti de techniques telles que l'exécution symbolique, l'analyse statique, l'interprétation abstraite ou la vérification bornée de modèles, pour chasser les bogues et vulnérabilités logicielles, conduisant à l'adoption de ces techniques par de grandes entreprises. Les bogues étant des points d'entrées pour les attaques, les éliminer est un premier pas vers une meilleure sécurité logicielle. Cependant, ces techniques reposent sur un modèle d'attaquant seulement capable de concevoir des entrées malicieuses, exploitant des cas particuliers dans le code lui-même, typiquement un fichier formaté de manière à déclencher un débordement mémoire dans les routines de traitement. Un modèle d'attaquant plutôt faible, alors qu'un attaquant avancé est capable de perturber l'exécution d'un programme en exploitant des vecteurs d'attaques tels que les injections de fautes matérielles, les attaques micro-architecturales, les attaques matérielles contrôlées par logiciel, et n'importe quelle combinaison de ces vecteurs d'attaques.

Dans cette thèse, notre objectif est de concevoir une technique automatique et efficace pour raisonner sur l'impact d'un attaquant avancé sur les propriétés de sécurité d'un programme. Nous proposons l'atteignabilité adversariale, un formalisme qui étend la notion d'atteignabilité en y incluant les capacités de l'attaquant. Nous avons construit un nouvel algorithme, l'exécution symbolique adversariale, pour répondre au problème de l'atteignabilité adversariale sous l'angle de la recherche de bogues (vérification bornée). Notre algorithme évite l'augmentation significative de l'espace d'états à analyser due nouvelles capacités de l'attaquant, grâce à un nouvel encodage non branchant de celles-ci sous forme de fautes injectées. Nous le montrons correct et  $k$ -complet pour l'atteignabilité adversariale. De plus, nous avons imaginé deux optimisations visant à réduire le nombre de fautes : la détection précoce de saturation et l'injection à la demande.

Nous proposons une implémentation de l'exécution symbolique adversariale au niveau binaire intégrée à l'outil BINSEC. Notre évaluation expérimentale repose sur des programmes usuels du domaine des fautes matérielles et des cartes à puce. Nos expériences montrent un gain de performance significatif par rapport à l'état de l'art, en moyenne d'un facteur  $\times 20$  et  $\times 400$  pour 1 et 2 fautes respectivement, avec un gain similaire en termes de chemins explorés. De plus, notre approche passe à l'échelle jusqu'à considérer 10 fautes alors que l'état de l'art atteint le temps limite pour 3 fautes. Nous montrons également l'utilisabilité de notre méthode dans différents scénarios de sécurité tels que reproduire une attaque BellCoRe sur CRT-RSA, chercher des attaques sur un programme protégé par la contre-mesure SecSwift, ou évaluer la robustesse d'un réseau de neurones. D'autre part, nous avons exploré le cas du programme de démarrage de WooKey en rejouant des attaques et en évaluant des contre-mesures proposées. En particulier, nous avons trouvé une attaque non rapportée précédemment et avons proposé un nouveau correctif aux développeurs.

**Mots-clefs:** sécurité logicielle · méthodes formelles · analyse de code · code binaire · vulnérabilités.

---

# Abstract

## Moving Code Analysis from Safety to Security: the Attacker Model

Major works have delved into program analysis over the last decades, leveraging techniques such as symbolic execution, static analysis, abstract interpretation or bounded model checking, to hunt for software vulnerabilities and bugs in programs, or to prove their absence, leading to industrial adoption in some leading companies. As bugs are an attack entry point, removing them is a first step towards better software security. Yet, they are based on the standard concept of reachability and represent an attacker able to craft smart, legitimate input, through legitimate input sources of the program, exploiting corner cases in the code itself, a rather weak threat model, typically a file formatted to trigger a buffer overflow in data processing. Tools only looking for bugs and software vulnerabilities may deem a program secure while the bar remains quite low for an advanced attacker, able for example to take advantage of attack vectors such as (physical) hardware fault injections, micro-architectural attacks, software-based hardware attacks like Rowhammer, or any combination of vectors.

In this thesis, our goal is to devise a technique to automatically and efficiently reason about the impact of an advanced attacker onto a program security properties. We propose adversarial reachability, a formalism extending standard reachability to reason about a program execution in the presence of an advanced attacker. As equipping the attacker with new capabilities significantly increases the state space of the program under analysis, we build a new algorithm based on symbolic techniques, named adversarial symbolic execution, to address the adversarial reachability problem from the bug-finding point of view (bounded verification). Our algorithm prevents path explosion thanks to a new forkless encoding of attacker capabilities, modeled as faults. We show it correct and  $k$ -complete with respect to adversarial reachability. To improve the performance further, we design two new optimizations to reduce the number of injected faults while keeping the same attacker power: Early Detection of fault Saturation and Injection On Demand.

We propose an implementation of our techniques for binary-level analysis, on top of the BINSEC framework. We systematically evaluate its performances against prior work, using a standard fault injection benchmark from physical fault attacks and smart cards. Experiments show a very significant performance gain against prior approaches, for example up to  $\times 20$  and  $\times 400$  times on average for 1 and 2 faults respectively - with a similar reduction in the number of adversarial paths. Moreover, our approach scales up to 10 faults whereas the state-of-the-art starts to timeout for 3 faults. We also show the use of our method in a number of different security scenarios such as reproducing a BellCoRe attack with one reset fault on CRT-RSA or looking for attacks on a program protected by the SecSwift countermeasure. In addition, we perform a security analysis of the well-tested WooKey bootloader and demonstrate the ability of our analysis to find attacks and evaluate countermeasures in real-life security scenarios. We were especially able to find an attack not reported before on a recently proposed patch, and proposed a new patch to the developers.

**Keywords:** software security · formal methods · code analysis · binary code · vulnerabilities.

# 91 Contents

92	Résumé . . . . .	iii
93	Abstract . . . . .	iv
94	<b>List of Figures</b>	<b>ix</b>
95	<b>List of Tables</b>	<b>xi</b>
96	<b>1 Introduction</b>	<b>1</b>
97	1.1 Context . . . . .	1
98	1.2 Problem . . . . .	2
99	1.3 Goal & Challenges . . . . .	3
100	1.4 Proposal . . . . .	3
101	1.5 Contributions . . . . .	4
102	1.6 Impact and Perspectives . . . . .	5
103	1.7 Manuscript Outline . . . . .	5
104	<b>2 Context and Motivation</b>	<b>7</b>
105	2.1 Security of Information Technology Systems . . . . .	8
106	2.1.1 Definition . . . . .	8
107	2.1.2 Security Properties . . . . .	8
108	2.1.3 Wide-spread Use . . . . .	9
109	2.1.4 Attack Surface and Attackers . . . . .	11
110	2.1.5 Ensure Security . . . . .	13
111	2.1.6 Example of Attack on a Security System . . . . .	15
112	2.1.7 Conclusion . . . . .	17
113	2.2 Capabilities of a Powerful Attacker . . . . .	17
114	2.2.1 Side Channel Attacks . . . . .	17
115	2.2.2 Software Attacks . . . . .	18
116	2.2.3 Hardware Fault Injection Attacks . . . . .	19
117	2.2.4 Software-Implemented Hardware Attacks . . . . .	21
118	2.2.5 Micro-Architectural Attacks . . . . .	22
119	2.2.6 Man-At-The-End (MATE) Attacks . . . . .	23
120	2.2.7 Summary . . . . .	24
121	2.3 Conclusion . . . . .	24
122	<b>3 Background</b>	<b>25</b>
123	3.1 Program Analysis Techniques . . . . .	26
124	3.1.1 Overview . . . . .	26

125	3.1.2	Program Analysis is Undecidable . . . . .	27
126	3.1.3	Reachability Property . . . . .	29
127	3.1.4	Link to the Thesis . . . . .	30
128	3.2	Symbolic Execution . . . . .	30
129	3.2.1	Overview . . . . .	30
130	3.2.2	Symbolic Execution Algorithm . . . . .	31
131	3.2.3	Limitations . . . . .	32
132	3.2.4	Link to the Thesis . . . . .	33
133	3.3	Binary Code Analysis . . . . .	33
134	3.3.1	Binary VS Source Code Analysis . . . . .	33
135	3.3.2	Challenges of Binary Analysis . . . . .	34
136	3.3.3	Link to the Thesis . . . . .	34
137	3.4	Program Analysis Techniques for Fault Injection . . . . .	34
138	3.4.1	Simulation . . . . .	34
139	3.4.2	Mutant Generation . . . . .	35
140	3.4.3	Forking Techniques . . . . .	35
141	3.4.4	Related Work . . . . .	36
142	3.4.5	Link to the Thesis . . . . .	36
143	<b>4</b>	<b>Adversarial Reachability</b> . . . . .	<b>37</b>
144	4.1	Attacker Model . . . . .	39
145	4.1.1	Advanced Attacker . . . . .	39
146	4.1.2	Attacker Actions . . . . .	39
147	4.1.3	Fault Budget . . . . .	39
148	4.1.4	Attacker Goal . . . . .	40
149	4.2	Adversarial Reachability . . . . .	40
150	4.2.1	Reminder: Standard Reachability . . . . .	40
151	4.2.2	Adversarial Reachability Definition . . . . .	40
152	4.2.3	Properties . . . . .	45
153	4.2.4	Discussion . . . . .	46
154	4.2.5	Conclusion . . . . .	46
155	4.3	Forkless Adversarial Symbolic Execution (FASE) . . . . .	47
156	4.3.1	Overview . . . . .	47
157	4.3.2	Forkless Fault Encoding . . . . .	47
158	4.3.3	Fault Injection Algorithm . . . . .	50
159	4.4	Optimizations . . . . .	54
160	4.4.1	Early Detection of Fault Saturation (EDS) . . . . .	54
161	4.4.2	Injection on Demand (IOD) . . . . .	56
162	4.4.3	Combination of Optimizations . . . . .	58
163	4.5	Discussion . . . . .	59
164	4.5.1	Fault Model Support - Formalization VS Algorithm . . . . .	59
165	4.5.2	Forkless Faults and Multi-Fault Analysis . . . . .	60
166	4.5.3	Forkless Encoding for Other Properties . . . . .	60
167	4.5.4	Forkless Encoding for Other Formal Methods . . . . .	61
168	4.5.5	Forkless Encoding and Instrumentation . . . . .	62
169	4.6	Related Work . . . . .	62
170	4.6.1	Fault Model Support . . . . .	62
171	4.6.2	Multiple Fault Analysis . . . . .	62
172	4.6.3	The Attacker in Different Security Fields . . . . .	64

173	4.6.4	Extending Existing Formalisms . . . . .	64
174	4.7	Conclusion . . . . .	65
175	<b>5</b>	<b>The BINSEC/ASE Prototype</b>	<b>67</b>
176	5.1	Overview . . . . .	69
177	5.2	Background: the BINSEC Tool . . . . .	69
178	5.2.1	BINSEC Presentation . . . . .	70
179	5.2.2	General Work-Flow . . . . .	70
180	5.2.3	Summary . . . . .	73
181	5.3	BINSEC/ASE Implementation . . . . .	73
182	5.3.1	BINSEC/ASE Overview . . . . .	73
183	5.3.2	ASE Implementation . . . . .	74
184	5.3.3	Forkless Fault Models . . . . .	77
185	5.3.4	Early Detection of Fault Saturation (EDS) . . . . .	81
186	5.3.5	Injection On Demand (IOD) . . . . .	82
187	5.3.6	Sub-fault Simplification . . . . .	82
188	5.3.7	Forking Fault Models . . . . .	83
189	5.3.8	Conclusion . . . . .	84
190	5.4	User Guide: a Methodology to Analyse a New Program . . . . .	86
191	5.4.1	Running Example . . . . .	86
192	5.4.2	Analysis Goal . . . . .	88
193	5.4.3	Configuration . . . . .	93
194	5.4.4	Reading BINSEC/ASE Output . . . . .	93
195	5.4.5	Analysis Process . . . . .	96
196	5.4.6	Summary . . . . .	96
197	5.5	Developer Guide: a Methodology to Add a New Fault Model . . . . .	97
198	5.5.1	Defining the New Fault Model . . . . .	97
199	5.5.2	Implementation . . . . .	98
200	5.5.3	Dedicated Metrics . . . . .	99
201	5.5.4	Testing . . . . .	99
202	5.5.5	Summary . . . . .	99
203	5.6	Discussion . . . . .	100
204	5.6.1	BINSEC/ASE Limitations . . . . .	100
205	5.6.2	Faults on Intermediate Representation . . . . .	101
206	5.6.3	Permanent VS Transient Faults . . . . .	101
207	<b>6</b>	<b>Experimental Evaluation</b>	<b>103</b>
208	6.1	Evaluation Overview . . . . .	104
209	6.1.1	Research Questions . . . . .	104
210	6.1.2	Experimental Setting . . . . .	105
211	6.1.3	Artifact Availability . . . . .	107
212	6.2	Correctness and K-completeness ( <b>RQ1</b> ) . . . . .	107
213	6.3	FASE Evaluation for Arbitrary Data Faults ( <b>RQ2</b> ) . . . . .	108
214	6.3.1	Scalability ( <b>RQ2.1</b> ) . . . . .	108
215	6.3.2	Impact of Optimizations ( <b>RQ2.2</b> ) . . . . .	110
216	6.3.3	Comparison of the Different Forkless Encodings ( <b>RQ2.3</b> ) . . . . .	116
217	6.4	FASE Evaluation of Other Fault Models ( <b>RQ3</b> ) . . . . .	118
218	6.4.1	FASE Evaluation of Reset Faults ( <b>RQ3.1</b> ) . . . . .	119
219	6.4.2	FASE Evaluation of Bit-Flip Faults ( <b>RQ3.2</b> ) . . . . .	120
220	6.4.3	FASE Evaluation of Test Inversion Faults ( <b>RQ3.3</b> ) . . . . .	121

221	6.4.4	FASE Evaluation of Instruction Skip Faults ( <b>RQ3.4</b> ) . . . . .	122
222	6.4.5	Summary . . . . .	123
223	6.5	Forkless Faults in Instrumentation ( <b>RQ4</b> ) . . . . .	124
224	6.5.1	Experimental Settings . . . . .	125
225	6.5.2	Scalability . . . . .	125
226	6.5.3	Conclusion . . . . .	126
227	6.6	Security Scenarios . . . . .	127
228	6.6.1	CRT-RSA . . . . .	128
229	6.6.2	Secret-keeping Machine . . . . .	129
230	6.6.3	SecSwift Countermeasure . . . . .	130
231	6.6.4	Neural Network . . . . .	131
232	6.6.5	Security Scenarios Feedback . . . . .	132
233	6.7	Case Study: WooKey Bootloader . . . . .	132
234	6.7.1	Presentation of WooKey . . . . .	132
235	6.7.2	Security Scenario and Goal of our Study . . . . .	133
236	6.7.3	Analyze Key Parts of Wookey . . . . .	133
237	6.7.4	Analyze a Security Patch of WooKey . . . . .	136
238	6.7.5	Propose a New Patch and Evaluate It . . . . .	136
239	6.7.6	Other Attacks on WooKey . . . . .	137
240	6.7.7	Case Study Conclusion . . . . .	138
241	<b>7</b>	<b>Conclusion</b> . . . . .	<b>139</b>
242	7.1	Conclusion . . . . .	139
243	7.2	Perspectives . . . . .	140
244	<b>A</b>	<b>Éléments de traduction en français</b> . . . . .	<b>I</b>
245	A.1	Chapitre 1 : Introduction . . . . .	I
246	A.2	Résumés de chaque chapitre . . . . .	III
247	A.2.1	Chapitre 2 : Contexte et Motivation . . . . .	III
248	A.2.2	Chapitre 3 : Préambule . . . . .	IV
249	A.2.3	Chapitre 4 : Atteignabilité Adversariale . . . . .	IV
250	A.2.4	Chapitre 5 : le Prototype BINSEC/ASE . . . . .	IV
251	A.2.5	Chapitre 6 : Evaluation Expérimentale . . . . .	IV
252	A.3	Chapitre 7 : Conclusion . . . . .	V
253	<b>B</b>	<b>Additional Experimental Data</b> . . . . .	<b>VII</b>
254	B.1	FASE Evaluation of Reset Faults ( <b>RQ B1</b> ) . . . . .	VII
255	B.1.1	Impact of Optimizations ( <b>RQ B1.1</b> ) . . . . .	VII
256	B.1.2	Comparison of the Different Forkless Encodings ( <b>RQ B1.2</b> ) . . . . .	X
257	B.2	FASE Evaluation of Bit-Flip Faults ( <b>RQ B2</b> ) . . . . .	XI
258	B.3	FASE Evaluation of Instruction Skip Faults ( <b>RQ B3</b> ) . . . . .	XIV
259	B.4	FASE Optimizations Summary . . . . .	XV
260	B.5	Influence of Solver on Encoding Operators ( <b>RQ B4</b> ) . . . . .	XV
261	<b>C</b>	<b>Instrumentation Details</b> . . . . .	<b>XVII</b>
262		<b>Bibliography</b> . . . . .	<b>XXI</b>

# 263 List of Figures

264	2.1	Motivating example, inspired by VerifyPIN [DPP+16]	16
265	2.2	Example of a buffer overflow vulnerability.	18
266	3.1	Illustration [Dan21] of an analysis $\mathcal{A}$ that over-approximates the behaviors of a program $\mathbb{P}$	27
267			
268	3.2	Illustration [Dan21] of an analysis $\mathcal{A}$ that under-approximates the behaviors of a program $\mathbb{P}$	28
269			
270	3.3	Mutant generation transformation in pseudo-code	35
271	3.4	Forking technique transformation in pseudo-code	36
272	4.1	Forkless injection technique	48
273	5.1	Overview of BINSEC workflow for symbolic execution	70
274	5.2	Overview of BINSEC/ASE workflow	73
275	5.3	Illustration of sub-fault simplification	83
276	5.4	User guide workflow	86
277	5.5	Running example, inspired by VerifyPIN [DPP+16]	87
278	5.6	Running example: disassembly of the main function	89
279	5.7	Running example: disassembly of the verifyPIN function	91
280	5.8	Running example: disassembly of the byteArrayCompare function	92
281	5.9	Running example: BINSEC/ASE configuration file	93
282	5.10	Running example: BINSEC/ASE statistics output	95
283	5.11	Running example: BINSEC/ASE attack path	96
284	6.1	FASE-IOD and forking analysis time comparison for arbitrary data faults (RQ2.1)	109
285			
286	6.2	FASE-IOD and forking number of explored paths comparison for arbitrary data faults (RQ2.1)	110
287			
288	6.3	FASE optimizations, analysis time for arbitrary data faults (RQ2.2)	112
289	6.4	FASE optimizations, average solving time per query for arbitrary data faults (RQ2.2)	114
290			
291	6.5	Forkless and forking analysis time for arbitrary data faults in instrumentation with Klee (RQ4)	126
292			
293	6.6	FASE and forking number of explored paths for arbitrary data faults in instrumentation with Klee (RQ4)	127
294			
295	6.7	functions of WooKey’s bootloader, with [LFBP21] fixes and <b>our patch</b>	135
296	6.8	CM1 double test pattern	137

## LIST OF FIGURES

---

297	C.1 Instrumented VerifyPIN main function . . . . .	XVIII
298	C.2 Instrumentation utility functions . . . . .	XIX

## List of Tables

300	2.1	Attacker capabilities given by software attacks . . . . .	19
301	2.2	Attacker capabilities given by hardware fault injection attacks . . . . .	21
302	2.3	Attacker capabilities given by software-implemented fault injection attacks . . . . .	22
303	2.4	Attacker capabilities given by micro-architectural attacks . . . . .	23
304	4.1	Forkless encoding variants for arbitrary data faults . . . . .	48
305	4.2	Forkless encodings for various fault models . . . . .	49
306	4.3	Fault model support . . . . .	63
307	5.1	Example of C and Intel x86 instructions translated to DBA . . . . .	72
308	5.2	Simplification rules added to BINSEC/ASE . . . . .	77
309	5.3	Arbitrary data encodings and their associated activation constraint . . . . .	79
310	5.4	Reset encodings and their associated activation constraint . . . . .	79
311	5.5	Bit-flip encoding and its activation constraint . . . . .	79
312	5.6	Instruction skip encoding for each dba instruction . . . . .	81
313	5.7	Ocaml interfaces . . . . .	99
314	6.1	Benchmarks characteristics and statistics of a standard SE analysis . . . . .	106
315	6.2	FASE vulnerability results compared to benchmark ground truth (RQ1) . . . . .	107
316	6.3	FASE-IOD and forking analysis time comparison for arbitrary data faults (RQ2.1) . . . . .	109
317	6.4	FASE-IOD and forking number of explored paths comparison for arbitrary data faults (RQ2.1) . . . . .	110
318	6.5	FASE optimizations, analysis time for arbitrary data faults (RQ2.2) . . . . .	111
319	6.6	FASE optimizations, number of explored paths for arbitrary data faults (RQ2.2) . . . . .	112
320	6.7	FASE optimizations, number of queries created and sent to the solver for arbitrary data faults (RQ2.2) . . . . .	113
321	6.8	FASE optimizations, average solving time per query for arbitrary data faults (RQ2.2) . . . . .	114
322	6.9	FASE optimizations average number of ite operator per query for arbitrary data faults (RQ2.2) . . . . .	115
323	6.10	Analysis time for sub-fault simplification for arbitrary data faults (RQ2.2) . . . . .	115
324	6.11	Average solving time per query for sub-fault simplification for arbitrary data faults (RQ2.2) . . . . .	116
325	6.12	Average number of ite per query for sub-fault simplification for arbitrary data faults (RQ2.2) . . . . .	116
326			
327			
328			
329			
330			
331			
332			
333			

334	6.13	Arbitrary data encodings and their associated activation constraint . . .	117
335	6.14	Encodings, analysis time for arbitrary data faults (RQ2.3) . . . . .	117
336	6.15	Encodings, number of queries created and sent to the solver for arbitrary	
337		data faults (RQ2.3) . . . . .	118
338	6.16	Encodings, average solving time per query for arbitrary data faults (RQ2.3)	118
339	6.17	Analysis time for reset faults (RQ3.1) . . . . .	119
340	6.18	Number of explored paths for reset faults (RQ3.1) . . . . .	120
341	6.19	Analysis time for bit-flip faults (RQ3.2) . . . . .	120
342	6.20	Number of explored paths for bit-flip faults (RQ3.2) . . . . .	121
343	6.21	Analysis time for test inversion faults (RQ3.3) . . . . .	121
344	6.22	Number of explored paths for test inversion faults (RQ3.3) . . . . .	121
345	6.23	Analysis time for instruction skip faults (RQ3.4) . . . . .	122
346	6.24	Number of explored paths for instruction skip faults (RQ3.4) . . . . .	123
347	6.25	Summary table comparing FASE and the forking technique (RQ3) . . .	124
348	6.26	Analysis time for arbitrary data faults in instrumentation with Klee (RQ4)	126
349	6.27	Number of explored paths for arbitrary data faults in instrumentation	
350		with Klee (RQ4) . . . . .	127
351	6.28	Benchmarks characteristics and statistics of a standard SE analysis . .	128
352	6.29	Summary of the CRT-RSA security scenario . . . . .	128
353	6.30	Summary of the Secret keeping machine security scenario . . . . .	130
354	6.31	Summary of the SecSwift security scenario . . . . .	131
355	6.32	Summary of the Neural Network security scenario . . . . .	132
356	6.33	Benchmarks characteristics and statistics of a standard SE analysis . .	134
357	6.34	Table summarizing the effects of countermeasures . . . . .	135
358	6.35	Summary of the WooKey bootloader case study . . . . .	138
359	B.1	Analysis time for reset faults (RQ B1.1) . . . . .	VIII
360	B.2	Number of queries created and sent to the solver for reset faults (RQ B1.1)	IX
361	B.3	Number of queries created and sent to the solver comparison for reset	
362		fault encodings (RQ B1.2) . . . . .	IX
363	B.4	Average solving time per query for reset faults (RQ B1.1) . . . . .	X
364	B.5	Analysis time for reset faults (RQ B1.2) . . . . .	X
365	B.6	Average solving time per query for reset faults (RQ B1.2) . . . . .	XI
366	B.7	Analysis time for bit-flip faults (RQ3.2) . . . . .	XII
367	B.8	Average solving time per query for bit-flip faults (RQ B2) . . . . .	XII
368	B.9	Number of queries created and sent to the solver for bit-flip faults (RQ	
369		B2) . . . . .	XIII
370	B.10	Number of queries created and sent to the solver comparison for instruc-	
371		tion skip faults (RQ B3) . . . . .	XIII
372	B.11	Analysis time for instruction skip faults (RQ B3) . . . . .	XIV
373	B.12	Average solving time per query for instruction skip faults (RQ B3) . . .	XV
374	B.13	Analysis time and average solving time per query of arbitrary data faults	
375		with different solvers (RQ B4) . . . . .	XVI
376	C.1	Instrumentation examples . . . . .	XVII

377 Chapter **1**

378 Introduction

379 **Contents**

---

380	<b>1.1 Context</b> . . . . .	<b>1</b>
381		
382	<b>1.2 Problem</b> . . . . .	<b>2</b>
383	<b>1.3 Goal &amp; Challenges</b> . . . . .	<b>3</b>
384	<b>1.4 Proposal</b> . . . . .	<b>3</b>
385	<b>1.5 Contributions</b> . . . . .	<b>4</b>
386	<b>1.6 Impact and Perspectives</b> . . . . .	<b>5</b>
387	<b>1.7 Manuscript Outline</b> . . . . .	<b>5</b>

---

388  
389  
390

391 **1.1 Context**

392 **Programs.** *Programs* are everywhere in our modern lives: our phones in our pockets,  
393 the computers we work on, our TVs at home, our local grocery store systems, but also  
394 the chips in our credit cards, systems in our cars, buses, trains, boats and planes, our  
395 hospital system, our water delivery system and our power grids, etc. Programs receive  
396 input from the world (or us), then perform a function and have an output, either  
397 as new data stored somewhere, as information display or as commands for actuators.  
398 Some features can be *security-oriented*, meaning that their failure can have important  
399 consequences, material-, money-, human- or information-wise.

400 **Vulnerabilities.** Because programs are human-written, some errors are made. The  
401 specification for the developers can have holes or inconsistencies. The developers can  
402 make mistakes in their development work, or make incorrect assumptions about li-  
403 braries, interpreters or compilers behaviors. This leaves programs with *bugs* that can  
404 be exploited. Program protections can be added to reduce the attack surface, but  
405 are often incomplete when faced with an advanced attacker. Furthermore, *advanced*  
406 *attackers* can not only leverage existing bugs present in a program but also perform  
407 active fault injection attacks, injecting faulty behavior directly into the execution of  
408 the program. New attack techniques are regularly discovered, increasing an advanced  
409 attacker’s capabilities.

410 **Attackers.** Some entities have an interest in exploiting a program’s incorrect behav-  
411 iors. Malicious groups such as government agencies, mafias or companies are known  
412 for, or suspected of, cyber-attacks like ransomware, blackmail, theft and corporate es-  
413 pionage. But they are not the only ones testing programs. Script kiddies do it for fun  
414 and learning. Security researchers challenge programs and further the state-of-the-art  
415 knowledge in security. Companies test the security of their products with the help of  
416 auditors before market release.

417 **Program Analysis.** To prevent attackers from exploiting a program, major works have  
418 delved into *automated program analysis* over the last decades, leveraging techniques  
419 such as symbolic execution [CS13, GLM12, BGM13], static analysis [Fac], abstract in-  
420 terpretation [CGJ+03] or bounded model checking [CBRZ01], to hunt software vulner-  
421 abilities and bugs in programs, or to prove their absence [CCF+05, KKP+15], leading  
422 to industrial adoption in some leading companies [BGM13, Fac, BCLR04, KKP+15,  
423 LRV+22]. Ensuring a program works as intended, i.e. without bugs, is the notion of  
424 *safety*. As bugs are also an attack entry point, removing them ensures program safety  
425 and is a first step towards better software *security*.

## 426 1.2 Problem

427 **Source Code VS Binary Analysis.** Many program analysis techniques are performed  
428 at the source code level, to allow better developer adoption and understanding of the  
429 analysis results. However, what is actually executed by a computer is the binary  
430 program (for compiled languages), which can differ from the source code, binary code  
431 is where the real vulnerabilities lay. Binary analysis is a harder problem due to the  
432 loss of information such as types or data structures, and requires about ten times more  
433 lines than source code to express equivalent behavior.

434 **Advanced Attackers Across Security Fields.** In addition to bugs, a program has other  
435 weaknesses during its execution. The more resourceful attackers, which we call ad-  
436 vanced attackers, are capable of using any attack vector to exploit a program, and per-  
437 form multiple actions through the course of an attack. For instance, micro-architectural  
438 attacks misuse complex micro-architectural behaviors such as cache systems or branch  
439 prediction. The running hardware can even be physically disturbed with hardware fault  
440 injection means. Those complex vulnerabilities, exploitable by an advanced attacker,  
441 are rarely taken into account in automated program analysis techniques.

442 **Security Program Analysis.** The existing program analysis techniques focus on finding  
443 bugs. It appears that all these methods consider a rather weak threat model, where  
444 the attacker can only craft smart “inputs of death” through legitimate input sources of  
445 the program, exploiting corner cases in the code itself. In order to perform a security  
446 evaluation, meaning replay advanced attacks at will, varying different parameters, per-  
447 form security evaluation and program protection, there is a need for security program  
448 analysis that would combine automated program analysis and attacker model. Tools  
449 only looking for bugs and software vulnerabilities may deem a program secure while  
450 the bar remains quite low for an advanced attacker.

451 **Efficient Program Analysis Techniques.** Taking into account all possible actions of an  
452 advanced attacker and their multiplicity in an analysis represents an added complexity,  
453 which leads to a state explosion in existing techniques. This limits the ability to replay  
454 attacks at will and to consider different attacker models.

### 455 1.3 Goal & Challenges

456 *Our goal is to devise a technique to automatically and efficiently reason about the*  
 457 *impact of an advanced attacker onto a program security properties, where the standard*  
 458 *program analysis techniques only support an attacker crafting smart legitimate inputs.*

459 **Represent an Advanced Attacker in a Formal Framework.** To represent the impact  
 460 of an advanced attacker on a program, the first step is to devise a model of such an  
 461 attacker. Then, we need to provide a *formal framework* to study what an advanced  
 462 attacker can do to attack a program under study. Interestingly, while it has been little  
 463 studied for program-level analysis, such frameworks are routinely used in cryptographic  
 464 protocol verification [Cer01, BCL14]. This formal framework can then be used as a  
 465 theoretical base for an automated program analysis technique for security.

466 Existing techniques tend to be domain specific and hence, represent an attacker  
 467 usually able to exploit one attack vector, with only limited attack power. Our goal  
 468 is to build a *generic model* able to represent an advanced attacker able to leverage  
 469 multiple attack vectors.

470 **Efficient and Generic Code Analysis Techniques for Security.** The second challenge  
 471 is to design an efficient algorithm to assess the vulnerability of a program to a given  
 472 attacker model, when adding capabilities to the attacker naturally gives rise to a sig-  
 473 nificant explosion of program states to consider – especially in the case of an attacker  
 474 able to perform multiple actions.

475 The rare prior work in the field, mostly focused on encompassing physical fault in-  
 476 jections for high-security devices, rely mostly on *mutant generation* [CCG13, RG14,  
 477 GWJLL17, CDFG18, GWJL20] or *forking analysis* [PMPD14, BBC<sup>+</sup>14, BHE<sup>+</sup>19,  
 478 LFBP21, Lan22], yielding scalability issues. Moreover, most of them are limited to  
 479 a few predefined *fault models*, i.e. patterns of the injected faulty behavior, and do not  
 480 propose a formalization of the underlying problem.

### 481 1.4 Proposal

482 **Adversarial Reachability.** We propose a model of an advanced attacker in terms of its  
 483 capability to change program behavior and the goal of the attack as the violation of a  
 484 security property of the program. We extend the usual transition system representing  
 485 the execution of a program with new transitions modeling the attacker actions. Then,  
 486 we extend the standard concept of reachability to reason about a program execution in  
 487 the presence of an advanced attacker. *Adversarial reachability* is a formalism expressing  
 488 the reachability of the attacker goal in this modified transition system.

489 **Adversarial Symbolic Execution.** We build a new algorithm based on symbolic tech-  
 490 niques, named *adversarial symbolic execution* (ASE), to address the adversarial reach-  
 491 ability problem from the bug-finding point of view (bounded verification). Our algo-  
 492 rithm is generic in terms of supported attacker capabilities and prevents path explosion  
 493 thanks to a new *forkless* encoding of attacker capabilities, represented as faults in the  
 494 program. We show it correct and k-complete with respect to adversarial reachabil-  
 495 ity. To improve the performance further, we design two new optimizations to reduce  
 496 the number of injected faults: Early Detection of fault Saturation and Injection On  
 497 Demand.

498 **Tool Prototype: BINSEC/ASE.** We implemented ASE by modifying an existing sym-  
499 bolic execution engine for binary programs, BINSEC, building BINSEC/ASE.

## 500 1.5 Contributions

501 As a summary, we claim the following contributions:

- 502 – We propose a model for a software-level advanced attacker and we formalize its  
503 impact on a program’s security properties in Chapter 4. We named it the *adver-*  
504 *sarial reachability problem*, extending standard reachability to take into account  
505 an advanced attacker in a transition system augmented with transitions repre-  
506 senting attacker actions. We define the associated correctness and completeness  
507 definitions;
- 508 – We describe a new symbolic exploration method, *adversarial symbolic execution*,  
509 to answer adversarial reachability in Chapter 4. The goal of this analysis tech-  
510 nique is to automatically and efficiently reason about the impact of an advanced  
511 attacker on a program, and be generic in the supported attacker model. It fea-  
512 tures a novel forkless fault encoding to prevent path explosion and two optimiza-  
513 tion strategies to reduce the number of fault injections without loss of generality  
514 with respect to the attacker model. We establish adversarial symbolic execution  
515 correctness and completeness for adversarial reachability;
- 516 – We propose an implementation of adversarial symbolic execution for binary-level  
517 analysis, on top of the BINSEC framework [DBT<sup>+</sup>16] in a *tool called BIN-*  
518 *SEC/ASE*, described in Chapter 5. In Chapter 6, we systematically evaluate  
519 its performances against prior work, using a standard SWIFI benchmark from  
520 physical fault attacks and smart cards. We highlight the interest and feasibility  
521 of our technique by exploring different security scenarios. Finally, we perform  
522 a case study on the WooKey bootloader, with a vulnerability assessment of two  
523 implementations and we propose our own patch for an attack not reported before.

## 524 Publication

525 Part of the work presented in this thesis has been published in an paper at ESOP  
526 2023: Soline Ducouso, Sébastien Bardin, and Marie-Laure Potet, "Adversarial Reach-  
527 ability for Program-level Security Analysis", *European Symposium on Programming*,  
528 2023 [DBP23].

## 529 Talks & presentations

530 Part of this work has been presented at:

- 531 – *Journées du Groupe de Travail Méthodes Formelles pour la Sécurité* (GT MFS)  
532 in 2022 and 2023,
- 533 – the seminar *La Cybersécurité sur un plateau* (CoaP) in February 2023,
- 534 – *Rendez-Vous de la Recherche et de l’Enseignement de la Sécurité des Systèmes*  
535 *d’Information* (RESSI) in 2023,
- 536 – *Journée thématique sur les attaques par injection de fautes* (JAIF) in 2023.

## 537 **Software Package and Artefacts**

538 Our benchmark infrastructure, case studies and the executable of BINSEC/ASE have  
539 been made available through artifacts for reproducibility purposes on GitHub<sup>1</sup> and  
540 Zenodo<sup>2</sup>. The BINSEC/ASE source code will be open-sourced.

## 541 **1.6 Impact and Perspectives**

542 This work is a first step in designing efficient program analysis techniques able to take  
543 into account advanced attackers. The approach is generic enough to accommodate  
544 many common fault models, including the bit flips from RowHammer, test inversion,  
545 instruction skip or arbitrary data modification; still, generic instruction modifications  
546 are currently out of reach. Also, while we investigate the bug-finding side of the problem  
547 (under-approximation), the verification side (over-approximation) is interesting as well.  
548 These are exciting directions for future research.

## 549 **1.7 Manuscript Outline**

550 The remainder of the manuscript is organized as follows:

- 551 – In Chapter 2, we describe the context of systems with security features and  
552 motivate the need for dedicated program analyses. We list a number of attack  
553 vectors and related capabilities available to an advanced attacker;
- 554 – In Chapter 3, we propose an overview of program analysis techniques, with a focus  
555 on symbolic execution and binary analysis. We also explore existing techniques  
556 to automatically reason about an attacker able to inject faults into a program;
- 557 – In Chapter 4, we detail our formalization of an advanced attacker and the concept  
558 of adversarial reachability, as well as an algorithm, adversarial symbolic execution,  
559 to verify adversarial reachability in practice;
- 560 – In Chapter 5, we present our implementation of adversarial symbolic execution  
561 for binary-level analysis inside the BINSEC tool that we call BINSEC/ASE. We  
562 propose a user and a developer guide for BINSEC/ASE;
- 563 – In Chapter 6, we evaluate BINSEC/ASE performance on standard SWIFI bench-  
564 marks and propose few security scenarios to showcase the interest and feasibility  
565 of our technique;
- 566 – Finally, in Chapter 7 we conclude this thesis and propose future research direc-  
567 tions.

---

<sup>1</sup>[https://github.com/binsec/esop2023\\_artefact](https://github.com/binsec/esop2023_artefact)

<sup>2</sup><https://zenodo.org/record/7507112>



569 Context and Motivation

570 **Contents**

---

571	<b>2.1 Security of Information Technology Systems . . . . .</b>	<b>8</b>
572		
573	2.1.1 Definition . . . . .	8
574	2.1.2 Security Properties . . . . .	8
575	2.1.3 Wide-spread Use . . . . .	9
576	2.1.4 Attack Surface and Attackers . . . . .	11
577	2.1.4.1 A Wide Attack Surface . . . . .	11
578	2.1.4.2 Attack Surface Through Time . . . . .	12
579	2.1.4.3 Protections Adding Attack Surface . . . . .	12
580	2.1.4.4 Malicious Attackers . . . . .	13
581	2.1.5 Ensure Security . . . . .	13
582	2.1.5.1 Standard System Protections . . . . .	13
583	2.1.5.2 Product Validation . . . . .	13
584	2.1.5.3 Security Evaluation . . . . .	14
585	2.1.5.4 Certification Process . . . . .	14
586	2.1.5.5 Our Scope . . . . .	15
587	2.1.6 Example of Attack on a Security System . . . . .	15
588	2.1.7 Conclusion . . . . .	17
589	<b>2.2 Capabilities of a Powerful Attacker . . . . .</b>	<b>17</b>
590	2.2.1 Side Channel Attacks . . . . .	17
591	2.2.2 Software Attacks . . . . .	18
592	2.2.3 Hardware Fault Injection Attacks . . . . .	19
593	2.2.4 Software-Implemented Hardware Attacks . . . . .	21
594	2.2.5 Micro-Architectural Attacks . . . . .	22
595	2.2.6 Man-At-The-End (MATE) Attacks . . . . .	23
596	2.2.7 Summary . . . . .	24
597	<b>2.3 Conclusion . . . . .</b>	<b>24</b>

---

601 This chapter aims to motivate the work of this thesis, that is, designing a technique  
602 to automatically and efficiently reason about the impact of an advanced attacker on a  
603 program's security properties. First, in Section 2.1, we will have a glance at what are  
604 security-oriented IT systems, their usages and their properties. They are prevalent in  
605 modern life and their attack can have important consequences. Then, in Section 2.2,  
606 we will go over different kinds of attack techniques an advanced attacker can leverage  
607 against an IT system and what capabilities it gives them over a program.

## 608 2.1 Security of Information Technology Systems

609 In this section, we present IT (Information Technology) systems and their security-  
610 oriented features, we stress their prevalence. IT systems are targets for advanced at-  
611 tackers and can have a wide attack surface. We introduce techniques aimed at reducing  
612 this attack surface, going toward better system safety and security.

### 613 2.1.1 Definition

614 We call security-oriented IT systems any kind of IT system providing a security feature  
615 whose malfunction has bad consequences, knowledge, human, material or monetary-  
616 wise for instance.

617 An IT system is at least composed of a hardware part, the chips themselves, and the  
618 software part, the program running on the hardware. It is often combined with input  
619 and output systems such as cabled connections, sensors, buttons, antennas, displays  
620 or engines, for interactions with the world. Such a system can be on its own as an  
621 embedded system like an IoT (Internet of Things) device, or part of a larger IT system  
622 like a cloud server.

623 Many of today's IT systems integrate some form of cryptography to protect the  
624 confidentiality of communications, which is one of the most famous security features.  
625 This gives an idea of the prevalence of systems with security properties.

### 626 2.1.2 Security Properties

627 In the security field, we typically identify three main security properties a system can  
628 have. Knowing they are not always compatible, the strength with which each should  
629 be enforced depends on the system application.

630 – **Integrity:** the system is not to be modified except by legitimate means. In  
631 particular, the hardware shouldn't suffer damage, the code shouldn't be altered  
632 and the data shouldn't be corrupted. For instance, data is encoded with extra  
633 information to detect errors in network transmissions. An example of a malicious  
634 integrity breach is the defacement of a website, where an attacker changes the  
635 content of a web page against the owner's wishes.

636 – **Confidentiality:** protected information is to be revealed only to legitimate users.  
637 An IT system contains a lot of information, from code to data or even hardware  
638 design. Some of that information is labeled as 'public' which means it can be  
639 exposed without endangering the confidentiality of the system. Other parts of

640 the information contained are labeled ‘private’ and should remain secret. They  
 641 can be intellectual property like algorithms and hardware design, or secret data  
 642 such as keys used for encryption or authentication, or personal data. A man-  
 643 at-the-middle attack consists in intercepting, decoding, observing, re-encoding  
 644 and sending on its way an information paquet on a network, threatening the  
 645 confidentiality of those exchanges.

646 – **Availability:** some systems shouldn’t stop working, even in unforeseen conditions  
 647 like internal errors or when subjected to attack attempts. Availability is mea-  
 648 sured as a minimum percentage of a system’s uptime. Systems integrate recovery  
 649 mechanisms to restore a working state, fail soft modes or be able to restart the  
 650 system so it continues its purpose. Denial of service (DoS) attacks target the  
 651 availability of a system by bombarding it with requests until it cannot handle the  
 652 number and stops answering requests.

653 Other types of properties are sought for in security devices. We list some in the  
 654 following.

655 – **Authentication:** the use of an IT system’s features may be restricted to autho-  
 656 rized users only, possibly according to their roles. Each user should register their  
 657 identity and only perform actions they are allowed to on the system. An unau-  
 658 thorized user is not to forge an authorized user’s identity. A brute-force attack  
 659 can be used to try common passwords until one is accepted and the attacker is  
 660 allowed entry on a system.

661 – **Tracability:** each action on a system should be logged in order to be audited.  
 662 What type of action, on what object and performed by whom should be saved.  
 663 This is particularly important to detect traces of attackers on a large computer  
 664 system, determine what has been compromised and where was the exploited  
 665 breach.

666 Those are general, system-level security properties. They can be refined into a wide  
 667 range of other properties at different abstraction levels. For instance, one underlying  
 668 property of integrity is memory safety. It consists in ensuring memory accesses only  
 669 affect legitimate objects and does not corrupt or leak other objects.

### 670 2.1.3 Wide-spread Use

671 Many areas of our life use IT systems containing security-oriented features. We provide  
 672 here a non-exhaustive list of fields and attack examples to stress the need for high  
 673 security requirements.

674 **Military.** The first domain to come to mind may be defense, where IT systems<sup>1,2</sup> have  
 675 life-and-death consequences on humans. Moreover, the confidentiality and integrity of  
 676 information, or commands, from the enemy are necessary. There is also a need for  
 677 preserving intellectual property as a military advantage.

678 However, a wide variety of fields, in our everyday life, also takes advantage of  
 679 systems with security-oriented features.

<sup>1</sup><https://www.wired.com/story/dire-possibility-cyberattacks-weapons-systems/>

<sup>2</sup><https://www.brookings.edu/techstream/hacked-drones-and-busted-logistics-are-the-cyber-future-of-warfare/>

680 **Aeronautics.** Airplanes contain many cyber-physical devices to help pilots, anti-  
681 collision systems or pressure sensors for instance. They require communication with  
682 the airports and control towers. An airplane malfunction can put at risk the lives  
683 of hundreds of people. In 2015, a denial-of-service (DoS) attack<sup>3</sup> on the Polish airline  
684 company LOT resulted in the cancellation of 10 flights, and 15 more were delayed. The  
685 company systems were unable to send flight plans and other information necessary for  
686 take off such as weather forecasts. This attack left 1.400 people stranded at Warsaw  
687 Chopin airport.

688 **Aerospace.** The aerospace field is rich with security-oriented IT systems. They play  
689 a big part in communication systems, GPS, weather forecasts and cartography, for  
690 civilian and military usage. In addition, humans are sent to space, which is a hostile  
691 environment they need to be appropriately shielded from. In 1999<sup>4</sup>, a teenager found  
692 a flaw in a system of the American Defense Threat Reduction Agency (DTRA) and  
693 gained access to messages, including authentication information belonging to NASA.  
694 He used it to perform a privilege escalation attack, ultimately downloading the source  
695 code of the software running on board the International Space Station (ISS), including  
696 temperature and humidity control within living quarters.

697 **Banking.** Maybe the first example that comes to mind with security-oriented IT sys-  
698 tems present in our pockets is the credit card. It contains private encryption keys to  
699 authenticate to a bank and allow payments. Credit card terminals are equally vul-  
700 nerable. In 2015 and 2016, the SWIFT banking network suffered an attack<sup>5</sup> where  
701 attackers, using a malware, gained access to legitimate SWIFT credentials they used  
702 to send transfer requests to other banks, who trusted the messages and transferred  
703 funds, leading to the theft of more than 100 millions of dollars. Maybe less known  
704 are crypto wallets, devices providing encryption keys to access a digital wallet of cryp-  
705 tourrencies and allow trade. The cryptocurrencies themselves rely on blockchains, and  
706 smart contracts where a bug has devastating consequences and is hard to fix. This is  
707 a current topic of research [WZ18].

708 **Energy.** The power grid is monitored and regulated by a network of devices. It can be  
709 attacked to cause blackouts. The first acknowledge successful attack on a power grid  
710 targeted Ukraine in 2015<sup>6</sup>. Corporate networks were infiltrated with the BlackEnergy  
711 malware using spear-phishing emails. The attack then progressed on to SCADA devices  
712 switching them off remotely, and disabling IT infrastructure and emergency power  
713 systems, causing a wide power blackout affecting more than 200.000 people. In addition,  
714 a DoS attack on the call center kept consumers in the dark.

715 The water management system is equally vulnerable. In 2021, Florida was the  
716 stage of an attack on its water system<sup>7</sup>. An attacker took advantage of a remote access  
717 software on a computer not up-to-date with an insufficient firewall rules to try and  
718 poison the water supply. Fortunately, it was discovered before any harm could be done  
719 to the population.

720 **Entertainment.** While the entertainment industry doesn't immediately come to mind  
721 when one thinks about security-oriented IT systems, their systems still record per-

---

<sup>3</sup><https://www.cnbc.com/2015/06/22/hack-attack-leaves-1400-passengers-of-polish-airline-lot-grounded.html>

<sup>4</sup><https://cybernews.com/editorial/how-a-florida-teenager-hacked-nasas-source-code/>

<sup>5</sup>[https://en.wikipedia.org/wiki/2015-2016\\_SWIFT\\_banking\\_hack](https://en.wikipedia.org/wiki/2015-2016_SWIFT_banking_hack)

<sup>6</sup>[https://en.wikipedia.org/wiki/2015\\_Ukraine\\_power\\_grid\\_hack](https://en.wikipedia.org/wiki/2015_Ukraine_power_grid_hack)

<sup>7</sup><https://edition.cnn.com/2021/02/10/us/florida-water-poison-cyber/index.html>

722 sonal data like account and password (sometimes reused elsewhere), but also banking  
 723 information. In 2012, the PlayStation Network underwent an attack<sup>8</sup> compromising  
 724 the account of around 77millions of people, including children. Among that personal  
 725 information was payment data. To limit the leak, the PlayStation network was down  
 726 for 23 days preventing users from accessing the gaming service.

727 **Health.** Another card in our wallet (in France at least) is the social security smart  
 728 card. More and more health devices, remotely controllable to some extent, are used  
 729 in medicine to help patients, like pacemakers or insulin pumps. They can be hacked  
 730 to hurt the people wearing them. In general, hospitals rely on IT systems but are  
 731 notorious for their lack of cybersecurity investments. For instance, in 2017 in the  
 732 United Kingdom<sup>9</sup>, the NHS was infected by the ransomware WannaCry, blocking access  
 733 to computers and medical equipment, appointments and operations were canceled and  
 734 patient information was lost.

735 **Telecommunications.** The field of telecommunication includes the Internet and the  
 736 many routers and relais paving it. They are outside, in the world, and susceptible to  
 737 disrupting the network or stealing sensitive data. Mobile phone telecommunications  
 738 can also be targeted, as well as television, radio frequencies, etc. For instance, in  
 739 1995, a radio contest in which the 102<sup>nd</sup> caller would win a Porsche was hijacked<sup>10</sup>.  
 740 Telephone lines were blocked to make sure only the attacker call could succeed. In  
 741 2015, the French television group, TV5Monde fell victim to an attack<sup>11</sup> that targeted  
 742 the broadcast system. All their channels when black which put the company at risk  
 743 of losing distribution contracts. Their social media, website and internal systems also  
 744 went down.

745 **Transport.** The rise of autonomous cars [KKJ<sup>+</sup>21] gives growing responsibilities to  
 746 manufacturers in traffic accidents and overall road safety. A famous example is a  
 747 recognition software being derailed by a speed limit sign interpreted as a stop sign.  
 748 Many other transport systems rely heavily on IT systems, such as boats, trains, traffic  
 749 control systems, automated subways, drones and UAVs. In 2017, boat navigation was  
 750 disturbed<sup>12</sup> by an attack on the GPS system, placing some boats inland.

## 751 2.1.4 Attack Surface and Attackers

752 Security-oriented IT systems are a prime target for attackers, who would exploit vul-  
 753 nerabilities from any component and at any stage of the product life-cycle.

### 754 2.1.4.1 A Wide Attack Surface

755 IT systems are complex systems, embedding many different layers which can all have  
 756 vulnerabilities: the hardware, the micro-architecture, software-controlled hardware  
 757 mechanisms, kernels, operating systems, network, software, algorithms or protocols. IT  
 758 systems can reside in physically secure server rooms, providing their services through

<sup>8</sup>[https://en.wikipedia.org/wiki/2011\\_PlayStation\\_Network\\_outage](https://en.wikipedia.org/wiki/2011_PlayStation_Network_outage)

<sup>9</sup><https://www.nao.org.uk/reports/investigation-wannacry-cyber-attack-and-the-nhs/>

<sup>10</sup><https://www.industrialcybersecuritypulse.com/threats-vulnerabilities/throwback-attack-kevin-poulsen-wins-a-porsche-and-hacks-the-u-s-government/>

<sup>11</sup>[https://en.wikipedia.org/wiki/TV5Monde#April\\_2015\\_Cyberattack\\_and\\_resulting\\_disruption](https://en.wikipedia.org/wiki/TV5Monde#April_2015_Cyberattack_and_resulting_disruption)

<sup>12</sup><https://www.newscientist.com/article/2143499-ships-fooled-in-gps-spoofing-attack-suggest-russian-cyberweapon/>

759 a cabled connection, usually over the Internet. As getting proximity access to such de-  
760 vices is difficult, attackers favor software-based attacks. Security-oriented IT systems  
761 can also be embedded devices, that is to say, small systems scattered in the world where  
762 they are needed for their features. Since embedded devices are in the open, an attacker  
763 can have direct contact with the hardware. It enables them to perform, in addition  
764 to the wide spectrum of attacks possible on physically secure systems, hardware fault  
765 injection attacks. We detail different types of attack vectors and what capabilities they  
766 provide an attacker in section 2.2.

767 Some but not all attack vectors rely on existing *vulnerabilities, or bugs*, in an IT  
768 system. A vulnerability is a weakness in a system that can be exploited. It can be a  
769 design flaw or a development error for instance.

### 770 2.1.4.2 Attack Surface Through Time

771 In addition to being complex systems, IT systems also have a rich life cycle, each  
772 step possibly introducing vulnerabilities. The Common Criteria [CC2] detail different  
773 moments in the life-cycle of a product where vulnerabilities may arise through failures  
774 in IT:

- 775 – *requirements*: an IT product may possess all the functions and features required  
776 of it and still contain vulnerabilities that render it unsuitable or ineffective with  
777 respect to security;
- 778 – *design*: an IT product has been poorly designed. Building a secure product,  
779 system, or application requires not only the implementation of functional re-  
780 quirements but also an architecture that allows for the effective enforcement of  
781 specific security properties the product, system, or application is supposed to  
782 enforce. The ability to withstand attacks the product, system, or application  
783 may be face in its intended operational environment is highly dependent on an  
784 architecture that prohibits those attacks or, if they cannot be prohibited, allows  
785 for detection of such attacks and/or limitation of the damage such an attack can  
786 cause;
- 787 – *development*: an IT product does not meet its specifications and/or vulnerabili-  
788 ties have been introduced as a result of poor development standards or incorrect  
789 design choices;
- 790 – *delivery, installation and configuration*: an IT product has vulnerabilities intro-  
791 duced during the delivery, installation and configuration of the product;
- 792 – *operation*: an IT product has been constructed correctly to a correct specification,  
793 but vulnerabilities have been introduced as a result of inadequate controls upon  
794 the operation.
- 795 – *maintenance*: an IT product is maintained in such a way that new vulnerabilities  
796 are introduced.

### 797 2.1.4.3 Protections Adding Attack Surface

798 Interestingly, removing vulnerabilities in a device by adding protections can lead to  
799 adding components, micro-architectural or software, in the systems, which can them-  
800 selves suffer from vulnerabilities or induce new ones in other parts. Hence, adding  
801 countermeasures also increases the attack surface of a system.

802 This stresses the high potential for a wide range of vulnerabilities to exist in a  
803 security product.

#### 804 **2.1.4.4 Malicious Attackers**

805 There is a wide diversity of malicious attackers with various motivations from monetary  
806 to politics. The most serious threats to security-oriented systems are powerful entities  
807 such as mafias and state agencies. They have enormous resources in terms of time,  
808 money to buy equipment and in terms of expertise. It is reasonable to assume such  
809 attackers will try every attack vector at their disposal and target any stage of the  
810 product life-cycle. They may even use some new vulnerabilities or attack techniques  
811 not yet known to the public. Those are called zero-day vulnerabilities [BD12]. They  
812 are hard to detect and impossible to protect against before they are publicly disclosed.  
813 Zero-day vulnerabilities can exist for years and are traded between malicious parties.

814 Malicious attackers pose a real threat to security-oriented IT systems. Techniques  
815 and processes have been developed to harden system security and add considerable  
816 difficulties for malicious attackers.

#### 817 **2.1.5 Ensure Security**

818 To harden security-oriented IT systems against attackers, different *complementary* ap-  
819 proaches have been developed.

##### 820 **2.1.5.1 Standard System Protections**

821 The first step to producing a secure product is to produce a ‘good’ product. Developers  
822 exchange good practices for efficient but also readable and self-explanatory code and  
823 specifications. Companies put in place development processes like code review to reduce  
824 the risk of bugs.

825 The design choice can have a big impact on reducing the risk and number of vulner-  
826 abilities. For instance, the choice of programming language can prevent some classes of  
827 bugs. A type-safe language contributes to program correctness and prevents type er-  
828 rors like using operations with incorrect data types. A memory-safe language enforces  
829 memory isolation.

830 When the chosen programming language is compiled, a generic tool is used, GCC<sup>13</sup>  
831 or Clang<sup>14</sup> to compile C for instance. They integrate by default a number of security  
832 protections, like stack canaries, and can be configured more precisely.

833 **Limits.** Those standard security practices permit the removal of some vulnerabilities,  
834 but are mostly human-based, hence fallible, or take advantage of default configurations  
835 that are rather weak and not tailored to a particular system.

##### 836 **2.1.5.2 Product Validation**

837 To complement standard development processes, various test phases have been designed  
838 to improve functional correctness and detect bugs early. Unit tests are written or  
839 generated, for each functional unit, verifying correct output for a limited number of  
840 inputs. This help make sure of the normal behavior of the system, i.e. functional

---

<sup>13</sup><https://gcc.gnu.org/>

<sup>14</sup><https://clang.llvm.org/>

841 correctness. Then, integration tests check that each component interacts with others  
842 as expected, as many imprecisions or bugs lay at system interfaces, which can then be  
843 exploited by attackers. Regression testing ensures new developments haven't impacted  
844 the nominal behaviors of the other parts of the system.

845 **Limits.** This validation procedure allows to remove most vulnerabilities in a system.  
846 However, one of the main difficulties with tests is the coverage of all functional behav-  
847 iors. If a set of paths is not explored by a test, its correctness isn't checked. Exhaust-  
848 tiveness is a hard problem and can result in very heavy test suites, too heavy for some  
849 development processes. Incomplete coverage often happens for corner cases, that an  
850 attacker can exploit. For instance, what happens if an element of a data structure is  
851 empty? In addition, test cases are often selected with functional testing in mind, that  
852 is to say, what the system should do, and not security, i.e. what the system shouldn't  
853 do. In particular, incorrect inputs are not always exhaustively tested.

### 854 2.1.5.3 Security Evaluation

855 The next step in strengthening an IT system is to perform *penetration testing*. It is a  
856 security-oriented testing phase to *evaluate the security by design in practice*. Penetra-  
857 tion testing consists, for security practitioners, in attacking the system as a malicious  
858 attacker would to identify the remaining vulnerabilities and available attack vectors.  
859 Different entities perform penetration testing on IT systems.

860 **Manufacturer.** It is the entity designing new IT systems they wish to market. They  
861 may externalize the hardware manufacturing and even the software development. How-  
862 ever, they are responsible for the end product security. Manufacturers test their prod-  
863 ucts and look for vulnerabilities to remove them as early as possible in the production  
864 process as making patches becomes more and more expensive. This includes trying to  
865 attack their product, to ensure a certain level of robustness to their clients.

866 **Evaluators.** Penetration testing can be externalized to evaluators. They can work  
867 for manufacturers or certification authorities (see Section 2.1.5.4). Evaluators receive  
868 an IT system, perform its vulnerability assessment including penetration testing, and  
869 produce a report.

870 **Security Researcher.** Researchers study the security of IT systems to improve the state  
871 of knowledge, discover new types of vulnerabilities, develop new tools to find them and  
872 devise new and more powerful countermeasures. Their work empowers the security  
873 community, helping improve their assessments. It is likely similar to the research made  
874 by malicious attackers, in particular, to discover zero-day vulnerabilities in systems.

### 875 2.1.5.4 Certification Process

876 Each company has its own development process, making it hard to compare systems  
877 quality and security. To bridge that gap, independent authorities receive product appli-  
878 cations and attribute to them a standardized label. Those are *certification authorities*,  
879 delivering *certifications*.

880 **Certification Authorities.** At an international level, countries have gathered around  
881 Common Criteria. It takes the form of a procedure to obtain the Common Criteria  
882 certificate, with different security levels. Other certifications exist, at the national level  
883 for instance, in France, the CSPN can be delivered by the French National Cyberse-  
884 curity Agency (Agence Nationale de la Sécurité de Systèmes d'Information, or ANSSI

885 for short).

886 **Certification.** The application consists of extensive documentation of the product,  
 887 development and test processes followed by the company. It also integrates a pen-  
 888 etration testing report from an external evaluation center. For instance, among the  
 889 documentation required for the Common Criteria certification, the class AVA relates  
 890 to vulnerability assessment, where the evaluator is expected to report on penetra-  
 891 tion testing activities to ensure the non-exploitability of the product’s vulnerabilities  
 892 [CC2]. The certification authority reviews the application and evaluates the product  
 893 made available to them to decide whether or not to give the certification to the new  
 894 product.

### 895 2.1.5.5 Our Scope

896 The work presented in this thesis focuses on *program-level evaluation for security*.  
 897 It aims to identify vulnerabilities in an IT system that an advanced attacker could  
 898 exploit. It could assist penetration testers identify sensitive code areas and possible  
 899 attack paths. We propose for instance the following use cases:

- 900 – the evaluation of the strength of different cryptographic implementations;
- 901 – the evaluation of security features of security-oriented IT systems such as a secure  
 902 bootloader;
- 903 – the vulnerability evaluation of a program with various countermeasures.

### 904 2.1.6 Example of Attack on a Security System

905 As a motivating example, we consider a pin verification program. It can be used  
 906 for different applications, for instance, for credit cards, building access control, secure  
 907 storage devices, or authentication of any kind. Here, an attacker may want to leak the  
 908 secret pin, or they can confuse the system into believing the attacker has the secret pin  
 909 even if not, in order to continue and use authenticated features. More formally, the  
 910 attacker’s goal is to be authenticated despite not knowing the correct pin.

911 The first step for an attacker would be to retrieve the program, either from public  
 912 sources or by leaking it. The motivating example is illustrated in Figure 2.1, it is an  
 913 unprotected version of a VerifyPIN program [DPP+16], from the domain of hardware  
 914 fault injection and smart cards. The user pin digits, stored in *g\_userPin*, are checked  
 915 against the reference digits stored in *g\_cardPin*, using the *byteArrayCompare* function.  
 916 The attacker seeks to be authenticated (satisfy the assert l.29) without knowing, i.e.  
 917 entering, the right digits.

918 We propose hereafter different attack scenarios for the attacker to exploit the Ver-  
 919 ifyPIN program.

920 **Attack Scenario 1.** Once the attacker has access to the program code, they would ana-  
 921 lyze it for known *Common Vulnerabilities and Exposures* (CVEs) and classic *exploitable*  
 922 *vulnerabilities*, checking for a buffer-overflow vulnerability for instance.

923 **Attack Scenario 2.** During their analysis of the program, the attacker can notice this  
 924 code is not constant-time. That is to say, a program branch depends on a secret value  
 925 (l.10) leading to different execution times and information leakage about the secret.  
 926 Here the program will take longer to execute, i.e. execute one more loop iteration,

```

1 #define PIN_SIZE 4
2
3 bool g_authenticated = 0;
4 int g_userPin [PIN_SIZE];
5 int g_cardPin [PIN_SIZE];
6
7 bool byteArrayCompare(int* a1, int* a2, int size) {
8     int i;
9     for(i = 0; i < size; i++) {
10         if(a1[i] != a2[i]) {
11             return 0;
12         }
13     }
14     return 1;
15 }
16
17 void verifyPIN() {
18     g_authenticated = 0;
19     if(byteArrayCompare(g_userPin, g_cardPin,
20         PIN_SIZE) == 1)
21     {
22         g_authenticated = 1; // Authentication();
23     }
24     return;
25 }
26
27 void main() {
28     verifyPIN();
29     assert(g_authenticated == 1);
30 }

```

Figure 2.1: Motivating example, inspired by VerifyPIN [DPP<sup>+</sup>16]

927 for each correct digit. By performing a timing side-channel attack, the attacker can  
928 leak the secret pin by learning one digit at a time, testing all 10 possibilities. This  
929 would require at most  $10 \times PIN\_SIZE$  attempts which is much lower than brute force  
930 ( $10^{PIN\_SIZE}$  attempts). Then the attacker can use this information to get legitimately  
931 authenticated by the program.

932 **Attack Scenario 3.** If they come up empty-handed, attackers could try and determine  
933 sensitive code areas, that is to say, code regions where inducing a fault in the execution  
934 is likely to result in a different program behavior, closer to what the attacker aims for.  
935 More details on means to inject faults into a program are presented in Section 2.2. If the  
936 injection of the fault requires precise timing, for *hardware fault injections* or software-  
937 induced hardware fault injections for instance, the attacker would need to observe the  
938 running program, through some form of side-channel leakages like power consumption  
939 or electromagnetic emission, to link processor activity with program execution in order  
940 to synchronize the injection with the execution of the vulnerable area. The attacker  
941 may need multiple attempts, changing the injection vector or the vulnerable code region

942 before performing a successful injection.

943 **Attack Scenario 3A.** The attacker could use fault injection to change the *control*  
944 *flow* of the program, by inverting the test l.19 which checks the return value of the  
945 *byteArrayCompare* function, so that when it returns 0 (pin incorrect), the execution  
946 would in fact execute l.22 and authenticate the user.

947 **Attack Scenario 3B.** The attacker could target *data in memory*. In this example,  
948 injecting a non-controlled fault into memory to corrupt *g\_cardPin* would not help the  
949 attacker. If they can inject a precise fault, resetting the whole memory array associated  
950 with *g\_cardPin*, then entering a "0 0 0 0" pin would authenticate them.

951 **Attack Scenario 3C.** The attacker can target the *data flow* of the program. By cor-  
952 rupting the initialization of *g\_authenticated* l.18 with a fault injection setting it to  
953 1, the program will detect the pin entered is wrong and forward the original value of  
954 *g\_authenticated*. Hence the attacker will be authenticated with an incorrect pin.

## 955 2.1.7 Conclusion

956 In this section, we had a look at various fields relying heavily on IT systems with  
957 security features and what efforts were put in place to harden their security against  
958 attackers, for instance through a Common Criteria certification. The VerifyPIN moti-  
959 vating example illustrated our focus on program-level security evaluation.

960 Cyber-security is a cat-and-mouse game, where attackers and security researchers  
961 develop new attack techniques while manufacturers and other security researchers try  
962 to detect them and design countermeasures. Hence protections are imperfect and still  
963 leave ways for an advanced attacker to build successful attacks. In the next section, we  
964 will delve into what attack techniques are available to an advanced attacker and what  
965 attack capabilities they provide the attacker to reach their goal.

## 966 2.2 Capabilities of a Powerful Attacker

967 Security research has delved into many aspects of system security, leveraging a growing  
968 variety of attack techniques and attack surfaces that a powerful attacker can take  
969 advantage of.

970 We aim in this section to touch on what impact an advanced attacker can have on  
971 a program. We consider many attack vectors from software to micro-architecture and  
972 hardware. We list the attack capability given to an attacker by each kind of attack.  
973 It is by knowing against what to defend a program that it is possible to evaluate its  
974 security and strengthen it.

### 975 2.2.1 Side Channel Attacks

976 A system running a program leaks a lot of information about its internal state. It  
977 can be through external observables like power consumption [KJJ99], electromagnetic  
978 emissions [AARR03, DOL<sup>+</sup>10], execution time [Koc96] or light emission [SNK<sup>+</sup>12], and  
979 more recently through sound emissions [GST17]. An attacker able to run a program  
980 on the target processor can also observe internal leakages like cache usage [Ber05,  
981 LM18] or branch prediction [AKS06]. An attacker observing those leaks deduces private  
982 information, like interesting addresses or secret keys. As the attacker does not modify

983 the system's behavior, it is a *passive attack*. We will not detail side channel attacks  
984 further and focus on active attacks in this thesis.

### 985 2.2.2 Software Attacks

986 Many tools have been developed to ensure system's security, mostly taking into account  
987 an attacker able to craft malicious inputs to give to legitimate input sources of a  
988 program. Software attacks exploit existing software vulnerabilities. A number of them  
989 are recorded in the CVE database<sup>15</sup>. They can be performed remotely. We list below  
990 the most common ones and the attack primitive they give to an attacker.

991 **Buffer Overflows.** It is probably the most famous one. This vulnerability appears  
992 when the program allows to write in memory outside the bound of a data structure,  
993 arrays in particular. This is an undefined behavior in the C language that is usually  
994 resolved by allowing the write outside of bounds. This vulnerability is often due to a  
995 code pattern where a write operation is performed inside a loop and there end up being  
996 too many iterations of the loop. It can also appear when copying a data structure into  
997 a smaller one. Figure 2.2, from OWASP<sup>16</sup>, illustrates a code relying on external data  
998 to control its behavior. The code uses the *gets()* function to read an arbitrary amount  
999 of data into a stack buffer. Because there is no way to limit the amount of data read  
1000 by this function, the safety of the code depends on the user to always enter fewer than  
*BUFSIZE* characters.

```
1 char buf [BUFSIZE];  
2 gets (buf);
```

Figure 2.2: Example of a buffer overflow vulnerability.

1001  
1002 A buffer overflow attack [CWP+00] gives the attacker write capabilities. The speci-  
1003 ficities of the code will constrain what can be overwritten and how much of it. A  
1004 common usage is to overwrite the return pointer at the end of a function and start a  
1005 return-oriented type of attack.

1006 **Use-after-free.** This type of attack [CGMN12] consists in dereferencing a pointer after  
1007 it has been released by the code. An attacker can use that pointer to perform illegal  
1008 memory accesses, on data or function pointers, retrieving old values or new ones al-  
1009 located after the space had been freed. A use-after-free attack enables an attacker to  
1010 control values from memory loads and jump to attacker control code.

1011 **Integer overflow.** Integer overflow [DLRA15] consists in providing inputs such that  
1012 an arithmetical computation will overflow –or underflow, that is to say, the theoretical  
1013 result cannot be written with the chosen integer representation. This is an undefined  
1014 behavior in the C language. This vulnerability can be used to corrupt data and even  
1015 control flow when targeting loop iterations.

1016 **Conclusion.** Software vulnerabilities provide the attacker with a number of interesting  
1017 attack primitives, mostly based on data corruption with extended effects. Table 2.1  
1018 summarizes attacker capabilities. Here we only consider the impact of exploiting such a  
1019 vulnerability and presume attack prerequisites have been met. Interestingly, works on

---

<sup>15</sup><https://cve.mitre.org/>

<sup>16</sup>[https://owasp.org/www-community/vulnerabilities/Buffer\\_Overflow](https://owasp.org/www-community/vulnerabilities/Buffer_Overflow)

1020 Control Flow Integrity (CFI) mechanisms [ABEL09, BCN<sup>+</sup>17] also define hypothetical  
 1021 attackers by their capabilities, such as “write anything anywhere” or “write anything  
 1022 somewhere”, in order to prove that their countermeasure is indeed able to thwart such  
 1023 an opponent.

Table 2.1: Attacker capabilities given by software attacks

Type of attack	Attacker capability
Buffer overflow	Overwrite data in memory
Use-after-free	Corrupt data loads from memory
Integer overflow	Corrupt input-related data

### 1024 2.2.3 Hardware Fault Injection Attacks

1025 When an attacker has direct access to a device, they can put it in extreme environmental  
 1026 conditions, past the operating limits of the components, to induce hardware faults in  
 1027 them. The faults occur at the circuit level, either setting a signal to zero (bit reset), to  
 1028 one (bit set) or flipping the value (bit-flip). The consequences of that circuit-level fault  
 1029 will manifest at the micro-architectural level depending on the location of the faults,  
 1030 for instance at the fetch, decode, execute or store stage of the microprocessor pipeline,  
 1031 in registers, cache or memory, etc. Finally, faults propagate to the software layer,  
 1032 changing the behavior of the executed program. Faults can be *transient*, with effects  
 1033 lasting only until they are naturally overwritten, or *permanent*, when they cause lasting  
 1034 damage to the circuit. We will only consider transient fault in this thesis. Security  
 1035 practitioners have identified various injection means [BECN<sup>+</sup>06, KSV13, BBKN12].  
 1036 We list the most common ones in the following, with what types of capabilities they  
 1037 are known to provide an attacker.

1038 **Laser Beam.** A photoelectric effect happens when a material is hit by electromagnetic  
 1039 radiation such as a laser beam, the material releases electrons creating electron holes  
 1040 that induce currents and cause faults [DDCS<sup>+</sup>14]. This is the most accurate hard-  
 1041 ware fault injection mean, with a timing precision of a few picoseconds [LBC<sup>+</sup>15] at  
 1042 a location precise at a few micrometers [BMPPM13]. A laser beam injection can affect  
 1043 each stage of the pipeline (instruction fetch, decode, execute, write) [VTM<sup>+</sup>17]; for  
 1044 instance corrupting operands, corrupting instructions or writing wrong results back in  
 1045 a possibly different register. A laser beam can also inject a bit-flip in a register, a flag  
 1046 value [VTM<sup>+</sup>17] or in flash memory [CMD<sup>+</sup>19] independently of program execution.  
 1047 A concrete attack could be flipping a precise bit in the secure configuration register  
 1048 (SCR) to bypass a secure boot [VTM<sup>+</sup>17]. A laser beam can inject multiple faults in  
 1049 one attack:

- 1050 – at different physical locations [CGV<sup>+</sup>22], for instance, precisely faulting each  
 1051 check of a pin verification code [DRPR19];
- 1052 – affecting consecutive instructions [RNR<sup>+</sup>15]. Interestingly, if we consider an at-  
 1053 tacker capable of an infinite number of precise skips, they have a Turing com-  
 1054 plete computing power [PCHR20]. In practice, up to 300 instructions have been  
 1055 skipped [DRPR19].

1056 **Power glitch.** The critical path is the signal taking the longest to get from one flip-  
 1057 flop or memory unit, through some logic gates and to the next flip-flop or memory  
 1058 unit. Under-powering a system will slow down signals until some signal in the critical

1059 path cannot reach the next sampling unit in time. Those are called *timing faults*.  
1060 Power glitching allows for attacks with precise timing but no spatial control. The  
1061 power supply of a component must be accessible. A power glitch can skip instructions  
1062 or influence a branch decision, corrupt memory locations, or alter the result of an  
1063 instruction or its side effects [BFP19]. Power glitches can be used for privilege escalation  
1064 on a Linux system [TM17]. Multiple power glitching faults are possible [BFP19] and  
1065 can be leveraged to extract a protected firmware with more than a hundred faults to  
1066 force memory accesses.

1067 **Clock glitch.** The clock gives the pulse at which rate flip-flops and memory units  
1068 sample signals. When the frequency is increased passed the maximal frequency safe  
1069 for the critical path, timing faults occur. The location of a clock glitch is architecture-  
1070 dependent but precise in timing. The device must have an external clock that can be  
1071 changed to an attacker-controlled one. A clock glitch in the CPU pipeline can corrupt  
1072 instructions by affecting an operand fetch or skipping an instruction [YGS<sup>+</sup>16], with or  
1073 without replaying the previous one [CPHR21]. A glitch injection platform [CPHR21]  
1074 has skipped a hundred consecutive instructions and performed multiple bursts of skips.  
1075 One clock glitch can bypass unprotected authentication and multiple faults, a protected  
1076 one [CPHR21].

1077 **Electromagnetic Pulse (EMP).** The electromagnetic field of an EM probe generates  
1078 by induction eddy currents in the chip. This also creates timing faults in the circuit  
1079 [ZDT<sup>+</sup>14, OGSM15]. EMP attacks are temporally and spatially precise [MDP<sup>+</sup>20].  
1080 With a low-intensity electromagnetic disturbance, instructions can be corrupted to  
1081 other instructions, and at high-intensity, instruction's op-code are zeroed [BGV11].  
1082 EMP has been shown [MDH<sup>+</sup>13] to inject faults either on the instruction bus, for  
1083 instance targeting branch instructions, ALU instructions or load-store instructions,  
1084 and on the data bus independently of the execution, bringing each faulted bit closer to  
1085 the precharge bit value (architecture dependent). Multiple bit-flips or most-significant  
1086 half-word reset on register values, operand substitutions and instruction skips or replay  
1087 were also observed [PHB<sup>+</sup>19]. EMP can induce vulnerabilities such as a buffer overflow  
1088 and enables control flow hijacking, covert backdoor insertion and Return Oriented  
1089 Programming [BLLL18]. Multiple consecutive instructions can be corrupted [RNR<sup>+</sup>15].

1090 **Conclusion.** Hardware fault injection attacks have various effects, usually described  
1091 as modifications of instructions, or data corruption. We list the most common ones  
1092 in Table 2.2. In addition, *multiple faults* can be injected with a hardware attack.  
1093 Capabilities provided by hardware fault injection are the building blocks of exploitation:

- 1094 – breaking cryptography [BS97, Gir05, BSGD09, BBB<sup>+</sup>10, BBKN12, VKS11, DDRT12,  
1095 HS14],
- 1096 – leaking sensitive data [BDL01, DMM<sup>+</sup>13, BFP19],
- 1097 – circumventing authentication [QS02, DPP<sup>+</sup>16, DRPR19, CPHR21],
- 1098 – bypassing software protections [NHH<sup>+</sup>17, BLLL18], sometimes in parallel with a  
1099 software attack exploiting the created vulnerability as a *combined attack* [TS16,  
1100 NHH<sup>+</sup>17].
- 1101 – bypassing secure boot [TS16, VTM<sup>+</sup>17, BFP19],
- 1102 – performing privilege escalations [TS16, TM17];
- 1103 – disrupting a system whose output is studied by side-channel analysis [AVFM07].

Table 2.2: Attacker capabilities given by hardware fault injection attacks

Type of attack	Attacker capability
Hardware fault injection attack: laser beam, power glitch, clock glitch, temperature, EMP	Skip an instruction
	Skip and replay the previous instruction
	Change operator of an instruction
	Change source register of an instruction
	Corrupt load from memory
	Change destination register of an instruction
	Corrupt a store instruction's value
	Corrupt a memory cell
	Corrupt a register
Corrupt a flag	

## 1104 2.2.4 Software-Implemented Hardware Attacks

1105 We distinguish software-implemented hardware attacks from traditional hardware at-  
 1106 tacks due to the source of the perturbation. While hardware perturbations are caused  
 1107 by extreme environmental conditions (external) in hardware attacks, they are caused  
 1108 by software-controlled mechanisms (internal) in software-implemented hardware at-  
 1109 tacks. The goal is to push the hardware into unstable states using legitimate soft-  
 1110 ware commands. While hardware-controlled injections require proximity to the target,  
 1111 software-controlled injections can be done remotely. We list some examples of software-  
 1112 controlled mechanisms that can be misused to inject faults into a program and their  
 1113 associated fault models.

1114 **Rowhammer.** In 2014, Kim [KDK+14] first discovered that high-intensity memory  
 1115 access to the same location in memory can induce a bit-flip in the adjacent memory row,  
 1116 in DRAM (Dynamic Random Access Memory). The trend of circuit miniaturization  
 1117 leads to memory cells closer to one another and their charges being able to disturb each  
 1118 other. This was named the Rowhammer vulnerability. It was then exploited in 2015  
 1119 to escape a sandbox and to perform a privilege escalation [SD15] by corrupting the  
 1120 page table entry (PTE) to point to the attacker's one and control all physical memory,  
 1121 hence the entire system. Since then, a number of works have developed other exploits,  
 1122 on many different platforms [MK19]. It provides the attacker with a bit-flip capability  
 1123 at a precise location anywhere in flash memory [RGB+16].

1124 **Memory Delays.** The time a signal takes to travel between the memory and the CPU  
 1125 mostly depends on the length of the wire connecting them. However, that delay can  
 1126 also change with voltage and temperature. To accommodate for the delays changing  
 1127 from one board to another and due to working conditions, hardware timing control  
 1128 mechanisms have been implemented, called delay-lines. They are usually configured  
 1129 at boot but remain programmable. After being used for power side-channel attacks  
 1130 [GDTM21b], exploiting the link between voltage fluctuations and memory delays, this  
 1131 software-controlled mechanism was turned into a fault injection mean, named Fault-  
 1132 Line, by Gravelier *et al.* [GDTM21a]. By modifying the delay-line to a shorter period,  
 1133 an attacker can induce faults in memory transfers (loads and stores) of a victim appli-  
 1134 cation. This provides the attacker with the capability to induce a precise number of  
 1135 bit-flips either in memory loads or memory writes, or both. For instance, it allowed  
 1136 Gravelier *et al.* to perform a BellCore attack on OpenSSL signatures.

1137 **Dynamic Voltage and Frequency Scaling (DVFS).** Hardware manufacturers have de-  
 1138 veloped CPU with energy management controls, adapting frequency and voltage to  
 1139 increase performance in high-intensity computation phases and reduce consumption in  
 1140 low-intensity phases. Those configurations are stored in registers that can be controlled  
 1141 by kernel drivers, leading to potential misuse. Tang *et al.* devised the first exploit tech-  
 1142 nique, CLKSCREW, based on the DVFS vulnerability in 2017 [TSS17]. They show  
 1143 the microprocessor can be put in an unstable state, with the same consequences as  
 1144 under-powering or overclocking. The attack is imprecise in location and induces faults  
 1145 in flip-flops. DVFS has then been shown in the Plundervolt attack [MOG<sup>+</sup>20] to induce  
 1146 memory safety vulnerabilities and recover secret keys from cryptographic programs on  
 1147 an SGX enclave.

1148 **Conclusion.** As systems become more complex, they often require more modularity  
 1149 to adapt to different boards and functioning conditions. Hence a number of software-  
 1150 implemented hardware mechanisms to keep signal synchronization and improve energy  
 1151 efficiency. However, they often remain programmable hence vectors of attacks.

1152 Software-implemented hardware attacks can yield the same types of faults than  
 1153 hardware fault injection, summarized in Table 2.3, while lifting the direct access re-  
 1154 striction. They can be used for the same classes of attacks, from memory corruption  
 1155 to privilege escalation and cryptography attacks.

Table 2.3: Attacker capabilities given by software-implemented fault injection attacks

Type of attack	Attacker capability
Rowhammer	flip a bit in memory
Faultline	bit-flips in memory transfers
DVFS	flip a bit somewhere

## 1156 2.2.5 Micro-Architectural Attacks

1157 Manufacturers have design a wide variety of acceleration mechanisms to improve the  
 1158 performance of their CPUs in parallel with the miniaturisation trend. Those mecha-  
 1159 nisms can be misused to yield attacks. We present some of them in this section.

1160 **Spectre Attack.** To improve modern processors’ performance and, in particular, to  
 1161 avoid having to wait for a conditional branching resolution, branch prediction and  
 1162 speculative execution have been implemented. Branch prediction consists, for the pro-  
 1163 cessor, in making a guess as to which branch will be taken, based on previous decisions.  
 1164 For instance, if the last hundred of time, the branch true was taken for a particular  
 1165 conditional jump, the branch predictor can guess the true branch will be taken again  
 1166 this time. The register state is saved and the processor enters a speculative execution  
 1167 mode, where it continues executing the guessed branch instructions instead of doing  
 1168 nothing, waiting for the branch result. Most of the time, the branch predictor is correct  
 1169 and the processor gains time. Sometimes it is wrong, the processor state is restored  
 1170 and the execution starts again on the other branch, which is not slower than simply  
 1171 waiting for the branch result.

1172 The Spectre attack [KHF<sup>+</sup>20] trains the branch predictor to its advantage to exploit  
 1173 misprediction, i.e. when the branch predictor was wrong and started executing instruc-  
 1174 tions it shouldn’t have, in addition to the observable effects left after the roll-back, like

1175 cache content. In effect, the attacker is able to invert a test during the speculative win-  
 1176 dows. Many variants exist. Variant 1 allows an attacker to read from arbitrary memory,  
 1177 variant 2 corrupts indirect branch target computation to deviate the execution to an  
 1178 arbitrary location.

1179 **Load Value Injection (LVI).** This attack [VBMS<sup>+</sup>20] also exploits speculation execu-  
 1180 tion, but from the data perspective. When a load operation fails due to a page fault,  
 1181 the processor can transiently use data from a micro-architectural buffer to increase  
 1182 execution speed. The attacker controlling the buffer can inject arbitrary data from any  
 1183 load operation, including implicit load micro-op like the x86 *ret* instruction.

1184 **Race attacks.** Modern CPU cores can run multiple processes at once, interleaving their  
 1185 execution with a scheduler. If those processes share some part of memory, they can  
 1186 interfere in the other’s data flow [GSV03]. For instance, in a Time-Of-Check-To-Time-  
 1187 Of-Use (TOCTOU) attack [WP05], the victim program will first sanitize some user  
 1188 input, then the malicious program will overwrite said data with malicious and mal-  
 1189 formed content, which will be processed by the victim as if correctly formed, causing  
 1190 unintended behavior the attacker can exploit. Race attacks give write capabilities to  
 1191 the attacker, on data in memory which includes saved register states.

1192 **Conclusion.** To increase the execution speed of processors, micro-architectural behav-  
 1193 iors have been developed like scheduling or speculative execution. Unfortunately, those  
 1194 added behaviors also increase the attack surface of the target and offer new possibilities  
 1195 for the attacker to induce unintended behavior. A summary of the capabilities due to  
 1196 micro-architectural attacks is presented in Table 2.4.

Table 2.4: Attacker capabilities given by micro-architectural attacks

Type of attack	Attacker capability
Spectre	Invert a test (transiently)
	Read from arbitrary memory (transiently)
	Jump to arbitrary location (transiently)
LVI	Write data load results, from any, even implicit load (transiently)
Race attack	Write data in memory and registers

## 1197 2.2.6 Man-At-The-End (MATE) Attacks

1198 In the context of a MATE attack, the attacker has unlimited access to the target,  
 1199 meaning full observability and control over a software code and its execution [ASA<sup>+</sup>15].  
 1200 The attacker can look at what is executed, inspect each instruction and possibly modify  
 1201 it or any architectural element (register, memory) before executing it. For the victim,  
 1202 considering MATE attacks is equivalent to executing code on an untrusted environment.  
 1203 The goal of a MATE attack is usually to steal intellectual property such as sensitive  
 1204 data or code through reverse engineering attacks. This can be done with direct access to  
 1205 the device or remotely. Another type of MATE attacker is a malicious trusted user, an  
 1206 unhappy system administrator for instance. They can use their knowledge of a system  
 1207 and through legitimate access corrupt or erase data, or crash services, machines or an  
 1208 entire IT system. A MATE attacker is also in a unique position to observe protections  
 1209 and monitoring systems in place. Bypassing them is part of the attack. They can for  
 1210 instance redirect a hardware interrupt signal triggered by their attack actions.

1211 The associated attacker model is hence very powerful, with capabilities such as halt-  
1212 ing the execution, modifying data in memory and registers, and modifying instructions,  
1213 at any point of the execution.

### 1214 **2.2.7 Summary**

1215 We saw in this section that a wide variety of attacks vectors are available to an ad-  
1216 vanced attacker. Those attacks can be performed multiple time in one attack and in  
1217 combination. We described several real software-level security scenarios where the at-  
1218 tacker goes beyond crafting legitimate input to abuse the system under attack. The  
1219 effects of the attacks described can be represented using fault models.

## 1220 **2.3 Conclusion**

1221 Security-oriented IT systems are omnipresent. They need to be able to resist pow-  
1222 erful attackers able to exploit various kinds of attack techniques and attack surfaces.  
1223 We presented software attacks exploiting programs' inputs, hardware fault injection  
1224 attacks attacking the physical layer, software-implemented hardware attacks targeting  
1225 software-controlled hardware mechanisms, micro-architectural attacks misusing micro-  
1226 architecture behaviors, and MATE attacks with full control over the program and its  
1227 execution.

1228 It is important to be able to describe in a generic way the capabilities given to an  
1229 attacker by all those attack vectors. We do this by using fault models, inspired by the  
1230 hardware fault injection field. It allows to reason about attacks in an abstract way,  
1231 focusing on the impact an attacker can have on a program and evaluate a program's  
1232 resistance. In particular, a unified description of attacker capabilities is easier to inte-  
1233 grate into generic tools to assist security practitioners in their program-level security  
1234 assessment.

1235 To evaluate the vulnerability of programs to such an attacker model, from a program  
1236 hardening or vulnerability detection point of view, code analysis techniques exist as  
1237 detailed in the next chapter, Chapter 3.

## Background

## Contents

**3.1 Program Analysis Techniques . . . . . 26**

## 3.1.1 Overview . . . . . 26

## 3.1.1.1 Bugs in Programs . . . . . 26

## 3.1.1.2 Program Analysis . . . . . 26

## 3.1.2 Program Analysis is Undecidable . . . . . 27

## 3.1.2.1 Over-Approximation . . . . . 27

## 3.1.2.2 Under-Approximation . . . . . 28

## 3.1.3 Reachability Property . . . . . 29

## 3.1.4 Link to the Thesis . . . . . 30

**3.2 Symbolic Execution . . . . . 30**

## 3.2.1 Overview . . . . . 30

## 3.2.2 Symbolic Execution Algorithm . . . . . 31

## 3.2.3 Limitations . . . . . 32

## 3.2.3.1 Path Explosion . . . . . 32

## 3.2.3.2 Constraint Solving . . . . . 33

## 3.2.4 Link to the Thesis . . . . . 33

**3.3 Binary Code Analysis . . . . . 33**

## 3.3.1 Binary VS Source Code Analysis . . . . . 33

## 3.3.2 Challenges of Binary Analysis . . . . . 34

## 3.3.3 Link to the Thesis . . . . . 34

**3.4 Program Analysis Techniques for Fault Injection . . . . . 34**

## 3.4.1 Simulation . . . . . 34

## 3.4.2 Mutant Generation . . . . . 35

## 3.4.3 Forking Techniques . . . . . 35

## 3.4.4 Related Work . . . . . 36

## 3.4.4.1 Robustness Analysis . . . . . 36

## 3.4.4.2 Mutation Testing . . . . . 36

1269  
1270  
1272

3.4.5	Link to the Thesis	36
-------	--------------------	----

---

1273 We saw in Chapter 2 a program can include security properties that should with-  
1274 stand a certification procedure, its production environment and attackers. In this back-  
1275 ground chapter, we introduce program analysis techniques used to verify programs in  
1276 Section 3.1, with a stronger focus on symbolic execution (Section 3.2) and binary anal-  
1277 ysis (Section 3.3) that the work described in this manuscript builds upon. In addition,  
1278 we present in Section 3.4 the general topic of software-implemented fault injection,  
1279 which consists of tools to evaluate program properties in the presence of an attacker  
1280 able to inject faults.

## 1281 3.1 Program Analysis Techniques

1282 In this section, we introduce the concept of automated program analysis and present  
1283 the main families of analyzer properties such as under- and over-approximation, and  
1284 correctness and completeness. Lastly, we present the notion of reachability which we  
1285 will build upon in this thesis.

### 1286 3.1.1 Overview

1287 First, we provide an overview of why program analysis techniques are needed to prevent  
1288 bugs in programs and present some general program properties that can be verified.

#### 1289 3.1.1.1 Bugs in Programs

1290 Writing code without vulnerabilities or bugs is a hard problem, especially with systems  
1291 ever growing in complexity. Techniques such as good coding practices or code review  
1292 and manual testing help in the matter but do not provide strong guarantees. For  
1293 instance, the OSS-Fuzz platform performs continuous fuzzing on over 300 open-source  
1294 projects. It has discovered 23,907 bugs in 4 years [DLG21]. There is a need for program  
1295 analysis techniques able to detect vulnerabilities in programs.

#### 1296 3.1.1.2 Program Analysis

1297 To analyze a program is to determine if the *program behaviors satisfy a property*. The  
1298 terms ‘program behavior’ and ‘property’ can mean different things. Here, we consider  
1299 a program behavior to be the sequence of states it goes through following an execution  
1300 path, and a property to be an oracle over the program behaviors. Properties are mainly  
1301 divided into two classes, safety and liveness.

1302 **Safety.** A safety property prescribes that *something bad will never happen*. Some  
1303 examples are:

- 1304 – There is no division by 0;
- 1305 – The program never crashes;
- 1306 – *Memory safety*: no illegitimate, out-of-bounds memory accesses are performed;
- 1307 – The program does not deviate from specifications.

1308 **Liveness.** A liveness property prescribes that *something good will eventually happen* in  
1309 the program execution. Here are a few examples:

- 1310 – *Availability*: data provided to a program will eventually be treated;  
 1311 – *Termination*: a program will eventually terminate.

1312 **Formalization.** More formally, a program  $\mathbb{P}$  satisfy a property  $\mathcal{P}$ , written  $\mathbb{P} \models \mathcal{P}$  if  
 1313 and only if  $\mathbb{P} \subseteq \mathcal{P}$ . That is to say, all behaviors of  $\mathbb{P}$  are in the set  $\mathcal{P}$  of all possible  
 1314 program behaviors satisfying the property. Conversely, a program  $\mathbb{P}$  violates a property  
 1315  $\mathcal{P}$ , written  $\mathbb{P} \not\models \mathcal{P}$ , if and only if  $\exists c \in \mathbb{P}. c \notin \mathcal{P}$ . A program behavior that violates a  
 1316 property is called a *vulnerability* or a *bug*.

1317 Formal methods [CW96] are mathematically based techniques developed to assess  
 1318 properties on programs, or find counter-examples, i.e. bugs.

### 1319 3.1.2 Program Analysis is Undecidable

1320 Unfortunately, Rice theorem [Ric53] states that the task of determining if a program  
 1321 satisfies a non-trivial property is undecidable. A program analysis cannot be at the  
 1322 same time precise, automatic, terminating and generic in supported programs. In  
 1323 practice, program analysis techniques relax one or more of those attributes to get  
 1324 interesting results.

1325 For instance, deductive verification [HH19] is not automatic. It can require manual  
 1326 annotations such as pre-conditions, post-conditions, invariants, or to fully write proofs  
 1327 in some techniques.

1328 A number of program analysis techniques relax precision to obtain results. We ex-  
 1329 plore hereafter the two main categories, under-approximation and over-approximations.

#### 1330 3.1.2.1 Over-Approximation

1331 An analysis over-approximates all possible program behaviors  $\mathbb{P}$  by  $\mathcal{A}$ , a set of computed  
 1332 behaviors that include  $\mathbb{P}$  ( $\mathbb{P} \subseteq \mathcal{A}$ ), as illustrated in Figure 3.1. It provides a strong  
 1333 guarantee that the program satisfies the property when the analysis proves the over-  
 1334 approximated behavior satisfies the property. In particular, it can prove the absence  
 1335 of bugs.

$$\mathbb{P} \subseteq \mathcal{A} \wedge \mathcal{A} \subseteq \mathcal{P} \Rightarrow \mathbb{P} \subseteq \mathcal{P} \Rightarrow \mathbb{P} \models \mathcal{P}$$

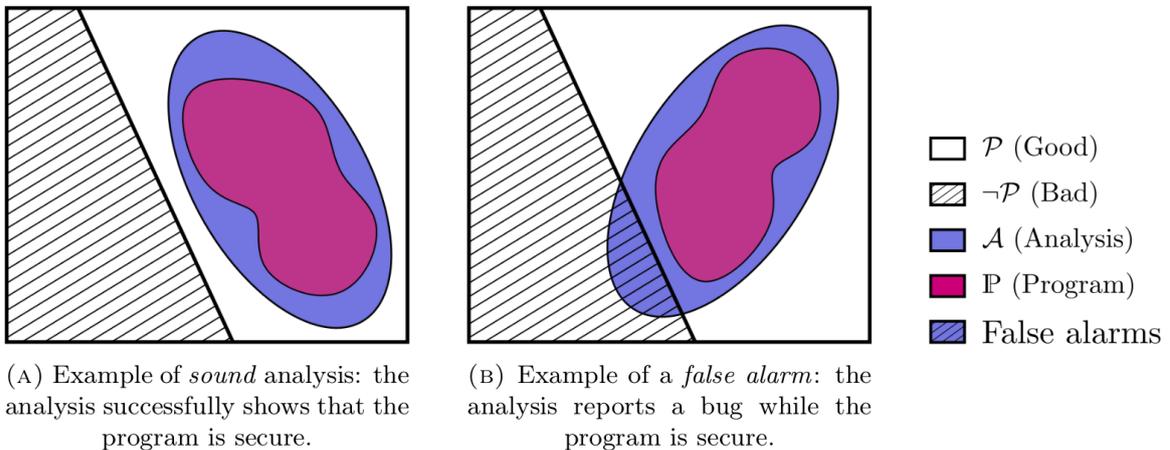


Figure 3.1: Illustration [Dan21] of an analysis  $\mathcal{A}$  that over-approximates the behaviors of a program  $\mathbb{P}$

1336 Because the analysis over-approximates program’s behaviors, when a bug is found,  
 1337 it may be a real one belonging to actual program behavior, or it can be a false positive,  
 1338 i.e. it is not part of program behavior.

1339 The notion of over-approximation relates to completeness. A program analysis  
 1340 technique is said *complete* when all program behaviors are studied. In particular, the  
 1341 analysis will find all violations of  $\mathcal{P}$ . An over-approximation-based technique is usually  
 1342 complete but not correct.

1343 An example of a program analysis technique based on over-approximation is *ab-*  
 1344 *stract interpretation* [Cou21, RY20]. The set of possible values for an object is over-  
 1345 approximated by a superset with interesting mathematical properties. A well-known  
 1346 example is representing possible integer values of a variable by an integer interval.

### 1347 3.1.2.2 Under-Approximation

1348 Another way to relax precision in an analysis is to under-approximate program be-  
 1349 haviors, i.e. by missing some of them. The set of computed program behaviors  $\mathcal{A}$   
 1350 is included in the actual program behaviors  $\mathbb{P}$  ( $\mathcal{A} \subseteq \mathbb{P}$ ) as illustrated in Figure 3.2.  
 1351 Here, if a bug is found by the analysis, it corresponds to a true problematic program  
 1352 behavior.

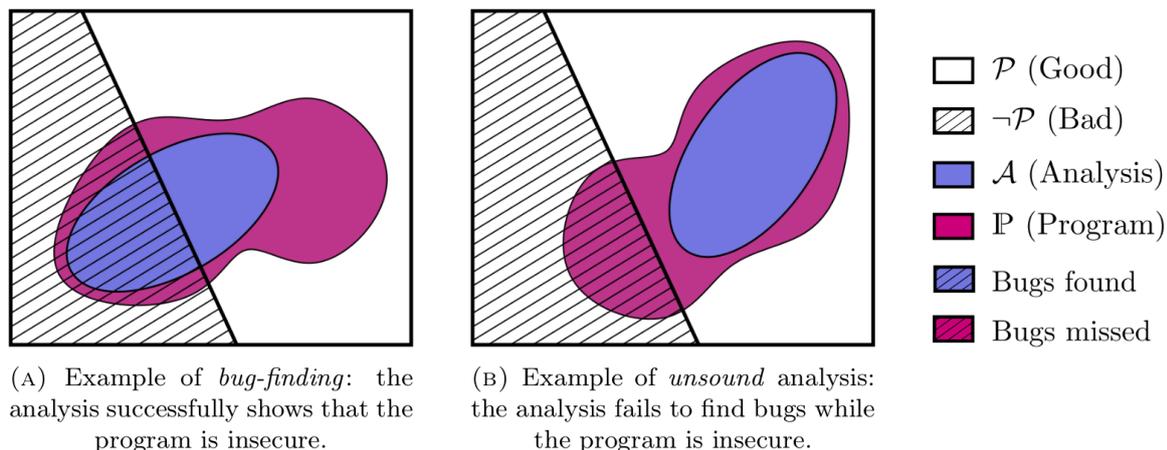


Figure 3.2: Illustration [Dan21] of an analysis  $\mathcal{A}$  that under-approximates the behaviors of a program  $\mathbb{P}$

$$\mathcal{A} \subseteq \mathbb{P} \wedge \exists c \in \mathcal{A}. c \notin \mathcal{P} \Rightarrow \exists c \in \mathbb{P}. c \notin \mathcal{P} \Rightarrow \mathbb{P} \neq \mathcal{P}$$

1353 A program analysis technique based on under-approximation will not have false  
 1354 positives but can miss bugs since some program behaviors elude it. Hence, such an  
 1355 analysis technique cannot prove the absence of bugs in a program.

1356 The notion of under-approximation is related to the correctness of an analysis. A  
 1357 program analysis is said *correct* when all bugs found are actual program behaviors. An  
 1358 under-approximation based technique is typically correct but not complete.

1359 There are some common program analysis techniques based on under-approximation.  
 1360 – An example of under-approximating program analysis technique is *fuzzing* [God20].  
 1361 Fuzzing consists in generating a large number of program inputs and detecting  
 1362 when the execution of the analyzed program violates a property using a monitor.  
 1363 A common property is the absence of crashes in a program. The coverage of such

1364 a technique is typically not complete, hence it misses some program behaviors,  
 1365 but correct, detected bugs are true ones.

- 1366 – *Model checking* [Cla97] consists in building a model of the studied program, and  
 1367 proving that properties hold for that model. As the model or paths can be infinite,  
 1368 *bounded model checking (BMC)* [CBRZ01] limits the depth of studied traces to a  
 1369 bound  $k$ . Bounded model checking is typically correct but not complete.
- 1370 – *Symbolic execution (SE)* [CS13] is a technique that interprets a program with  
 1371 symbolic (non-deterministic) inputs and explores reachability properties but can  
 1372 be non-terminating. As paths can be infinite in practice, symbolic execution is  
 1373 often bounded to finite paths of length lesser than a bound  $k$ . Symbolic execution  
 1374 in practice is typically correct but not complete. Another approximation angle  
 1375 consists in concretizing some part of the exploration, system calls in particular.  
 1376 This variant is called *dynamic symbolic execution (DSE)* or *concolic execution*  
 1377 [Sen07].

1378 The concept of completeness can be restricted to *bounded completeness*, also called  
 1379 *k-completeness*. An analysis is said to be  $k$ -complete when it explores all program  
 1380 behaviors constituting paths of length lesser than a bound  $k$ .

1381 Note that not all program analysis techniques are either under-approximation or  
 1382 over-approximations. They can be neither, like machine learning techniques for soft-  
 1383 ware vulnerability detection [LWH<sup>+</sup>20] for instance.

### 1384 3.1.3 Reachability Property

1385 We present in this section a formalization of the reachability property, that can be  
 1386 assessed by a program analysis technique. We will extend this definition of reachability  
 1387 to include the impact of an attacker in the work presented in this thesis.

1388 **Transition System.** Considering a program  $\mathbb{P}$ , we denote  $S$  the set of all possible states.  
 1389 A state is composed of the code memory, the data memory (i.e. the stack and heap),  
 1390 the state of registers and the location of the next instruction to execute. The set of  
 1391 input states of a program  $\mathbb{P}$  is noted  $S_0 \subseteq S$ . The set of transitions (or instructions)  
 1392 of the program is denoted  $T$ . The execution of an instruction  $t \in T$  is represented by  
 1393 a one-step transition relation  $\rightarrow_t \in S \times S$ . We denote  $s \rightarrow s'$  when  $s \rightarrow_t s'$  for some  
 1394 transition  $t \in T$ .

1395 **Execution Path.** We extend the transition relation over any finite path  $\pi \in T^*$  by the  
 1396 composition of a finite sequence of transitions, denoted  $\rightarrow_\pi$ .

**Reachability.** The transitive reflexive closure of  $\rightarrow$  is noted  $\rightarrow^*$ . A state  $s' \in S$   
 reachable in a program  $\mathbb{P}$  from a state  $s \in S$ , is noted  $s \rightarrow^* s'$ . We write for two states  
 $s, s' \in S$ :

$$s \rightarrow^* s' \Leftrightarrow \exists \pi \in T^*. s \rightarrow_\pi s'$$

1397 We use  $S \rightarrow s'$  as a shortcut for  $\exists s \in S. s \rightarrow s'$ , and  $\rightarrow_{\leq k}$  for reachability in at most  
 1398  $k$  steps.

1399 **Location Reachability.** We consider in the rest of the paper the case of *location reach-*  
 1400 *ability*: given a location  $l$  (instruction or code address) of the program under analysis,  
 1401 the question is whether we can reach any state  $s$  at location  $l$ . More formally,  $L$  is  
 1402 the finite set of locations of  $\mathbb{P}$ , and we consider a mapping  $loc : S \mapsto L$  from states to  
 1403 locations. For example,  $loc$  may return the program counter value. We write  $S \rightarrow^* l$   
 1404 as a shortcut for  $\exists s' \in S. S \rightarrow^* s' \wedge loc(s') = l$ .

1405 **Definition 3.1** (Standard reachability). *A location  $l$  is reachable in a program  $\mathbb{P}$  if*  
 1406  $S_0 \rightarrow^* l$ .

We extend this definition to location reachability in at most  $k$  steps, denoted:

$$S_0 \rightarrow_{\leq k}^* l$$

1407 We now define correctness and completeness for a program analyzer.

1408 **Definition 3.2** (Correctness, completeness). *Let  $\mathcal{V} : (\mathbb{P}, l) \mapsto \{1, 0\}$  be a verifier taking*  
 1409 *as input a program  $\mathbb{P}$  and a target location  $l$ .*

- 1410 –  $\mathcal{V}$  is correct when for all  $\mathbb{P}, l$ , if  $\mathcal{V}(\mathbb{P}, l) = 1$  then  $l$  is reachable in  $\mathbb{P}$ ;
- 1411 –  $\mathcal{V}$  is complete when for all  $\mathbb{P}, l$ , if  $l$  is reachable then  $\mathcal{V}(\mathbb{P}, l) = 1$ ;
- 1412 – if  $\mathcal{V}$  also takes an integer bound  $n$  as input,  $\mathcal{V}$  is  $k$ -complete when for all bound  
 1413  $n$  and  $\mathbb{P}, l$ , if  $l$  is reachable in at most  $n$  steps then  $\mathcal{V}(\mathbb{P}, l, n) = 1$ .

### 3.1.4 Link to the Thesis

Part of the work described in the thesis relates to automatic program analysis, which we leverage with the reachability property in mind.

## 1414 3.2 Symbolic Execution

1415 In this section, we delve deeper into one program analysis technique, symbolic exe-  
 1416 cution, as part of the work described in this thesis is based upon it. In particular,  
 1417 symbolic execution can assess location reachability in a program. After an overview,  
 1418 we present the main symbolic execution algorithm and its inherent limitations to keep  
 1419 in mind.

### 1420 3.2.1 Overview

1421 Symbolic execution (SE) [Kin76, GKS05, SMA05, CGP<sup>+</sup>08, CS13, BCD<sup>+</sup>18] is an auto-  
 1422 mated program analysis technique for bug finding and (bounded-)verification. Symbolic  
 1423 execution is a symbolic exploration technique for standard reachability. SE consists in  
 1424 simulating the execution of a program on symbolic inputs, i.e. of undetermined value,  
 1425 instead of concrete ones. The analysis keeps the program state as a symbolic state where  
 1426 program variables are expressed as terms of symbolic inputs. The analysis records a  
 1427 path explored as the conjunction of constraints associated with each conditional and  
 1428 indirect jump taken to follow the path, called a path predicate.

1429 SE has two main usages.

- 1430 – It can assess the reachability of a code location in a path by transforming the  
 1431 path predicate in an SMT formula discharged to an SMT solver [BT18] that we  
 1432 use as a black box.
  - 1433 – Either the solver answers ‘SAT’, meaning the formula is satisfiable and pro-  
 1434 vides a model for symbolic inputs such that executing the program with  
 1435 those inputs will reach the code location;
  - 1436 – Either the solver answers ‘UNSAT’ and the code location is not reachable  
 1437 from that path.

- 1438 – It is also possible to set a timeout to avoid having to wait a long time for  
 1439 the solver to answer in some edge cases. If the solver reaches the timeout  
 1440 then we cannot say anything about the reachability of the location.  
 1441 When used with an oracle detecting a bug, SE becomes a useful bug-finding  
 1442 technique, for instance, to detect use-after-free bugs [FMB<sup>+</sup>16].  
 1443 – It can generate test cases with the goal of maximizing the test suite coverage,  
 1444 each input covering one path.

### 1445 3.2.2 Symbolic Execution Algorithm

1446 Algorithm 3.1 gives a high-level view of a typical SE algorithm, adapted for location  
 1447 reachability. More complex properties can be verified with the same principles, such  
 1448 as local predicate reachability, trace properties or hyper-properties [DBR20].

---

**Algorithm 3.1:** Standard symbolic execution algorithm, taken from [GFB21]

---

**Input:** a program  $P$ , a bound  $k$ , a target location  $l$   
**Output:** Boolean value indicating whether  $l$  can be reached within  $k$  steps.

```

1 for path  $\pi$  in GetPaths( $k$ ) do
2   if  $\pi$  reaches  $l$  then
3      $\Phi :=$  GetPredicate( $\pi$ )
4     if  $\Phi$  is satisfiable then
5       return true
6     end
7   end
8 end
9 return false

```

---

1449 The analysis follows each possible path  $\pi$  of a program, enumerated by the function  
 1450 GetPaths, up to a depth bound  $k$ . If  $\pi$  reaches the target location  $l$ , then we check  
 1451 whether  $\pi$  is indeed feasible by computing its *path predicate*  $\Phi$  – a logical formula  
 1452 representing the path constraints over the input variables along  $\pi$ , and sending it to  
 1453 an SMT solver, that will try to answer whether the formula is satisfiable or not, and  
 1454 provide a model for free variables (e.g. inputs) if it is (omitted here for simplicity).

1455 **Property 3.1.** *SE is correct for location reachability, and  $k$ -complete if we assume a*  
 1456 *correct encoding of path predicates [God11].*

---

**Algorithm 3.2:** Assignment evaluation in SE

---

**Input:** path predicate  $\Phi$ , assignment instruction  $x := expr$   
**Output:** Updated  $\Phi$

```

1 Function eval_assign( $\Phi, x, expr$ ) is
2   return  $\Phi \wedge (x \triangleq expr)$ 
3 end

```

---

1457 In this thesis, we will focus on the evaluation of assignments and conditional jumps  
 1458 for SE, detailed in Algorithms 3.2 and 3.3 respectively, as this is where our adversarial  
 1459 symbolic execution will mainly differ from the standard one. It requires going slightly

**Algorithm 3.3:** Conditional jump evaluation in SE**Input:** path predicate  $\Phi$ , conditional jump instruction*if cdt then addr<sub>t</sub> else addr<sub>e</sub>***Data:** a worklist  $WL$  containing the pending path prefixes to explore – list of pairs (path predicate, next location)**Output:**  $WL$  updated in place

```

1 Function eval_conditional_jump( $\Phi$ ,  $cdt$ ,  $addr_t$ ,  $addr_e$ ) is
2   if  $\Phi \wedge cdt$  is satisfiable then
3     | Add ( $\Phi \wedge cdt$ ,  $addr_t$ ) to  $WL$ 
4   end
5   if  $\Phi \wedge (\neg cdt)$  is satisfiable then
6     | Add ( $\Phi \wedge \neg cdt$ ,  $addr_e$ ) to  $WL$ 
7   end
8 end

```

1460 deeper into details. In practice, the program paths are explored incrementally. A  
 1461 worklist  $WL$  records all pending paths together with their associated path predicate  
 1462 and their next instruction to explore.

- 1463 – Assignments (Algorithm 3.2) are dealt with straightforwardly, simply adding a  
 1464 new logical variable definition  $expr$  of the variable  $x$  to the path predicate  $\Phi$ ,  
 1465 written  $x \triangleq expr$ . Actually, a symbolic state usually comprises the path predicate  
 1466 itself plus a mapping from program variable names to logical variable names, and  
 1467 assignments involve both creating new logical names and updating the mapping.  
 1468 We abstract away from these details.
- 1469 – On conditional branches (Algorithm 3.3), the symbolic path is split in two, one  
 1470 for each branch ( $cdt$  and  $\neg cdt$ ), updating the path constraint  $\Phi$  accordingly. The  
 1471 satisfiability of each associated formula is checked before adding the updated path  
 1472 predicates to the worklist  $WL$ .

1473 **3.2.3 Limitations**

1474 Symbolic execution suffers two main limitations hindering performance and scalability:  
 1475 path explosion and constraint solving.

1476 **3.2.3.1 Path Explosion**

1477 SE explores all paths, which are possibly infinitely many, even those which will prove  
 1478 unfeasible, hence inducing a path explosion. A program typically contains various  
 1479 patterns generating new paths.

- 1480 – At each conditional jump, the path is split in two, one will follow the then branch,  
 1481 the other the else branch.
- 1482 – At indirect jump instructions, possible jump targets are enumerated and a new  
 1483 path is created for each one.
- 1484 – A loop whose termination is conditioned by a symbolic input becomes unbounded.  
 1485 SE unrolls the loop for each possible number of iterations, yielding possibly in-  
 1486 finitely many new paths.
- 1487 – Similarly, if a program contains a recursion whose termination is conditioned by a  
 1488 symbolic input, the analysis will call it for each possible recursion depth, yielding

possibly infinitely many new paths.

To tame path explosion during the exploration, different options are available:

- by identifying redundant paths that present a behavior already explored. Those can be pruned [BCE08];
- by merging paths [KP05] at the price of more complex formulas [KKBC12];
- by limiting the maximal depth of a path, hence performing bounded verification;
- by using search heuristics to guide the analysis toward exploring paths to maximize some coverage criteria [XTDHS09]. Only a subset of paths are then explored.

The last two techniques presented forgo completeness. By using such techniques, the analysis performs an under-approximation of the real program behavior and cannot be used to prove a property, but only for bug finding.

### 3.2.3.2 Constraint Solving

Symbolic execution generates formulas expressing the feasibility of a branch or the reachability of a location and discharges them to an SMT solver. SE analysis depends on the ability of the solver to answer those queries which tend to grow in complexity as the path explored is long and complex itself. SE performance is limited by the solver's performance, which is the main bottleneck of symbolic execution.

Analyses based on symbolic execution rely on different SMT theories, i.e. representations of structures more complex than boolean variables, such as integers or arrays, and their corresponding axioms. Some theories are notoriously hard to reason about, when they include quantifiers for instance. The work in this thesis extends a symbolic execution algorithm but does not need a more complex theory.

### 3.2.4 Link to the Thesis

Symbolic execution explores all program paths of bounded depth as a bug-finding technique for reachability. In this thesis, we adapt the symbolic execution engine to support fault injection and use it to assess the reachability of a code location in the presence of an advanced attacker.

## 3.3 Binary Code Analysis

In this section, we explore the specificities related to program analysis when applied to binary programs. We call binary code or binary program an executable composed of 0s and 1s, that the CPU will take as input. The assembly of a program, or its representation at ISA level, is binary code interpreted as instructions, in a human-readable form. The most well-known are Intel x86, ARM and RISC-V.

### 3.3.1 Binary VS Source Code Analysis

Binary code is a representation of what is executed on a machine, contrary to source code. Source code is a human-readable representation that contains details such as type information and data structures and hides others like memory VS register usage and flag updates. Moreover, to go from source code to binary code, compilers apply many transformation passes that:

- can remove part of the code considered to be dead code;
- can rewrite control and data flows, and rearrange the order of instructions;

- 1525 – can break properties that hold at source level like constant time or non-interference;
- 1526 – implements a specific behavior for each source-level undefined behavior;
- 1527 – can use optimizations not guaranteed to preserve the semantics of the source
- 1528 code. For instance, the optimization flag ‘-funsafe-math-optimizations’<sup>1</sup> for float-
- 1529 ing point arithmetic in gcc can violate language standards.

1530 This phenomenon is known as *What You See* (in source code) *Is Not What You eXecute*  
1531 (machine code), or WYSINWYX for short [BR10]. Hence, an analysis at the binary  
1532 level is closer to what is actually executed than an analysis at the source level.

1533 Another advantage of analyzing binary code is that it is available for any program  
1534 a user runs (except for remote reverse engineering). The associated source code is not  
1535 always provided, in particular when analyzing commercial, off-the-shelf executables, or  
1536 when studying malwares.

### 1537 3.3.2 Challenges of Binary Analysis

1538 Several challenges arise when analyzing programs at the binary level. A binary pro-  
1539 gram has more lines of code than the associated source code, challenging the scaling  
1540 capabilities of the analysis. Binary code is the encoding into ”0”s and ”1”s of an as-  
1541 sembly code (ISA), that encompasses implicit flag update. Those updates need to be  
1542 made explicit since they play an important role, for instance, in conditional jumps  
1543 where they determine the choice of the jump target. While source code analysis rea-  
1544 sons about typed variables, a binary code analysis requires precise reasoning about  
1545 the memory as assignments become load and store operations. It also needs to handle  
1546 explicit stack machinery. Finally, binary analysis tools tend to be harder to use than  
1547 source-level code, often requiring reverse engineering skills to configure the tool and  
1548 then understand the analysis results.

### 3.3.3 Link to the Thesis

In this thesis, we implemented our analysis technique at the binary level in order to be close to what is actually executed by the processor and how it is disturbed by fault injection.

## 1549 3.4 Program Analysis Techniques for Fault Injection

1550 In Chapter 2, we saw an advanced attacker could inject various kinds of faults into  
1551 a running program. Dedicated program analysis techniques have been designed to  
1552 account for those extra program behaviors in the assessment of a property. They are  
1553 called SoftWare-Implemented Fault Injection techniques (SWIFI). In this section, we  
1554 overview the main SWIFI techniques and their limitations.

### 1555 3.4.1 Simulation

1556 Dynamic analysis, also called *simulation* [DPdC<sup>+</sup>15, HSP20], consists in choosing con-  
1557 crete input and successively decoding and simulating the effects of decoded instructions  
1558 along the execution path. This technique can be used to simulate the effects of fault  
1559 injection. Typically, a golden run without fault is performed to identify all possible

---

<sup>1</sup><https://gcc.gnu.org/onlinedocs/gcc/Optimize-Options.html>

1560 injection locations. Then a simulation is launched for each concrete fault-induced in-  
 1561 struction modification. The final states are recorded and checked against the attacker’s  
 1562 goal.

1563 **Limitation.** Simulation only supports concrete inputs, forcing the analysis to only con-  
 1564 sider one reference execution path at the time, hence is not complete. Furthermore,  
 1565 simulation is limited by the explosion of simulations to run in particular when consid-  
 1566 ering an attacker able to perform multiple faults, faults with different effects, or as the  
 1567 program grows in size.

### 1568 3.4.2 Mutant Generation

1569 The *mutant generation* approach [CCG13, RG14, GWJLL17, CDFG18, GWJL20] con-  
 1570 sists in analyzing slightly modified versions of the program (named mutants), each of  
 1571 them embedding a different faulty instruction. For example, Figure 3.3 illustrates how  
 1572 a program containing few instructions (Figure 3.3a) can be faulted into three different  
 mutants 3.3b, 3.3c and 3.3d. Each mutant is then analyzed on its own using standard

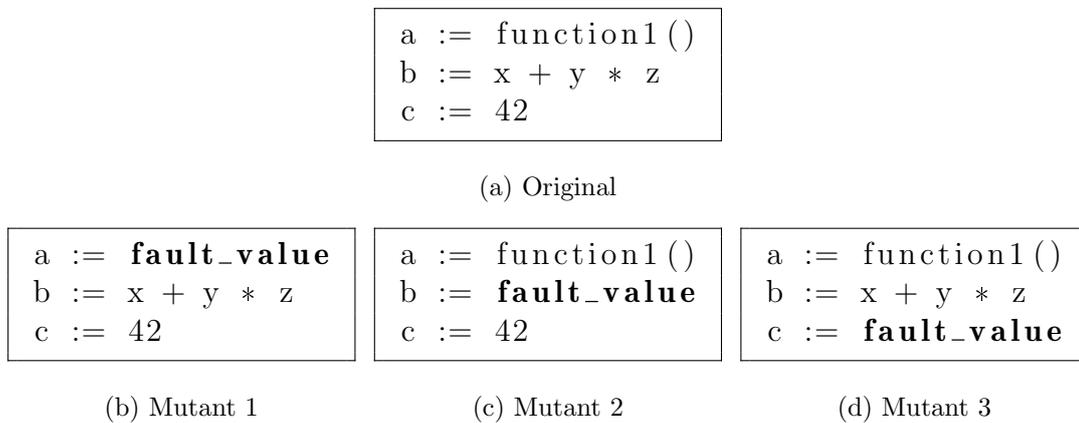


Figure 3.3: Mutant generation transformation in pseudo-code

1573 program analysis tools. Inputs can be symbolic or concrete. Mutant generation is  
 1574 usually used only for single-fault analysis.

1576 **Limitation.** The main limitation of mutant generation is the explosion of mutants, in  
 1577 particular when considering an attacker able to perform multiple faults, faults with  
 1578 different effects, or as the program grows in size. Also, as the different mutants differ  
 1579 only slightly, analyzing each of them separately wastes lots of time repeating similar  
 1580 reasoning.

### 1581 3.4.3 Forking Techniques

1582 The *forking approach* [PMPD14, BBC<sup>+</sup>14, BHE<sup>+</sup>19, LFBP21, Lan22] consists in in-  
 1583 strumenting the analysis (or the code, via instrumentation) to add all possible faults as  
 1584 forking points (branches) controlled by boolean values indicating whether a particular  
 1585 fault will be taken or not, plus constraints on the maximal number of faults allowed.  
 1586 A forking data fault is illustrated in Figure 3.4. The original statement (Figure 3.4a)  
 1587 is transformed into an if-then-else-like construct (Figure 3.4b). A standard program  
 1588 analysis technique is then launched – typically symbolic execution or bounded model  
 1589 checking. Compared with the mutant generation technique, this method allows sharing  
 1590 of parts of the analysis between the different possible faults.

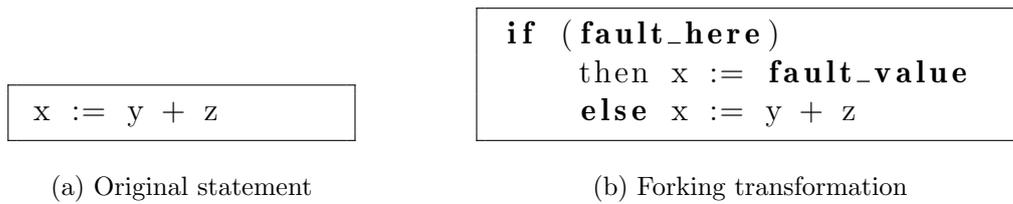


Figure 3.4: Forking technique transformation in pseudo-code

1591 **Limitation.** The number of paths explored by the analysis explodes with the number  
 1592 of possible faults (forking points).

### 1593 3.4.4 Related Work

1594 We list in the following the main topics related to software-implemented fault injection,  
 1595 namely robustness analysis and mutation testing.

#### 1596 3.4.4.1 Robustness Analysis

1597 SWIFI is also used for robustness evaluation [LH07, PNKI08, HTS+17, WTSS13,  
 1598 LHGD18, PLFP19, CDSL20, ZGWL+21], to consider the impact of *safety faults*  
 1599 instead of security related faults. This type of analysis aims to verify the correct  
 1600 behavior of error-handling mechanisms, implemented to recover from, or be resilient  
 1601 to, such safety faults. They can be generated by hardware damage, nuclear radiation or  
 1602 cosmic rays for instance. Robustness analyses rely also on mutant generation or forking  
 1603 techniques. The fault models are similar to hardware fault injection, yet multi-fault  
 1604 is not really an issue there, as faults are supposed to originate from safety issues and  
 1605 have no reason to accumulate unreasonably.

#### 1606 3.4.4.2 Mutation Testing

1607 Sometimes called software fault injection, mutation testing [PIK+18, CN13] aims to  
 1608 generate a comprehensive test suite by building test cases discriminating various mu-  
 1609 tants of a program, and is recognized as a very powerful testing criterion. As it fo-  
 1610 cuses on coverage, the mutant explosion cannot be avoided. Dedicated SE techniques  
 1611 [PM10, BCDK14, BKC14, MBK+18] have been designed.

### 3.4.5 Link to the Thesis

In this thesis, we present a new program analysis technique to represent the effect of fault injection on a program, with the aim of mitigating the state explosion faced by the existing techniques.

1613 Adversarial Reachability

1614 **Contents**

---

1615	<b>4.1 Attacker Model</b>	<b>39</b>
1616		
1617	4.1.1 Advanced Attacker	39
1618	4.1.2 Attacker Actions	39
1619	4.1.3 Fault Budget	39
1620	4.1.4 Attacker Goal	40
1621	<b>4.2 Adversarial Reachability</b>	<b>40</b>
1622	4.2.1 Reminder: Standard Reachability	40
1623	4.2.2 Adversarial Reachability Definition	40
1624	4.2.2.1 Adversarial Transitions	41
1625	4.2.2.2 Fault Expression	41
1626	4.2.2.3 State-of-the-art Fault Models	43
1627	4.2.2.4 Attacker Budget	44
1628	4.2.2.5 Adversarial Reachability	45
1629	4.2.2.6 Attacker Model Formalization	45
1630	4.2.3 Properties	45
1631	4.2.4 Discussion	46
1632	4.2.4.1 Other Properties	46
1633	4.2.4.2 Other Attacker Actions	46
1634	4.2.4.3 Location Reachability	46
1635	4.2.5 Conclusion	46
1636	<b>4.3 Forkless Adversarial Symbolic Execution (FASE)</b>	<b>47</b>
1637	4.3.1 Overview	47
1638	4.3.2 Forkless Fault Encoding	47
1639	4.3.2.1 Overview	47
1640	4.3.2.2 Different Forkless Encodings	48
1641	4.3.2.3 Supported Fault Models	48
1642	4.3.2.4 Forkless Faults in the Analysis	50

---

1643	4.3.2.5 Complexity Trade-off . . . . .	50
1644	4.3.3 Fault Injection Algorithm . . . . .	50
1645	4.3.3.1 Overview . . . . .	50
1646	4.3.3.2 Building the Adversarial Path Predicate . . . . .	51
1647	4.3.3.3 Properties . . . . .	53
1648	<b>4.4 Optimizations . . . . .</b>	<b>54</b>
1649	4.4.1 Early Detection of Fault Saturation (EDS) . . . . .	54
1650	4.4.1.1 Supported Fault Models . . . . .	54
1651	4.4.1.2 Algorithm . . . . .	55
1652	4.4.1.3 Properties . . . . .	55
1653	4.4.2 Injection on Demand (IOD) . . . . .	56
1654	4.4.2.1 Supported Fault Models . . . . .	56
1655	4.4.2.2 Algorithm . . . . .	56
1656	4.4.2.3 Properties . . . . .	58
1657	4.4.3 Combination of Optimizations . . . . .	58
1658	<b>4.5 Discussion . . . . .</b>	<b>59</b>
1659	4.5.1 Fault Model Support - Formalization VS Algorithm . . . . .	59
1660	4.5.2 Forkless Faults and Multi-Fault Analysis . . . . .	60
1661	4.5.3 Forkless Encoding for Other Properties . . . . .	60
1662	4.5.4 Forkless Encoding for Other Formal Methods . . . . .	61
1663	4.5.5 Forkless Encoding and Instrumentation . . . . .	62
1664	<b>4.6 Related Work . . . . .</b>	<b>62</b>
1665	4.6.1 Fault Model Support . . . . .	62
1666	4.6.2 Multiple Fault Analysis . . . . .	62
1667	4.6.3 The Attacker in Different Security Fields . . . . .	64
1668	4.6.4 Extending Existing Formalisms . . . . .	64
1669	<b>4.7 Conclusion . . . . .</b>	<b>65</b>

---

1673 In this chapter, we propose our model of an advanced attacker in Section 4.1, encom-  
1674 passing capabilities listed in Chapter 2. We describe a new formalization to represent  
1675 the impact of this advanced attacker on a program, called *adversarial reachability*, in  
1676 Section 4.2. Then, we design an algorithm to evaluate adversarial reachability, *forkless*  
1677 *adversarial symbolic execution* (FASE) in Section 4.3. The challenge is to prevent the  
1678 path explosion faced by the state of the art. We answer this challenge with a dedicated  
1679 *forkless encoding* of faults (Section 4.3.2) and optimizations (Section 4.4) aiming to re-  
1680 duce the number of injections without loss of generality. Discussions and comparisons  
1681 with related work are presented in Sections 4.5 and 4.6. The practical implementation  
1682 of the algorithms described in this chapter is detailed in Chapter 5.

1683 Part of the work presented in this chapter has been published at ESOP 2023  
1684 [DBP23].

## 1685 4.1 Attacker Model

1686 This section presents how we model an advanced attacker at program level. Advanced  
 1687 attackers can do more than carefully craft legitimate inputs to trigger vulnerabilities in  
 1688 software. They can use a wide variety of attack vectors (e.g. hardware fault injection  
 1689 attacks, software-implemented hardware attacks, micro-architectural attacks, software  
 1690 attacks, etc), in any combination, and multiple times. We suppose attack vectors  
 1691 prerequisites have been met, and only consider the impact of the generated faults on  
 1692 the program under attack.

### 1693 4.1.1 Advanced Attacker

1694 We start by providing a definition of what we mean by an advanced attacker at the  
 1695 program level.

1696 **Definition 4.1.** *We define an advanced attacker as an attacker able to perform multiple*  
 1697 *program-level actions during one attack, independently of the actual attack vector used.*

1698 We model an advanced attacker by an attacker model composed of the following  
 1699 components:

- 1700 – a set of attacker actions;
- 1701 – a fault budget;
- 1702 – an attacker goal.

1703 We describe those components in more detail in the following sections.

### 1704 4.1.2 Attacker Actions

1705 The first element of our attacker model is a set of attacker actions. Its purpose is to  
 1706 establish clear boundaries of the power of the attacker model on a program so it can  
 1707 fit a wide variety of security scenarios. An attacker action is composed of the following  
 1708 elements:

- 1709 – It encompasses the *type of modification* brought by the attack action, such as  
 1710 flipping a bit, changing a word or resetting a variable;
- 1711 – An action is also defined by the *object* it is applied to. It can be a register, an  
 1712 instruction op-code or a cell in memory for instance.

1713 Actions can also be described with a negative definition, that is to say, they can  
 1714 be explicitly restricted. For instance, if a strong memory isolation is supposed to be in  
 1715 place, it can be interesting to restrict a memory write capability to a certain range of  
 1716 addresses.

1717 We suppose an attacker can leverage multiple distinct attacker actions, possibly  
 1718 originating from different attack vectors, hence the attacker model includes a *set* of  
 1719 attacker actions.

1720 Attacker actions are similar to fault models used to describe the effects of a hard-  
 1721 ware fault injection on a program. Hence, in the rest of this manuscript, we will use  
 1722 interchangeably *attacker capability*, *attacker action* and *fault model*.

### 1723 4.1.3 Fault Budget

1724 Now that our attacker has different actions at their disposal, they can perform them  
 1725 at will on the program. To restrict the power of the attacker model we introduce a

1726 maximum number of actions that can be performed during one attack. We call it the  
1727 *fault budget* of the attacker model. An attacker with infinite power can do anything  
1728 with the program, including disarming potential protections, which is not always very  
1729 interesting to consider or realistic in a threat assessment.

#### 1730 4.1.4 Attacker Goal

1731 Finally, the attacker is trying to achieve something with their attack. We model this  
1732 attack goal by an oracle over the state of the program, and the aim of the attacker is  
1733 to reach a state of the program satisfying the oracle. A simple oracle is a code location  
1734 the attacker may want to reach. More complex functions of the program state can be  
1735 defined such as properties of the program's variables or trace properties.

1736 The ability to reach a state satisfying a property is the standard notion of *reachability*.  
1737 However, here, the attacker can disrupt the execution of the program with their  
1738 capabilities in order to reach such a state. Hence we define the ability to reach a pro-  
1739 gram state satisfying a property in the presence of an advanced attacker as *adversarial*  
1740 *reachability*. This concept is detailed in Section 4.2.

## 1741 4.2 Adversarial Reachability

1742 In this section, we describe our formalism to reason about a program's execution in the  
1743 presence of an advanced attacker. We consider a program as a transition system, where  
1744 the attacker's capabilities are expressed as new transitions. Adversarial reachability is  
1745 standard reachability in this modified transition system. We start this section with a  
1746 reminder of the notations used in Section 3.1.3 for standard reachability, then we define  
1747 adversarial reachability and related properties.

### 1748 4.2.1 Reminder: Standard Reachability

1749 We use the following notations to describe a program as a transition system and express  
1750 standard reachability:

- 1751 – a program  $\mathbb{P}$ ;
- 1752 – the set  $S$  of all possible states,  $S_0$  the set of initial states of  $\mathbb{P}$ ;
- 1753 – the set  $T$  of transitions (or instructions) of the program  $\mathbb{P}$ ;
- 1754 – the execution of an instruction  $s \rightarrow s'$  for some transition  $t \in T$ , with  $s, s' \in S$ ;
- 1755 – an execution path  $\rightarrow_\pi$ , with  $\pi \in T^*$ ;
- 1756 – the standard reachability of a state  $s \in S$  is  $S_0 \rightarrow^* s$ , as a short way to write  
1757  $\exists s_0 \in S_0. s_0 \rightarrow^* s$ ;
- 1758 – the standard reachability of a location  $l \in L$  is  $S_0 \rightarrow^* l$  as a short way of writing  
1759  $\exists s \in S. S_0 \rightarrow^* s \wedge \text{loc}(s) = l$ .

1760 We refer the reader to Section 3.1.3 for the definition of correctness and completeness  
1761 for standard reachability.

### 1762 4.2.2 Adversarial Reachability Definition

1763 We now build upon the transition system and notations previously described to include  
1764 the presence of an advanced attacker and detail how attacker actions are encoded as  
1765 faults.

1766 **4.2.2.1 Adversarial Transitions**

1767 **Adversarial Transitions.** We extend the transition model previously described to include *adversarial transitions* denoted  $\rightsquigarrow_{\mathcal{A}} \in S \times S$ , related to an attacker model  $\mathcal{A}$ . The  
 1768 total set of possible transitions in the context of an attacked program is composed of  
 1769 legitimate program transitions and attacker transitions, noted  $T_{\mathcal{A}} = T \cup \rightsquigarrow_{\mathcal{A}}$ .  
 1770

1771 Then, the transition relation of  $\mathbb{P}$  under attack from an attacker model  $\mathcal{A}$  is denoted  
 1772 as  $\mapsto_{\mathcal{A}} = \rightarrow \cup \rightsquigarrow_{\mathcal{A}} = (\cup_{t \in T} t) \cup \rightsquigarrow_{\mathcal{A}}$ .

1773 **Remark.** The successor of a state in the execution of the program has no longer a  
 1774 unique solution, as either the standard transition  $\rightarrow$  corresponding to the legitimate  
 1775 program instruction, or one or more adversarial transitions  $\rightsquigarrow_{\mathcal{A}}$  can be taken.

1776 **Adversarial Path.** We denote  $\pi \in T_{\mathcal{A}}^*$ , an *adversarial path*, representing a succession  
 1777 of program instructions and adversarial actions.

1778 **Program States.** We consider a program state  $s \in S$  to be composed of:

- 1779 – the code memory  $M_c$ ,
- 1780 – the data memory (i.e. the stack and heap)  $M_d$ ,
- 1781 – the set of registers  $R$ ,
- 1782 – the location of the next instruction to execute  $ip$ .

1783 The location of the next instruction to execute is often contained in a register, we single  
 1784 it out in this model for clarity.

1785 **Transitions Granularity.** A transition, adversarial or not, can only modify the next  
 1786 instruction pointer and one other element of the state. For instance, we do not consider  
 1787 writing at multiple locations in the data memory as one transition.

1788 **4.2.2.2 Fault Expression**

1789 We now illustrate various types of faults and how they can be represented as adversarial  
 1790 transitions.

1791 **Faults In-Between Instructions.** To represent a fault happening between normal tran-  
 1792 sitions, in addition to legitimate program behavior, the adversarial transition corrupts  
 1793 the relevant part of the program state, but leaves the next instruction pointer un-  
 1794 changed.

1795 We take the example of data faults corrupting a value in the data memory, with  
 1796  $addr$  an address in the data memory and  $fault\_v$  a corrupted value.

$$s\{M_c, M_d, R, ip\} \rightsquigarrow s'\{M_c, M_d\{addr \leftarrow fault\_v\}, R, ip\}$$

1797 This fault model corresponds for instance to a Rowhammer bit-flip, happening in  
 1798 memory independently of the program execution.

1799 Depending on the grammar complexity of the language modeled by the transition  
 1800 system, faults in-between can have varying expressivity.

- 1801 – For instance, in assembly, where every load and intermediate operation result  
 1802 are stored in intermediary variables, typically in registers, a fault in-between  
 1803 assembly instructions can express all possible combinations of corruption of the  
 1804 end result;
- 1805 – On the opposite, if complex operation aggregations are allowed by the grammar,  
 1806 having a load operation as one element of a binary operation without intermediate  
 1807 storing for instance, then an in-between fault could not fault just the result of

1808 the load, hence cannot express corruptions between the memory and the data  
1809 treatment unit.

1810 **Specific Fault Behavior.** To represent a specific behavior more precisely, the corrupted  
1811 value can be expressed as a function. Here, we denote  $Fault$  such a function.

$$s\{M_c, M_d, R, ip\} \rightsquigarrow s'\{M_c, M_d\{addr \leftarrow Fault(M_d[addr])\}, R, ip\}$$

$$s_{i-1} \xrightarrow{t_i} s_i \xrightarrow{t_{i+1}} s_{i+1} \xrightarrow{t_{i+2}} s_{i+2}$$

$$s_{i-1} \xrightarrow{t_i} s_i \rightsquigarrow s'_i \xrightarrow{t_{i+1}} s'_{i+1} \xrightarrow{t_{i+2}} s'_{i+2}$$

$$s_{i-1} \xrightarrow{t_i} s_i \rightsquigarrow s'_{i+1} \xrightarrow{t_{i+2}} s'_{i+2}$$

To precisely express the bit-flip from Rowhammer, we take

$$Fault(v) = v \oplus (1 \ll bv)$$

1812 with  $v$  a value,  $\oplus$  denoting a bitwise xor operator and  $bv$  indicating the bit that is  
1813 flipped.

1814 We now discuss the impact of using a  $Fault$  function expressing the fault behavior  
1815 on in-between fault expressivity:

- 1816 – If the  $Fault$  function cannot access the inner representation of the expression be-  
1817 hind  $v$ , typically explore its syntax tree, but only use  $v$  as such, then expressivity  
1818 restrictions discussed above for in-between faults remain valid;
- 1819 – On the contrary, if the  $Fault$  function can access and modify the inner repre-  
1820 sentation of  $v$ , restrictions are lifted. Coming back to the load inside a binary  
1821 operation example, such a  $Fault$  function could fault only the load part, equiva-  
1822 lent to faulting a solitary intermediate load.

1823 **Faults Replacing an Instruction.** On the opposite of faults in-between instructions,  
1824 faults can be modeled as replacing a legitimate instruction, changing its original effect.  
1825 The adversarial transition ‘hijacks’ a legitimate transition, mimicking it but with faulty  
1826 behavior, and updating the instruction pointer to directly go to the next instruction,  
1827 in effect bypassing the legitimate transition.

1828 Staying with the data fault in memory example, a legitimate store instruction of  
1829 the value contained in the register  $r$ , at the address  $addr$ , written

$$s\{M_c, M_d, R, ip\} \rightarrow s'\{M_c, M_d\{addr \leftarrow R[r]\}, R, ip \leftarrow ip + 1\}$$

1830 can be replaced by the following adversarial transition

$$s\{M_c, M_d, R, ip\} \rightsquigarrow s'\{M_c, M_d\{addr \leftarrow fault\_value\}, R, ip \leftarrow ip + 1\}$$

1831 This type of fault represents for instance the corruption of a stored value before it  
1832 reaches the memory.

1833 Most in-between faults can be equivalently expressed as a replacing fault. The same  
1834 expressivity considerations apply.

- 1835 – As in-between faults are somewhat independent of the execution, as long as they  
1836 happen before the corrupted value usage, they can be pushed ahead to the pre-  
1837 vious store. A replacing fault can then be used instead;

- 1838 – Only data in the initial state cannot be corrupted with a replacing fault, as no  
 1839 instruction is storing them. However, each usage of this initial data can be faulted  
 1840 with a replacing fault;  
 1841 – Delaying the fault until the corrupted value usage and using an adversarial tran-  
 1842 sition of each usage of that value until it is overwritten, if at all (global initial  
 1843 variables are often constants), is equivalent to in-between faults. However, this  
 1844 requires much more faults in the attacker’s budget.

1845 **Transient & Permanent Faults.** Transient faults refer to faults whose effects last only  
 1846 until the affected value is overwritten. On the contrary, permanent faults are faults  
 1847 that cannot be overwritten during the execution of the program. This can represent  
 1848 permanent damage in the chip containing code or data memory, or faults in the code  
 1849 memory when it is not changed during the execution of the program, which is usually  
 1850 the case, except for self-modifying programs.

1851 The formalization presented is well-suited for transient faults, as illustrated by in-  
 1852 between and replacing faults previously described. We can also express permanent  
 1853 faults in the code memory, assuming it is unchanged during the execution.

$$s\{M_c, M_d, R, ip\} \rightsquigarrow s'\{M_c\{addr \leftarrow fault\_instr\}, M_d, R, ip\}$$

1854 The new instruction *fault\_instr* will be used instead of the original one for the  
 1855 remainder of the execution.

1856 However, the state model used does not include the memory physical layer and  
 1857 thus prevents us from expressing permanent fault on other parts of the program state  
 1858 like the data memory or registers. Hence, this formalism cannot express faults such as  
 1859 stuck-at bits (except having an in-between fault applying the fault effect each time).  
 1860 A possibility to include permanent faults on data would be to add a mapping of the  
 1861 physical layer to the values in the state, where a mask and fault effects would be applied  
 1862 to each read and write of values.

### 1863 4.2.2.3 State-of-the-art Fault Models

1864 To specify practical fault models, restrictions are applied onto  $\rightsquigarrow_{\mathcal{A}}$ , limiting what part  
 1865 of the state can be modified and how. In particular, relevant *Fault* functions specifying  
 1866 the fault behavior can be defined for each fault model available to the attacker. We  
 1867 now explore a few common fault models.

**Data Faults.** Data faults have been described above, for arbitrary data faults and  
 bit-flips. Extensions to other types of data faults can be derived from those. We write  
 here a reset fault on the register *r*.

$$s\{M_c, M_d, R, ip\} \rightsquigarrow s'\{M_c, M_d, R\{r \leftarrow 0\}, ip \leftarrow ip + 1\}$$

**Instruction Skip Fault.** An instruction skip consists of bypassing any sequence of  
 consecutive instructions with the following adversarial transition.

$$s\{M_c, M_d, R, ip\} \rightsquigarrow s'\{M_c, M_d, R, ip \leftarrow ip + n\}$$

1868 We write here *ip + n* to denote the *n*th next instruction address in the code memory  
 1869 layout. A single instruction skip is written for *n* = 1.

1870 Interestingly, this representation of instruction skips can model the corruption of the  
 1871 program counter, modeled by *ip* here. But it also encompasses replacing an instruction

1872 with a no-operation (NOP), as the proposed adversarial transition does not modify the  
1873 state and goes to the next instruction.

1874 **Replaying Instructions.** Some observations of instructions skips report a replay of the  
1875 previous instructions, in particular for attacks on the instructions prefetch buffer. For  
1876 instance, to skip 4 instructions  $n$  to  $n + 3$ , the instructions up to  $n - 1$  are executed  
1877 normally, and then instructions  $n - 4$  to  $n - 1$  are executed again, before going straight  
1878 to instruction  $n + 4$ . This can be represented with 2 adversarial transitions, one setting  
1879  $ip \leftarrow ip - 4$  just before the instruction  $n$  should be executed, and one adversarial  
1880 transition  $ip \leftarrow ip + 4$  just before the instruction  $n$  should again be played.

1881 However, for those two adversarial transitions to be considered as one fault, extra  
1882 elements need to be added to the program state to track the relation between instruc-  
1883 tions.

**Test Inversion.** We consider a conditional jump to be written as follows, with a con-  
ditional expression written  $(a ? b : c)$  where if  $a$  is true, then the expression evaluates  
to  $b$  and to  $c$  otherwise.  $ip_A$  and  $ip_B$  are two instructions pointers.

$$s\{M_c, M_d, R, ip\} \rightarrow s'\{M_c, M_d, R, ip \leftarrow R[r] ? ip_A : ip_B\}$$

The fault model consisting of inverting a test can be written multiple ways, cor-  
rupting the test condition before the conditional jump for instance.

$$s\{M_c, M_d, R, ip\} \rightsquigarrow s'\{M_c, M_d, R\{r \leftarrow \neg r\}, ip\}$$

Or the conditional jump can be replaced with an adversarial conditional jump  
inverting the instructions pointers.

$$s\{M_c, M_d, R, ip\} \rightsquigarrow s'\{M_c, M_d, R, ip \leftarrow R[r] ? ip_B : ip_A\}$$

1884 The test inversion fault model illustrates that one fault can be expressed with  
1885 different adversarial transitions.

1886 **Limits and Extensions.** We discuss now some limits of the current formalization and  
1887 possible extensions.

- 1888 – Complex, multi-effect fault models can be split into their atomic components,  
1889 modeled separately. However, the number of fault count needs to account for this  
1890 split. Moreover, synchronizations between those atomic adversarial transitions  
1891 may require extending the state with some new elements, counters for instance;
- 1892 – This formalization does not include any micro-architectural detail. As discussed  
1893 for instruction replay and permanent data faults, the representation of the pro-  
1894 gram state can be extended to account for more fault models.

#### 1895 4.2.2.4 Attacker Budget

1896 Still, we need to take into account the maximum number of faults the attacker can  
1897 perform along an execution path. Given a path  $\pi$  over  $T_{\mathcal{A}}^*$ ,  $\pi$  is said to be *legit* if it  
1898 does not contain  $\rightsquigarrow_{\mathcal{A}}$ , and *faulty* otherwise. The number of occurrences of transitions  
1899  $\rightsquigarrow_{\mathcal{A}}$  in  $\pi$  is its *number of faults*.

1900 Given a bound  $m_{\mathcal{A}}$  on the fault capability of  $\mathcal{A}$ , we define  $\mapsto_{(\mathcal{A}, m_{\mathcal{A}})}^*$  by limiting the  
1901 adversarial reachability relation to paths  $\pi$  with less than  $m_{\mathcal{A}}$  faults. We consider  $m_{\mathcal{A}}$   
1902 to be  $+\infty$  in case the attacker has no such limitation. For the sake of simplicity, in the  
1903 following, we will consider  $m_{\mathcal{A}}$  as an implicit parameter of  $\mathcal{A}$ , and simply write  $\mapsto_{\mathcal{A}}^*$   
1904 instead of  $\mapsto_{(\mathcal{A}, m_{\mathcal{A}})}^*$ .

### 1905 4.2.2.5 Adversarial Reachability

1906 We extend the notations for state reachability to the relation  $\mapsto_{\mathcal{A}}$ . Especially,  $S \mapsto_{\mathcal{A}}^* s'$   
 1907 means  $\exists s \in S. s \mapsto_{\mathcal{A}}^* s'$ . The adversarial transition relation up to  $k$  is denoted  $\mapsto_{\mathcal{A}, \leq k}$   
 1908 and reaching a location  $l$  from a set of states  $S$  will be denoted  $S_0 \mapsto_{\mathcal{A}}^* l$ .

1909 **Definition 4.2** (Adversarial reachability). *Given an attacker  $\mathcal{A}$  with a  $m_{\mathcal{A}}$  fault budget  
 1910 and a program  $\mathbb{P}$ , a location  $l \in L$  is adversarially reachable if  $S_0 \mapsto_{\mathcal{A}}^* s' \wedge \text{loc}(s') = l$   
 1911 for some  $s' \in S$ .*

1912 In the following, the adversarial reachability of location  $l$  from a set of states  $S$  will  
 1913 be denoted  $S_0 \mapsto_{\mathcal{A}}^* l$ .

1914 We extend the adversarial reachability definition to adversarial reachability in at  
 1915 most  $k$  steps, noted  $S_0 \mapsto_{\mathcal{A}, \leq k}^* l$ .

### 1916 4.2.2.6 Attacker Model Formalization

1917 Finally, we can write the formalization of our attacker model representing an advanced  
 1918 attacker using adversarial reachability.

**Definition 4.3** (Attacker model). *We define an attacker model in this formalism by*

$$\mathcal{A} = \{\{\rightsquigarrow\}, m_{\mathcal{A}}, \phi\}$$

- 1919 – the adversarial transitions  $\{\rightsquigarrow\}$  correspond to attacker capabilities;
- 1920 –  $m_{\mathcal{A}}$  is the maximum number of faults allowed;
- 1921 – the goal of the attacker is noted  $\phi \in S \times \{0, 1\}$ . If  $\exists s \in S$  such that  $s$  is ad-  
 1922 versarially reachable and  $\phi(s) = 1$ , then the attack is successful. For location  
 1923 reachability of a location  $l_{\mathcal{A}}$ , we take  $\phi(s) = (\text{loc}(s) == l_{\mathcal{A}})$ .

### 1924 4.2.3 Properties

1925 We now discuss the properties of adversarial reachability and of a verifier for adversarial  
 1926 reachability.

1927 **Proposition 4.1.** *Standard reachability implies adversarial reachability. The converse  
 1928 does not hold.*

1929 *Proof.* Standard reachability can be viewed as adversarial reachability with an attacker  
 1930 able to perform 0 faults. □

1931 We redefine what it means for an analysis answering adversarial reachability to be  
 1932 correct, complete and  $k$ -complete.

1933 **Definition 4.4.** *Let  $\mathcal{V}_{\mathcal{A}} : (\mathbb{P}, A, l) \mapsto \{1, 0\}$  be a verifier taking as input a program  $\mathbb{P}$ , an  
 1934 attacker  $A$  with  $m_{\mathcal{A}}$  fault budget and a target location  $l$ .*

- 1935 –  $\mathcal{V}_{\mathcal{A}}$  is correct given  $A$  when for all  $\mathbb{P}, l$ , if  $\mathcal{V}_{\mathcal{A}}(\mathbb{P}, A, l) = 1$  then  $l$  is adversarially  
 1936 reachable in  $\mathbb{P}$  for attacker  $A$ ;
- 1937 –  $\mathcal{V}_{\mathcal{A}}$  is complete given  $A$  when for all  $\mathbb{P}, l$ , if  $l$  is adversarially reachable for attacker  
 1938  $A$  then  $\mathcal{V}_{\mathcal{A}}(\mathbb{P}, A, l) = 1$ ;
- 1939 – if  $\mathcal{V}_{\mathcal{A}}$  also takes an integer bound  $n$  as input,  $\mathcal{V}_{\mathcal{A}}$  is  $k$ -complete given  $A$  when  
 1940 for all integer  $n$  and  $\mathbb{P}, l$ , if  $l$  is adversarially reachable in at most  $n$  steps then  
 1941  $\mathcal{V}_{\mathcal{A}}(\mathbb{P}, A, l, n) = 1$ .

## 1942 4.2.4 Discussion

1943 In this section, we take a step back and discuss what kind of properties can be verified  
1944 using our formalization and how it could be extended to other types of attacker actions.  
1945 We also discuss the advantages of studying location reachability.

### 1946 4.2.4.1 Other Properties

1947 Our formalism itself is quite generic and can accommodate a wide range of properties,  
1948 as we mainly keep the property unchanged but modify the underlying transition system.  
1949 We focus in this work on the reachability property in this extended transition system.  
1950 However, more complex properties than location reachability can be studied. We list  
1951 here some examples.

- 1952 – We can have an oracle over the state more complex, for instance looking for  
1953 use-after-free or buffer overflows.
- 1954 – The oracle could be a monitor for runtime errors.
- 1955 – We could imagine an attacker willing to activate a non-terminating execution  
1956 (denial of service).

1957 Any property studied based on a transition system could be studied in the presence  
1958 of an attacker with our extended transition system.

### 1959 4.2.4.2 Other Attacker Actions

1960 The proposed formalism only considers active faults, i.e. faults modifying the state.  
1961 It does not include the possibility of passive faults, i.e. an attacker read action, nor  
1962 an attacker able to execute multiple times a program and learn from it until they can  
1963 achieve their goal.

1964 While those are out of the scope of this thesis, we believe they could benefit from  
1965 our formalization, by extending the attacker model with a knowledge component, and  
1966 the transition system with a leakage model. Those extensions would allow to represent  
1967 an attacker performing combined passive and active attacks [AVFM07].

### 1968 4.2.4.3 Location Reachability

1969 We want to stress that while location reachability can be seen as a basic case, we  
1970 consider it sufficient here for two reasons:

- 1971 – first, it keeps the formalism light while still straightforward to generalize to  
1972 stronger reachability properties (e.g., local predicates of the form  $(l, \varphi)$ , sets of  
1973 finite traces, etc.);
- 1974 – second, it is already rather powerful on its own, as we can still instrument the code  
1975 to reduce some stronger forms of reachability to it (e.g., adding local assertions  
1976 or monitors).

## 1977 4.2.5 Conclusion

1978 We proposed a formalization of an advanced attacker and their impact on a program.  
1979 We detail the concept of adversarial reachability as an attacker’s goal. Modeling ad-  
1980 vanced attackers allows to reason about them and ultimately strengthen programs  
1981 against those attackers. We also defined what it means for a verifier to be correct and  
1982 complete with respect to adversarial reachability. In the following section, we propose  
1983 a technique to assess adversarial reachability in practice.

## 1984 4.3 Forkless Adversarial Symbolic Execution (FASE)

1985 In this section, we present our algorithm to verify adversarial reachability. It aims to  
 1986 represent in practice the impact of an advanced attacker on a program. We mitigate  
 1987 the path explosion faced by existing techniques (Section 3.4) with a forkless encoding  
 1988 of attacker actions.

### 1989 4.3.1 Overview

1990 We conceived our fault injection technique with the following design guidelines in mind.

- 1991 – We want a technique that *prevents the path explosion* faced by the state-of-the-  
 1992 art, and is *generic* in supported fault models. We introduce in Section 4.3.2 the  
 1993 forkless fault encoding that embedded the activation and effect of a fault in a  
 1994 non-forking manner, and that is able to support many different fault models;
- 1995 – We want an automated technique that is *correct* and *k-complete* for adversarial  
 1996 reachability. We achieve this by basing our technique on symbolic execution as  
 1997 described in Section 4.3.3, which is correct and k-complete for standard reacha-  
 1998 bility;
- 1999 – We want to *limit the added complexity of generated SMT queries*, as query reso-  
 2000 lution is the main bottleneck of symbolic execution, which we burden with more  
 2001 complexity due to the fault encodings. We designed two optimizations presented  
 2002 in Section 4.4, Early Detection of fault Saturation (EDS) and Injection On De-  
 2003 mand (IOD) to reduce the number of fault injections in SMT formulas.

### 2004 4.3.2 Forkless Fault Encoding

2005 We focus now on our forkless encoding of faults and the fault models it supports. It is  
 2006 using this encoding that we will integrate the attacker’s actions into our analysis.

#### 2007 4.3.2.1 Overview

2008 We use the term fault encoding to denote the replacement of a program instruction (or  
 2009 group of) by some other instruction(s) with a different, illegitimate, behavior.

2010 The forkless encoding aims to address the path explosion induced by the forking  
 2011 treatment of fault injection in prior works. It focuses on data corruption, which can  
 2012 be values or addresses, leading to a wide variety of supported fault models. The  
 2013 forkless encoding consists of wrapping arithmetically an expression and embedding the  
 2014 activation of the fault and its effect.

- 2015 – The **fault activation** is a boolean variable whose value conditions if the fault is  
 2016 active, i.e. modifies the normal program behavior, or if it is inactive and can  
 2017 be considered as not present. A new activation variable is generated for each  
 2018 possible fault location;
- 2019 – The **fault effect** is a value corresponding to the corruption applied to the program  
 2020 and depends on the fault model.

2021 We show in Figure 4.1 an arbitrary data fault on the assignment of  $x$  receiving  
 2022 the expression  $expr$ . The encoding uses an *ite* operator, an inlined form of if-then-  
 2023 else. The activation of this fault location is determined by the symbolic boolean value  
 2024  $fault\_here$ , and the corrupted value of  $x$ , i.e. the fault effect, is the fresh variable  
 2025  $fault\_value$ .

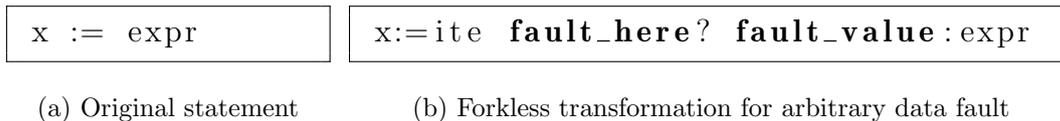


Figure 4.1: Forkless injection technique

2026 The point is to embed the fault injection as an arithmetical expression inside the  
 2027 path predicate, *without any explicit path forking* at the analysis level, in order to let  
 2028 the analyzer reason about both legit executions and faulty executions at the same time  
 2029 – this is akin to path merging in some ways, except that we do it only for the treatment  
 2030 of fault injection (we could also see the approach as avoiding undue path splits).

### 2031 4.3.2.2 Different Forkless Encodings

2032 Different forkless encodings are possible. We considered a few different encodings for  
 2033 arbitrary data faults in Table 4.1, leveraging various operators:

- 2034 – the inlined if-the-else operator (Ite),
- 2035 – the multiplication operator  $*$  (Mul),
- 2036 – the bitwise and operator  $\&$  (And),
- 2037 – the bitwise xor operator  $\oplus$  (Xor).

Table 4.1: Forkless encoding variants for arbitrary data faults

Fault model	Encoding
Original	$x := \text{expr}$
Ite	$x := \text{ite } \text{fault\_here} ? \text{fault\_value} : \text{expr}$
Mul	$x := \text{expr} + \text{fault\_here} * \text{fault\_value}$
And	$x := \text{expr} + (-\text{fault\_here}) \& \text{fault\_value}$
Xor	$x := \text{expr} \oplus ((-\text{fault\_here}) \& \text{fault\_value})$

2038 We use  $(-\text{fault\_here})$  as a shortcut for a signed extension of the boolean variable  
 2039  $\text{fault\_here}$  and then take its opposite. When the fault is active and  $\text{fault\_here} :=$   
 2040  $0b1$ , we consider  $(-\text{fault\_here})$  to be  $0xffffffff$  (only ones in the binary representation,  
 2041 equivalent to  $-1$  for signed integers), serving as a mask for the fault effect. The opposite  
 2042 of 0 remains 0.

2043 Experiments (see Appendix B.5) show that different solvers slightly favor different  
 2044 encodings in terms of performance. For the solver used (see Section 5.2.2.4), the *ite*  
 2045 operator proves to be the most efficient. Hence, we chose to mainly use the *ite* ex-  
 2046 pression operator, an inlined form of if-then-else at the expression level that does not  
 2047 induce forks in the analysis.

### 2048 4.3.2.3 Supported Fault Models

2049 We mainly illustrated the forkless fault encoding with arbitrary data faults, but other  
 2050 fault models are supported, summarized as pseudo code transformation in Tables 4.2.

2051 **Link With Adversarial Transitions.** In general, our forkless fault models are implemen-  
 2052 tations of replacing adversarial transitions (see Section 4.2.2.2), hijacking a legitimate  
 2053 instruction with a faulty behavior. By using in-between faults, an algorithm would  
 2054 require to consider faulting each possible memory cell and register between each in-  
 2055 struction. This is much more computationally heavy than replacing instructions with

Table 4.2: Forkless encodings for various fault models

Fault model	original instruction	Forkless encoding
Arbitrary data	$x := expr$	$x := ite\ fault\_here ? fault\_value : expr$
Variable reset	$x := expr$	$x := ite\ fault\_here ? 0x00000000 : expr$
Variable set	$x := expr$	$x := ite\ fault\_here ? 0xffffffff : expr$
Bit-flip	$x := expr$	$x := ite\ fault\_here ?$ $(expr\ xor\ 1 \ll\ fault\_value) : expr$
Test inversion	$if\ cdt$ $\quad then\ goto\ addr_1$ $\quad else\ goto\ addr_2$	$if\ (ite\ fault\_here ? \neg cdt : cdt)$ $\quad then\ goto\ addr_1$ $\quad else\ goto\ addr_2$
Instruction skip	$x := expr$	$x := ite\ fault\_here ? x : expr$
	$jump\ addr$	$if\ fault\_here\ then\ jump\ next$ $\quad else\ jump\ addr$

2056 predefined behavior expressing specific fault models. We detail, for each fault model,  
2057 the link with adversarial transitions (examples of fault models are detailed in Section  
2058 4.2.2.3).

2059 **Arbitrary Data Fault.** This fault model only targets assignment operations, with  
2060 an arbitrary effect, meaning *fault\_value* is a fresh variable the solver will propose  
2061 an interesting value for. Arbitrary data faults impact the data memory (load and  
2062 store operations) and registers. Arbitrary data faults implement adversarial transitions  
2063 described in Section 4.2.2.2.

2064 **Reset Fault.** We can use the forkless encoding to reset a variable instead of the legit-  
2065 imate value update. Here, the effect is simply to propagate the value 0. Reset faults  
2066 impact the data memory (load and store operations) and registers, as illustrated in  
2067 Section 4.2.2.3.

2068 **Set Fault.** Similarly to reset faults, the forkless encoding can instantiate a set fault by  
2069 propagating the value 0xffffffff (all bits at one in the binary representation).

2070 **Bit-flip Fault.** We encode a bit-flip as a xor of the original expression with a mask  
2071 containing only one set bit at a place determined by *fault\_value*. Bit-flip faults impact  
2072 the data memory (load and store operations) and registers, as described in Section  
2073 4.2.2.2.

2074 **Test Inversion Fault.** We consider the test *cdt* of a conditional jump as a data and  
2075 fault it with our forkless encoding. The effect of the fault is to take the negation of the  
2076 condition,  $\neg cdt$ , hence inverting the test. We write *goto addr<sub>1</sub>* and *goto addr<sub>2</sub>* to repre-  
2077 sent the execution continuing to different instructions depending on the branch taken.  
2078 Here we chose a compromise between the first adversarial transition proposed (cor-  
2079 rupting the condition beforehand) and the second (inverting the addresses) discussed  
2080 in Section 4.2.2.3. The first one could be confused with a data fault that, interestingly,  
2081 can have a test inversion effect. With the expressive grammar we have in practice, we  
2082 can express the negation of the condition inside the if, which is equivalent to inverting  
2083 jump addresses.

2084 **Instruction Skip Fault.** We can extend the forkless encoding to support the instruction  
2085 skip fault model by also considering jump addresses as data. With this representation,  
2086 we perform instruction skip as NOPs, and not as a corruption in the program counter.

2087 All categories of instructions have to be faulted to simulate their skipping. We take  
2088 our usual data fault encoding for assignments, with the old value of the left-hand side  
2089 as the fault effect. We transform jumps into conditional jumps where, if the fault is  
2090 active, the execution goes to the next address in the memory. More detail about the  
2091 implementation of the instruction skip fault model across instructions are available in  
2092 Section 5.3.3.3.

2093 **Other Fault Models.** The list of fault models above does not represent exhaustively  
2094 what can be encoded with a forkless encoding. It showcases the most well-known ones  
2095 and can be extended to other fault models. We discuss encoding other forkless fault  
2096 models in Section 4.5.1.

#### 2097 4.3.2.4 Forkless Faults in the Analysis

2098 For each instruction that could be targeted by the attacker, a forkless fault is injected  
2099 by the analysis. This allows to consider all possible placings simultaneously. Only later  
2100 will it be decided which faults are interesting and can lead to a successful attack path.  
2101 It also enables *a multiple fault analysis* by choosing how many faults are active at the  
2102 same time in a path, without any extra path created.

#### 2103 4.3.2.5 Complexity Trade-off

2104 While a forkless encoding indeed allows a significant path reduction compared to forking  
2105 approaches, the corresponding path predicates are more complicated than standard  
2106 path predicates, as they involve lots of extra-symbolic variables for deciding whether  
2107 the faults occur and emulating their effect. This is experimentally verified in Chapter 6.  
2108 We show later in Section 4.4 how to reduce these extra variables through two dedicated  
2109 optimizations.

2110 Also, note that we used the expressivity of the underlying SMT theory to encode  
2111 the possibility of faults and their effects. A more powerful, and hence more complex  
2112 to solve, theory is not needed here.

### 2113 4.3.3 Fault Injection Algorithm

2114 We present now how we use the forkless fault encoding to modify a symbolic execu-  
2115 tion algorithm into adversarial symbolic execution, an algorithm verifying adversarial  
2116 reachability.

#### 2117 4.3.3.1 Overview

2118 In the same way that we extended a transition system with adversarial transitions to go  
2119 from expressing standard reachability to adversarial reachability, we extend symbolic  
2120 execution, a correct and k-complete analysis technique for standard reachability, with a  
2121 fault encoding to build an analysis correct and k-complete for adversarial reachability.

2122 Our extension of the symbolic execution algorithm mainly consists of embedding our  
2123 fault encoding in the path predicate instead of the normal behavior of the program. Are  
2124 concerned the evaluation of assignments and conditional jumps. The various aspects  
2125 of symbolic execution not discussed in the following are kept as is. We named our  
2126 technique Forkless Adversarial Symbolic Execution, FASE for short, to express the  
2127 use of our forkless encoding (Forkless) to prevent path explosion in representing an  
2128 advanced attacker (Adversarial) in a reachability-oriented technique (SE). We also use

2129 the name Adversarial Symbolic Execution (ASE) to designate integrating faults into a  
 2130 symbolic execution algorithm, without explicitly mentioning whether they are forkless  
 2131 or forking faults.

### 2132 4.3.3.2 Building the Adversarial Path Predicate

2133 We start by considering data faults. FASE requires modifications to the evaluation of  
 2134 assignments and conditional jumps.

2135 **Fault Counter.** We keep track of the number of active faults with a fault counter  $nb_f$ ,  
 2136 containing the sum of all boolean activation values, extended to integers.

2137 **Adversarial Reachability Evaluation.** The main algorithm for adversarial symbolic ex-  
 2138 ecution (Algorithm 4.1) changes compared to standard SE (Algorithm 3.1) to include  
 2139 the impact of the attacker. The function `GetAdversarialPaths` (line 1) is the adver-  
 2140 sarial counterpart of `GetPaths`, enumerating all possible adversarial paths of bounded  
 2141 depth. The function `GetPredicate` (line 3) then computes the adversarial path pred-  
 2142 icate for each adversarial path. The adversarial reachability of a code location  $l$  also  
 2143 ensures the attacker has not exceeded their fault budget ( $nb_f \leq m_{\mathcal{A}}$ ) line 4.

---

#### Algorithm 4.1: Adversarial symbolic execution algorithm

---

**Input:** a program  $P$ , a bound  $k$ , a target location  $l$ , a maximal number of  
 faults  $m_{\mathcal{A}}$

**Data:** fault counter  $nb_f$

**Output:** Boolean value indicating whether  $l$  can be reached within  $k$  steps.

```

1 for path  $\pi$  in GetAdversarialPaths( $k$ ) do
2   if  $\pi$  reaches  $l$  then
3      $\Phi := \text{GetPredicate}(\pi)$ 
4     if  $\Phi \wedge (nb_f \leq m_{\mathcal{A}})$  is satisfiable then
5       return true
6     end
7   end
8 end
9 return false
```

---

2144 **Adversarial Assignment Evaluation.** The assignment evaluation, originally described  
 2145 in Algorithm 3.2 for standard SE, is transformed as illustrated in Algorithm 4.2. We  
 2146 consider an assignment of the form ‘a variable  $x$  receives an expression  $expr$ ’, written  
 2147  $x := expr$ . The assignment evaluation process embeds a wrapper, `FaultEncoding`,  
 2148 encoding the fault in a forkless manner. It involves the declaration of fresh symbolic  
 2149 variables for fault activation and fault effects – hence the update of the path predicate  
 2150  $\Phi$ . The fault counter  $nb_f$  is updated by adding to it an extension of the boolean  
 2151 activation variable  $nb_f := nb_f + \text{fault\_here}$ . A new expression containing a fault  
 2152 injection,  $expr'$ , is computed. The `eval_assign` function returns the updated path  
 2153 predicate  $\Phi$ , augmented with the assignment of  $x$  to the faulted expression.

2154 **Adversarial Conditional Jump Evaluation.** The conditional jump evaluation, originally  
 2155 described in Algorithm 3.3 for SE, is transformed as illustrated in Algorithm 4.3. We  
 2156 consider a conditional jump of the form ‘if a condition  $cdt$  evaluates to true, then  
 2157 the branch continuing at address  $addr_t$  is taken, otherwise, the branch at  $addr_e$  is

**Algorithm 4.2:** Forkless assignment evaluation for data faults**Input:** path predicate  $\Phi$ , assignment instruction  $x := expr$ **Data:** fault counter  $nb_f$ **Output:** updated  $\Phi$ ,  $nb_f$  updated in place

---

```

1 Function eval_assign( $\Phi, x, expr$ ) is
2   |    $\Phi', expr', nb_f := \text{FaultEncoding}(\Phi, expr, nb_f)$ 
3   |   return  $\Phi' \wedge (x \triangleq expr')$ 
4 end

```

---

2158 taken', written *if*  $cdt$   $addr_t$  *else*  $addr_e$ . Checking the feasibility of each branch ( $cdt$   
2159 and  $\neg cdt$ ) also includes making sure the attacker does not exceed their fault budget  
2160  $m_A$  to explore it. Hence the number of faults check ( $nb_f \leq m_A$ ) is appended to each  
2161 satisfiability query. The number of faults check can be performed at different places.  
2162 We found the best trade-off for forkless faults is to perform it at conditional jump  
2163 evaluation, as checking it at the end of a path often involves exploring many unfeasible  
2164 faulty paths.

**Algorithm 4.3:** Forkless conditional jump evaluation for data faults**Input:** path predicate  $\Phi$ , conditional jump instruction *if*  $cdt$   $addr_t$  *else*  $addr_e$ **Data:** fault counter  $nb_f$ , maximal number of faults  $max_f$ , worklist  $WL$ **Output:**  $WL$  updated in place

---

```

1 Function eval_conditional_jump( $\Phi, cdt, addr_t, addr_e$ ) is
2   |   if  $\Phi \wedge cdt \wedge (nb_f \leq max_f)$  is satisfiable then
3   |   |   Add ( $\Phi \wedge cdt, addr_t$ ) to  $WL$ 
4   |   end
5   |   /* Idem for else branch ( $\neg cdt$ )                                     */
6 end

```

---

2165 **Test Inversion Case.** A test inversion fault happens when considering a conditional  
2166 jump, leaving the standard SE process for assignments. We show in Algorithm 4.4 an  
2167 expression containing the activation of the fault and its effect is computed by the func-  
2168 tion `FaultEncoding` similarly to data faults. The resulting faulty condition,  $fault\_cdt$   
2169 is used in the rest of the conditional jump evaluation as the condition.

**Algorithm 4.4:** Forkless conditional jump evaluation for test inversions**Input:** path predicate  $\Phi$ , conditional jump instruction *if*  $cdt$   $addr_t$  *else*  $addr_e$ **Data:** fault counter  $nb_f$ , maximal number of faults  $max_f$ , worklist  $WL$ **Output:**  $WL$  and  $nb_f$  updated in place

---

```

1 Function eval_conditional_jump( $\Phi, cdt, addr_t, addr_e$ ) is
2   |    $\Phi', fault\_cdt, nb_f := \text{FaultEncoding}(\Phi, cdt, nb_f)$ 
3   |   if  $\Phi' \wedge fault\_cdt \wedge (nb_f \leq max_f)$  is satisfiable then
4   |   |   Add ( $\Phi' \wedge fault\_cdt, addr_t$ ) to  $WL$ 
5   |   end
6   |   /* Idem for else branch ( $\neg fault\_cdt$ )                               */
7 end

```

---

2170 **Instruction Skip Case.** This fault model changes each instruction handled by the SE  
 2171 to include the possibility of it having no effect and going to the next instruction in  
 2172 memory. It is hence implementation specific and we refer the reader to implementation  
 2173 details provided in Section 5.3.3.3. In particular for assignment evaluation, the process  
 2174 presented in Algorithm 4.2 is used.

2175 **Implementation Details.** We present our implementation of those algorithmic mod-  
 2176 ifications to a standard SE engine in Section 5.3 for binary-level analysis inside the  
 2177 BINSEC tool.

### 2178 4.3.3.3 Properties

2179 We now consider the properties of the FASE algorithm.

2180 **Proposition 4.2.** *The FASE algorithm is correct and  $k$ -complete for adversarial reach-*  
 2181 *ability.*

2182 *Proof.* If our algorithm finds an adversarial path reaching the target location  $l$ , by  
 2183 providing specific input values and a fault sequence, then an attacker executing the  
 2184 program with the provided inputs and performing the proposed faults will reach its  
 2185 goal. Our algorithm is based on symbolic execution with bounded path depth and  
 2186 explores all possible adversarial paths according to the considered attacker model,  
 2187 hence its  $k$ -completeness for adversarial reachability.  $\square$

2188 FASE will propose a set of adversarial paths if there are some, it is not guaranteed  
 2189 to provide all possible adversarial paths with respect to fault placements.

2190 **Definition 4.5** (Path). *We define a path by its control flow, that is to say, by the*  
 2191 *sequence of branch choices and jump addresses leading to the target location.*

2192 **Definition 4.6** (Equivalent Adversarial Paths). *We define an adversarial path equiva-*  
 2193 *lence class by the set of adversarial paths that follow the same control flow.*

2194 **Proposition 4.3.** *FASE will only report one adversarial path per equivalence class.*

2195 *Proof.* FASE injects faults in a symbolic way, that is to say, they do not have concrete  
 2196 placements during the path exploration performed by the SE analysis. Only at the end  
 2197 of one path reaching the target location will the analysis ask an SMT solver for the  
 2198 satisfiability of the adversarial path and a model for symbolic values. Hence, FASE  
 2199 provides only one model (sequence for faults) for each path reaching the attacker's  
 2200 goal.  $\square$

2201 *Remarque 4.1.* In a multiple fault analysis, FASE is not guaranteed to output minimal  
 2202 adversarial paths in the sense that there may be faults present that could be removed  
 2203 while the path still reaches the attacker's goal.

2204 **Proposition 4.4** (Tightness of FASE). *Consider a single path with no branching instruc-*  
 2205 *tion and a code location to reach at the end, together with  $f$  possible fault locations and*  
 2206 *a maximum of  $m_A$  faults. Then an SE injecting forking faults yields up to  $C_{m_A}^f$ <sup>1</sup> paths*  
 2207 *to analyze, and as many queries to send to the solver. In the same scenario, FASE*  
 2208 *(with forkless faults) will analyze only the original path, and send a single query to the*  
 2209 *solver.*

<sup>1</sup>Remind that  $k$  among  $n$ , written  $C_k^n$  is  $C_k^n = \binom{n}{k} = \frac{n!}{k!(n-k)!}$

2210 *Remarque 4.2.* The Forkless encoding increases query complexity, as described in Sec-  
2211 tion 4.3.2.5 and experimentally verified in Section 6.3.2. We present in the remainder  
2212 of this chapter two mitigation techniques.

## 2213 4.4 Optimizations

2214 As previously mentioned, FASE introduces many new symbolic variables which increase  
2215 the burden of SMT solving in the analysis. FASE is based on a trade-off, no path  
2216 explosion at the price of more complex queries. Our goal is to reduce this query  
2217 complexity in an effort to get the best of both worlds.

2218 In this section, we propose two optimizations, Early Detection of fault Saturation  
2219 (EDS) and Injection On Demand (IOD), that aim to produce queries with fewer faulty  
2220 expressions in them to alleviate the solver’s efforts, while remaining correct and k-  
2221 complete. The evaluation of each optimization’s effectiveness is performed in Section  
2222 6.3.2.

### 2223 4.4.1 Early Detection of Fault Saturation (EDS)

2224 Our first optimization is called Early Detection of Fault Saturation, or EDS for short.  
2225 The goal of EDS is to *stop fault injection as soon as possible*. The idea is to detect  
2226 when the attacker has necessarily spent all their fault budget in the currently explored  
2227 adversarial path, meaning they can’t inject more in the rest of that path. Fewer faults  
2228 injected in a path translates to fewer faulty expressions in the queries, hence, a reduced  
2229 query complexity.

#### 2230 4.4.1.1 Supported Fault Models

2231 The saturation check itself can be performed anywhere, at the price of more queries.  
2232 We discuss here how it can be applied to our various fault models.

2233 **Data Faults.** With forkless data faults in mind, we selected conditional jump eval-  
2234 uation as the saturation check location, taking advantage of the branch feasibility  
2235 computation already performed. Checking it at each fault location wouldn’t provide  
2236 more information, and only checking at the end would be useless.

2237 **Test Inversion Faults.** For this fault model, injection locations are already conditional  
2238 jump evaluation. EDS can be applied straightforwardly, like for data faults. We do not  
2239 believe EDS would provide much performance gain as the number of fault locations is  
2240 very limited and the effect of test inversion fault does not add a new symbolic variable  
2241 to be expressed, it is simply the opposite of the condition. Hence it would add queries  
2242 and computation to try and mitigate a limited added query complexity. For those  
2243 reasons, we believe that EDS could be used with test inversion faults, but with limited  
2244 benefits. It has not been implemented in this thesis.

2245 **Instruction Skip Faults.** We propose again to apply EDS at conditional jump evalu-  
2246 ation, to take advantage of the information of branch feasibility. We believe EDS can  
2247 be implemented for instruction skips. However, we made the choice in this thesis to  
2248 implement for instruction skip faults only the best performing optimization, Injection  
2249 On Demand, detailed in the next section. Section 6.3.2 shows the experimental results  
2250 for data fault that we believe can be extended to instruction skip since assignments

2251 are the most prevalent instructions in assembly. Hence, EDS for instruction skip faults  
2252 has not been implemented in this thesis.

#### 2253 4.4.1.2 Algorithm

2254 EDS modifications to FASE concern only the conditional jump evaluation, illustrated  
2255 in Algorithm 4.5.

---

#### Algorithm 4.5: FASE-EDS conditional jump evaluation

---

**Input:** path predicate  $\Phi$ , conditional jump instruction  
*if cdt then addr<sub>t</sub> else addr<sub>e</sub>*  
**Data:** fault counter  $nb_f$ , maximal number of faults  $max_f$ , worklist  $WL$   
**Output:**  $WL$  updated in place

```

1 Function eval_conditional_jump_EDS( $\Phi$ ,  $cdt$ ,  $addr_t$ ,  $addr_e$ ) is
2   if  $\Phi \wedge cdt \wedge (nb_f < max_f)$  is satisfiable then
3     | Add ( $\Phi \wedge cdt$ ,  $addr_t$ ) to  $WL$ 
4   else if  $\Phi \wedge cdt \wedge (nb_f == max_f)$  is satisfiable then
5     | Stop injection in this path
6     | Add ( $\Phi \wedge cdt$ ,  $addr_t$ ) to  $WL$ 
7   end
   /* Idem for else branch ( $\neg cdt$ )                               */
8 end

```

---

2256 Instead of checking whether a branch can be explored without exceeding the max-  
2257 imum number of faults ( $nb_f \leq max_f$ ), we double the check:

- 2258 1. First we check whether the branch can be explored with strictly fewer faults than  
2259 allowed ( $nb_f < max_f$ ). If the query is satisfiable, the analysis continues down  
2260 that branch as usual;
- 2261 2. If not satisfiable, we check whether the branch is feasible with exactly the maximal  
2262 number of faults allowed ( $nb_f == max_f$ ). If not, the branch is infeasible and we  
2263 stop as usual. Yet, if it is feasible, the saturation detection is triggered and we  
2264 know that we have spent all allowed faults. We can thus continue the exploration  
2265 *without injecting any new fault* in the corresponding search sub-tree, leading to  
2266 simpler subsequent queries.

#### 2267 4.4.1.3 Properties

2268 We now consider the properties of FASE augmented with EDS, which we refer to as  
2269 FASE-EDS.

2270 **Proposition 4.5.** *FASE-EDS is correct and  $k$ -complete for the adversarial reachability*  
2271 *problem.*

2272 *Proof.* FASE-EDS remains correct as it does not modify the path predicate compu-  
2273 tation, and it remains  $k$ -complete as it only prunes fault injections that are actually  
2274 infeasible – and would have been proven so by the solver, later in the solving pro-  
2275 cess. □

## 2276 4.4.2 Injection on Demand (IOD)

2277 Our second optimization is called Injection On Demand, or IOD for short. The goal of  
2278 IOD is to *inject fault on demand*, or as late as possible. The idea is to explore a path  
2279 without injecting faults, then, only when they become needed to continue the explo-  
2280 ration, they are added retrospectively to the path predicate. With this optimization,  
2281 there are most of the time fewer faulty expressions in queries than what the attacker  
2282 could do, since some injections have been delayed, hence reducing query complexity.

### 2283 4.4.2.1 Supported Fault Models

2284 We discuss here how it can be applied to our various fault models.

2285 **Data Faults.** This optimization was designed with forkless data faults in mind. As  
2286 they add complexity to the path predicate, it is particularly interesting to delay the  
2287 injection until it is truly needed. This is evaluated at conditional jump locations to  
2288 ensure all branches possible with faults are indeed explored.

2289 **Test Inversion Faults.** Again, test inversion fault injection locations are already condi-  
2290 tional jump evaluation. We do not believe IOD would provide much performance gain  
2291 as the number of fault locations is very limited, only to conditional jump instructions,  
2292 while data assignments are predominant, especially in assembly. Furthermore, the ef-  
2293 fect of test inversion fault does not add a new symbolic variable to be expressed, it is  
2294 simply the opposite of the condition. Hence it would add queries and computation to  
2295 try and mitigate a limited added query complexity. For those reasons, we believe that  
2296 IOD could be used with test inversion faults, but with limited benefits. It has not been  
2297 implemented in this thesis.

2298 **Instruction Skip Faults.** This optimization works well with the data-related part of  
2299 instruction skip, in a similar manner as for data faults. The algorithm retrospectively  
2300 adding faults needs to be extended to take into account the other types of instructions.  
2301 This mechanism is detailed in Section 5.3.5, at the implementation phase, when the  
2302 types of instructions considered have been described.

### 2303 4.4.2.2 Algorithm

2304 IOD modifies FASE assignment evaluation as illustrated in Algorithm 4.6 and condi-  
2305 tional jump as shown in Algorithm 4.7.

2306 **Dual Path Predicates.** In order to perform the retroactive injection of faults when  
2307 they are needed, we create a second path predicate.

- 2308 – The first path predicate  $\Phi$  is our main path predicate, upon which the analysis  
2309 reasons and builds queries.
- 2310 – The second path predicate  $\Phi_F$ , or faulty path predicate, is our backup path  
2311 predicate containing all the faults as the exploration progresses. It is used only  
2312 when the analysis detects it needs more faults. It is then switched with the main  
2313 path predicate, in effect, injecting faults retrospectively.

2314 **Assignment Process.** The assignment evaluation is duplicated as shown in Algorithm  
2315 4.6:

- 2316 – The normal symbolic assignment, with the original right-end-side expression *expr*,  
2317 is performed in  $\Phi$ ;

- 2318 – The fault injection is performed in  $\Phi_F$ , which is updated with the fault encoding  
 2319 of the assignment,  $expr'$ .

---

**Algorithm 4.6:** FASE-IOD assignment evaluation
 

---

**Input:** path predicate  $\Phi$ , faulted path predicate  $\Phi_F$ , assignment instruction  
 $x := expr$   
**Data:** fault counter  $nb_f$   
**Output:** Updated  $\Phi$ ,  $\Phi_F$ ,  $nb_f$  updated in place

```

1 Function eval_assign_IOD( $\Phi$ ,  $\Phi_F$ ,  $cdt$ ,  $x$ ,  $expr$ ) is
2   |  $\Phi_F$ ,  $expr'$ ,  $nb_f := \text{FaultEncoding}(\Phi_F, expr, nb_f)$ 
3   | return ( $\Phi \wedge (x \triangleq expr)$ ,  $\Phi_F \wedge (x \triangleq expr')$ )
4 end

```

---

2320 **Conditional Jump Process.** The conditional jump evaluation is extended to test  
 2321 whether more faults are needed. Queries are built according to the following rules:

- 2322 1. The first branch feasibility check is built with the main, simpler, path predicate  
 2323  $\Phi$ , encompassing the least number of faults. We continue this way as long as we  
 2324 can, meaning we rely on standard reachability as much as we can;
- 2325 2. When encountering a branch infeasible with  $\Phi$ , we then check whether this branch  
 2326 is feasible with all the possible faults seen so far, i.e. using  $\Phi_F$ . If not, that is a  
 2327 stop, otherwise, we know that  $\Phi$  does not encompass enough faults to go further.  
 2328 We then replace  $\Phi$  by  $\Phi_F$  (called a *switch*) at this stage, and thus continue with  
 2329 strictly more faults. The switch is straightforward as  $\Phi_F$  and  $\Phi$  only differ on  
 2330 fault injections. Then again, the new  $\Phi$  will not accumulate any fault (until a  
 2331 new switch) while  $\Phi_F$  continues accumulating all possible faults.

---

**Algorithm 4.7:** FASE-IOD conditional jump evaluation
 

---

**Input:** path predicate  $\Phi$ , conditional jump instruction *if cdt then  $l_t$  else  $l_e$*   
**Data:** fault counter  $nb_f$ , maximal number of faults  $max_f$ , under  
 approximation counter *under\_counter*, worklist  $WL$   
**Output:**  $WL$  updated in place

```

1 Function eval_conditional_jump_IOD( $\Phi$ ,  $\Phi_F$ ,  $cdt$ ,  $l_t$ ,  $l_e$ ) is
2   | if  $\Phi \wedge cdt \wedge (nb_f \leq max_f)$  is satisfiable then
3     | Add ( $\Phi \wedge cdt$ ,  $\Phi_F \wedge cdt$ ,  $l_t$ ) to  $WL$ 
4   | else if under_counter  $\leq max_f$  then
5     | if  $\Phi_F \wedge cdt \wedge (nb_f \leq max_f)$  is satisfiable then
6       |  $\Phi := \Phi_F$ 
7       | under_counter := under_counter + 1
8       | Add ( $\Phi \wedge cdt$ ,  $\Phi_F \wedge cdt$ ,  $l_t$ ) to  $WL$ 
9     | end
10  | end
11  | /* Idem for else branch ( $\neg cdt$ ) */
11 end

```

---

2332 **Under-Approximation Counter.** As a bonus, the number of path predicate switches  
 2333 gives us an under-approximation *under\_counter* of the number of faults already needed

2334 in the path under analysis. We use it to stop the injection early, when at least  $nb_f$   
 2335 faults have been used.

### 2336 4.4.2.3 Properties

2337 We now consider the properties of FASE augmented with IOD, which we refer to as  
 2338 FASE-IOD.

2339 **Proposition 4.6.** *FASE-IOD is correct and  $k$ -complete for the adversarial reachability*  
 2340 *problem.*

2341 *Proof.* FASE-IOD explores the same feasible paths as FASE, hence preserving its prop-  
 2342 erties.  $\square$

### 2343 4.4.3 Combination of Optimizations

2344 Since our two optimizations approach the problem of query complexity reduction from  
 2345 different angles, they can be combined, as illustrated in Algorithm 4.8. Taking FASE-  
 2346 IOD as a basis, saturation detection is added for the faulted path predicate  $\Phi_F$  queries  
 2347 at the conditional branch evaluation. If the saturation is detected ( $(nb_f \not\leq max_f) \wedge$   
 2348  $(nb_f == max_f)$ ), the main path predicate switches to  $\Phi_F$  but for the last time along  
 2349 this path and  $\Phi_F$  is not updated and queried upon anymore. This effectively stops  
 2350 fault injection, thus reducing subsequent query complexity.

---

**Algorithm 4.8:** FASE-IOD and FASE-EDS combination, conditional jump evaluation

---

**Input:** path predicate  $\Phi$ , faulty path predicate  $\Phi_F$ , conditional jump instruction *if cdt then addr<sub>t</sub> else addr<sub>e</sub>*  
**Data:** fault counter  $nb_f$ , maximal number of faults  $max_f$ , under approximation counter *under\_counter*, worklist *WL*  
**Output:** *WL* updated in place

```

1 Function eval_conditional_jump_EDS_IOD( $\Phi$ ,  $\Phi_F$ , cdt, addrt, addre) is
2   if  $\Phi \wedge cdt \wedge (nb_f \leq max_f)$  is satisfiable then
3     | Add ( $\Phi \wedge cdt$ ,  $\Phi_F \wedge cdt$ , addrt) to WL
4   else if under_counter  $\leq max_f$  then
5     | if  $\Phi_F \wedge cdt \wedge (nb_f < max_f)$  is satisfiable then
6       |  $\Phi := \Phi_F$ 
7       | under_counter := under_counter + 1
8       | Add ( $\Phi \wedge cdt$ ,  $\Phi_F \wedge cdt$ , addrt) to WL
9     | else if  $\Phi_F \wedge cdt \wedge (nb_f == max_f)$  is satisfiable then
10    |  $\Phi := \Phi_F$ 
11    | Stop switches and  $\Phi'$  update and queries
12    | Add ( $\Phi \wedge cdt$ ,  $\Phi_F \wedge cdt$ , addrt) to WL
13    | end
14  end
15 end
/* Idem for else branch ( $\neg cdt$ ) */

```

---

2351 **Proposition 4.7.** *The combination of FASE-EDS and FASE-IOD is correct and  $k$ -*  
 2352 *complete for the adversarial reachability problem.*

2353 *Proof.* This combination also explores all possible paths for the considered attacker  
 2354 models, like FASE, hence preserving its properties.  $\square$

## 2355 4.5 Discussion

2356 In this section, we discuss the applicability of forkless fault encoding, from formal-  
 2357 ization to verification algorithm and its advantages for multi-fault analysis. We also  
 2358 consider using the forkless fault encoding to verify different properties and with differ-  
 2359 ent formal techniques. General techniques used in prior work in SWIFI have already  
 2360 been discussed in Section 3.4.

### 2361 4.5.1 Fault Model Support - Formalization VS Algorithm

2362 We discuss in this section what types of faults are supported by our formalization and  
 2363 those supported by our algorithm, FASE.

2364 Our formalization is an extension of a transition system with adversarial transi-  
 2365 tions to represent the impact of an advanced attacker on a program. Our algorithmic  
 2366 approach based upon it, FASE, currently supports:

- 2367 – Data faults involving a single instruction,
- 2368 – Simple control-flow faults (test inversion and instruction skip).

2369 While our formalization supports the following types of faults, our algorithm ap-  
 2370 proach does not.

- 2371 – Advanced data faults, whose effects span over multiple instructions. We believe  
 2372 that once characterized, those could be implemented by extending the program  
 2373 state with some form of counters to link the different effects together;
- 2374 – Advanced control-flow faults such as arbitrary jumps. We have no algorithm more  
 2375 efficient in mind than a simple, and explosive, enumeration of possible effects for  
 2376 this class of fault models;
- 2377 – Instruction corruption permanently changes an instruction, while we modify com-  
 2378 putation results. It is related to self-modification, which is a hard problem for  
 2379 symbolic execution. Supporting instruction corruption would require modeling  
 2380 instruction decoding and an efficient algorithm to avoid state explosion.

2381 FASE can, in theory, support, or be extended to support, a vast number of different  
 2382 fault models. However, we do not guarantee our algorithm will be efficient for them,  
 2383 since they may involve new symbolic variables and an increased query complexity.

2384 **Representing Micro-Architectural Attacks.** While Rowhammer’s bit-flip effects are  
 2385 quite straightforward to model, modeling Spectre-type attacks is more difficult.

- 2386 – For Spectre v1 (BTB, branch-target-buffer), the effect itself of mispredicting a  
 2387 branch can be modeled by a test inversion, however, the duration of the effect is  
 2388 limited in time and a rollback mechanism needs to be devised;
- 2389 – In an LVI attack, transient execution is used to inject attacker-controlled values  
 2390 in the victim process. The effect of the attack is an arbitrary data fault, but with  
 2391 a limited effect duration to take into account;
- 2392 – In Spectre v4 (STL, store-to-load), the processor speculates if a memory load  
 2393 depends on previous stores that are buffered, or independent and the value in

2394 memory is the correct one. This can be seen as a restricted form of an LVI  
2395 attack, where the attack effect is a data corruption dependent on previous values  
2396 at the considered memory address. If those addresses are symbolic, the analysis  
2397 can suffer a state explosion. Again, the effect is of limited duration.

2398 While we believe FASE could be adapted to model temporary effects and state rollback,  
2399 the application of our optimizations is not straightforward.

2400 Other types of attacks exploit micro-architectural features and would require ex-  
2401 tending the program state with micro-architectural components [LBD<sup>+</sup>18, LBDPP19,  
2402 TAC<sup>+</sup>22] to represent:

- 2403 – permanent data faults,
- 2404 – faults in the instruction prefetch buffer,
- 2405 – faults in hidden registers,
- 2406 – faults in the multiplication unit,
- 2407 – faults in the forwarding mechanism,
- 2408 – skip faults on variable-length instruction set [ACD<sup>+</sup>22] with miss-aligned instruc-  
2409 tions.

2410 While FASE is architecture-independent, extending the program state as such would  
2411 make the analysis architecture dependent, and would have to be adapted for each  
2412 architecture considered.

2413 **Probabilistic Faults.** Some techniques [DPdC<sup>+</sup>15], mostly simulation-based, imple-  
2414 ment probabilistic fault models to account for the plausibility of a fault, depending on  
2415 injection parameters. We believe our formalism could be extended to support proba-  
2416 bilistic adversarial transitions. However, our current algorithm cannot support those  
2417 faults without a path explosion, forking to consider each fault having a non-zero prob-  
2418 ability.

2419 **Related Work.** Fault model support of other SWIFI works is presented in Section  
2420 4.6.1.

## 2421 4.5.2 Forkless Faults and Multi-Fault Analysis

2422 We now discuss the interest of the forkless fault encoding for a multi-fault analysis.

2423 The goal of the forkless encoding is to be able to represent attackers able to per-  
2424 form multiple faults in one attack while preventing the path explosion faced by existing  
2425 techniques. We do this by embedding the possibility of a fault at each possible location  
2426 without forking the path each time, at the price of more complexity in the path predi-  
2427 cate, which makes for more complex SMT queries, the bottleneck of symbolic execution.  
2428 SWIFI is a trade-off between exploring more paths or having more complex queries.  
2429 We chose to explore the side of more complex queries. Experiments (see Chapter 6)  
2430 will show that our approach yields the best performance in most cases, even without  
2431 our optimizations reducing the number of injection points in queries.

2432 While our algorithm is abstraction-level independent, we implement it in Chapter 5  
2433 for a binary-level analysis, which is harder than source-level analyses (see Section 3.3).

## 2434 4.5.3 Forkless Encoding for Other Properties

2435 The forkless encoding can surely benefit other classes of properties to be achieved  
2436 by the attacker, especially those known to be supported by (extensions of) symbolic  
2437 execution, for example:

- 2438 – Trace properties that can be evaluated based on a transition system, such as
- 2439 use-after-free, could also be evaluated under attack when evaluated based on
- 2440 our adversarial transition system instead. This corresponds to considering an
- 2441 attack goal other than adversarial reachability, with a more complex oracle. In
- 2442 particular, an extension to other trace properties allows to consider complex
- 2443 multi-step attacks where a vulnerability is used to trigger the next one;
- 2444 – K-hyperreachability properties (secret leakage, privacy leakage, violation of constant-
- 2445 time, etc.) [DBR20] are often evaluated with relational symbolic execution, where
- 2446 two traces following the same path are computed and then compared. Faults can
- 2447 be added to that computation similarly to how we included them in symbolic
- 2448 execution. This could be used to extend attacker goals beyond reachability, in
- 2449 particular towards information leakage, or serve as a guide for multi-step attacks;
- 2450 – The recent robust reachability proposal [GFB21] and its quantified version [BG22]
- 2451 consider the control an attacker needs, or not, over each program parameter to
- 2452 trigger a bug. This work aims to mitigate symbolic execution false positives
- 2453 and find highly replicable bugs. This technique and ours could be combined to
- 2454 quantify the control an attacker needs over the injected fault. For instance, does
- 2455 the attacker have to corrupt a value to 0xdeadbeef as the adversarial symbolic
- 2456 execution says, or any random value would work? In particular, this could help
- 2457 an attacker decide what injection mean provides an adequate level of control over
- 2458 the injected fault to perform the attack. However, robust reachability requires a
- 2459 more expressive, and harder-to-solve, SMT theory than we use. We believe the
- 2460 combination of the two techniques would suffer scalability issues.

#### 2461 4.5.4 Forkless Encoding for Other Formal Methods

2462 While in this thesis, we focus on symbolic execution, we believe the main optimization  
 2463 ideas developed here can be used with other formal techniques, e.g. Bounded Model  
 2464 Checking, Abstract Interpretation or CEGAR model checking. Note that for each of  
 2465 them, fault injection may result either in path explosion or precision loss. Still, our  
 2466 forkless encoding should be able to help at least all approaches based to some extent  
 2467 on path unrolling.

2468 Previous works have explored non-forking injection techniques with abstract in-  
 2469 terpretation [LFBP21], deductive verification [MKP22] and bounded model checking  
 2470 [TAC<sup>+</sup>22].

2471 Lacombe *et al.* [LFBP21] use abstract interpretation, through the Frama-C [KKP<sup>+</sup>15]  
 2472 plugin Eva [BBY17]. They propose a form of forkless faults, xoring the result of a com-  
 2473 putation with a free variable, manually added to a C program. However, they rely on  
 2474 abstract interpretation only to filter injection points, not to find adversarial paths  
 2475 directly.

2476 Martin *et al.* [MKP22] use deductive verification at C level through the Frama-  
 2477 C plugin LTEST [BCDK14] with test inversion faults. They use an uninterpreted  
 2478 function to simulate the activation of the fault and its behavior. Only the contract  
 2479 of this function is written, leaving Frama-C to verify if the security property can be  
 2480 violated.

2481 Tollec *et al.* [TAC<sup>+</sup>22] leverage bounded model checking for fault injection analysis,  
 2482 except they reason at the RTL level, modeling the CPU micro-architecture. The RTL  
 2483 model takes the place of the program, whose input is the binary program analyzed,  
 2484 with possible faults on each wire on the RTL model. The analysis can be considered

2485 completely forkless as the possibility of faults and their effect are encoded directly into  
2486 the model of the bounded model checker, which never forks.

### 2487 **4.5.5 Forkless Encoding and Instrumentation**

2488 Several prior works use code-level instrumentation [MKP22] or llvm level instrumen-  
2489 tation [PMPD14, LFBP21, LHGD18] in order to leverage standard program analyzers  
2490 as is. The forkless encoding we propose can also be used this way, for more flexibility.  
2491 However, since using code-level instrumentation takes a step back from the symbolic  
2492 execution engine, not all optimizations still make sense. While we believe Early Detec-  
2493 tion of Fault Saturation could be adapted, Injection On Demand requires heavy path  
2494 predicate transformation that does not seem feasible in code instrumentation. Actu-  
2495 ally, we performed some experiments with Klee and C-level forkless instrumentation in  
2496 Section 6.5, and do observe significant performance improvement compared to a forking  
2497 instrumentation.

## 2498 **4.6 Related Work**

2499 In this section, we present the related work regarding fault model support and multi-  
2500 fault analysis. We also explore other fields considering attackers and other formalisms  
2501 that have been proposed.

### 2502 **4.6.1 Fault Model Support**

2503 We list in Table 4.3 the main SWIFI works and the fault models they support. Most  
2504 only support one or very few fault models that are typically simple data faults or  
2505 control-flow faults (test inversion and instruction skip). To our knowledge, only one  
2506 [TAC+22] includes micro-architecture details in their fault injection analysis which  
2507 allows them to reason about complex faults.

### 2508 **4.6.2 Multiple Fault Analysis**

2509 We now present how existing techniques support multi-fault in their analyses. Among  
2510 SWIFI techniques, half rely on the mutant approach [CCG13, RG14, GWJLL17, CDFG18,  
2511 GWJL20], which is not designed for multi-fault analysis has the number of mutants  
2512 explodes quickly. The other half relies on forking techniques [PMPD14, BBC+14,  
2513 BHE+19, LFBP21, Lan22] whose scalability in the number of faults considered is hin-  
2514 dered by path explosion. Very few works consider multi-faults [PMPD14, LFBP21,  
2515 MKP22, Lan22, GHHR23] (see Table 4.3).

2516 Potet *et al.* [PMPD14] only considers test inversion faults, which are much less  
2517 frequent than data faults. They use a preprocessing consisting of a basic block coloring  
2518 algorithm based on reachability in a graph to filter necessary fault locations, locations  
2519 where faults should be avoided and uncertain locations. That information is then  
2520 embedded in a meta-mutant which is analyzed by symbolic execution with KLEE  
2521 [CDE+08]. Hence, the only forking faults handled by the symbolic engine are uncertain  
2522 fault locations of a fault model having few in total.

2523 Lacombe *et al.* [LFBP21] use abstract interpretation, through the Frama-C [KKP+15]  
2524 plugin Eva [BBY17], to perform a data dependancy analysis reducing injection points

Table 4.3: Fault model support

Reference	Analysis level	Arbitrary data	Bit-flip	Other data fault	Test inversion	Instruction skip	Other	Multi-fault
Christofi [CCG13]	C	✗	✗	✓	✗	✗	✗	✗
LAZART [PMPD14, LFBP21]	llvm	✓	✗	✗	✓	✗	✗	✓
Rauzy [RG14]	Custom	✓	✗	✓	✗	✗	✗	✗
EFS [BBC+14]	Binary	✗	✗	✗	✗	✓	✗	✗
Given-Wilson [GWJLL17]	ISA	✗	✗	✓	✓	✓	✗	✗
Carré [CDFG18]	Binary	✗	✓	✗	✗	✗	✗	✗
RobustB [BHE+19]	Binary	✓	✗	✗	✗	✓	✗	✗
Given-Wilson [GWJL20]	Binary	✗	✓	✓	✗	✓	✗	✗
Lacombe [LFBP21]	C/llvm	✓	✗	✗	✗	✗	✗	✓
Martin [MKP22]	C	✗	✗	✗	✓	✗	✗	✓
Tollec [TAC+22]	RTL	✗	✗	✗	✗	✗	✓	✗
Lancia [Lan22]	ISA	✓	✓	✓	✓	✗	✗	✓
SAMVA [GHHR23]	ISA	✗	✗	✗	✗	✓	✗	✓
FASE	ISA	✓	✓	✓	✓	✓	✗	✓

2525 on C programs. The remaining injection points are then provided to a more recent ver-  
 2526 sion of Lazart [PMPD14], able to reason on data faults, to compute adversarial paths.  
 2527 Their fault injection pruning technique based on data-flow analysis is complementary  
 2528 to our own method – still, static analysis at binary level is known to be hard.

2529 Martin *et al.* [MKP22] use an uninterpreted function whose contract embeds the  
 2530 activation of test inversion faults and a counter to limit them to the attacker’s budget.  
 2531 This technique discharges the multi-fault reasoning to a solver, similar to our approach,  
 2532 except that we have a concrete fault encoding, allowing us to reason about and prune  
 2533 fault injection locations. Furthermore, they reason at the C level.

2534 Gicquel *et al.* [GHHR23] have developed an injection technique based on graph  
 2535 reasoning for the instruction skip fault model at the binary level. All possible faults  
 2536 are encoded by adding edges on the graph representing the program with basic blocks,  
 2537 each labeled as ‘neutral’, ‘skip’ or ‘execute’. Their algorithm then builds attack paths  
 2538 placing instructions skips of arbitrary width to cover all ‘skip’ nodes while avoiding  
 2539 all the ‘execute’ nodes. Here, increasing the attacker fault budget doesn’t change the  
 2540 extended graph, it only offers more freedom to the placing algorithm.

2541 In summary, those techniques are able to perform multiple fault analyses by:  
 2542 – using preprocessing filtering fault locations to limit the path explosion [PMPD14,  
 2543 LFBP21];  
 2544 – reasoning at a higher abstraction level than ISA, such as C [LFBP21, MKP22]  
 2545 or llvm [PMPD14];  
 2546 – fine-tuning their multi-fault technique to a single fault model [PMPD14, MKP22,  
 2547 GHHR23];  
 2548 – limiting multi-fault considerations to two faults in practice [LFBP21, MKP22,  
 2549 Lan22].

### 2550 4.6.3 The Attacker in Different Security Fields

2551 In this section, we overview how attackers are represented in various fields for security  
2552 analysis.

2553 **Protocol Verification.** It is common in the field of automated formal verification of  
2554 cryptographic protocols to consider models of attackers, typically extensions of the  
2555 “Dolev-Yao” model [DY83]. In the Dolev-Yao model, the attacker can see, intercept  
2556 and generate messages exchanged over the network. Cryptographic primitives are rep-  
2557 resented as abstract operators. The attacker only knows what has been previously  
2558 exchanged [AC04] to try and guess the secret key. Conversely, attackers can also be  
2559 represented by what they cannot do [BCL14].

2560 **Control-Flow Integrity.** In software security, control-flow integrity attacks have been  
2561 categorized by the capability an attacker needs [BCN<sup>+</sup>17] to perform an attack or  
2562 against which a countermeasure will hold. Those capabilities are of the form ‘write  
2563 something anywhere’, ‘write anything anywhere’, etc. These efforts have been restricted  
2564 to manual reasoning.

2565 **Bug Triaging.** With the advent of massive fuzzer usage, the number of bugs discovered  
2566 in programs becomes much higher than the capabilities of developers to fix them. A  
2567 rising area of research is bug triaging, where the severity and impact of bugs are eval-  
2568 uated to prioritize the most dangerous ones. Bugs are represented by the capabilities  
2569 [JGH<sup>+</sup>22] they provide to the attacker willing to exploit them, such as the number of  
2570 controlled bits.

2571 **High-Level Attacker.** Some works intend to model an actual attacker, a “hacker”  
2572 in action [CTB<sup>+</sup>17, ASA<sup>+</sup>15, BCR<sup>+</sup>19, CTB<sup>+</sup>19]. They aim to categorize high-level  
2573 attackers, understand what capabilities they have or not, and when and why they  
2574 favor one over another in the course of an attack. Those capabilities are of the form  
2575 ‘tamper with the execution environment’, ‘deobfuscate code’, ‘brute force attack’, etc.  
2576 While at a very high-level (workflow of the attack), these efforts are indeed relevant to  
2577 better understand which capabilities should be granted to a low-level formal model of  
2578 an attacker.

2579 **Precise Hardware Fault Models.** Determining the most appropriate fault models rep-  
2580 resenting hardware faults is an ongoing research question. Fault models [GMA22,  
2581 RBSG22] can be derived from experiments. Methodologies have been proposed [DPdC<sup>+</sup>15].  
2582 However, software fault models miss architecture-specific behavior related to control  
2583 signals in the pipeline, able to bypass countermeasures [LBD<sup>+</sup>18]. Faults on data are  
2584 independent of hardware, and faults on instructions depend on the Instruction Set  
2585 Architecture (ISA).

### 2586 4.6.4 Extending Existing Formalisms

2587 We highlight the main works that extended existing formalisms to include the capabil-  
2588 ities of an attacker.

2589 **Adversarial Logic.** Here, Vanegue [Van22] extends incorrectness logic [O’H19], propa-  
2590 gating accumulated errors to determine a bug exploitability. This bug-finding (under-  
2591 approximation) technique integrates an attacker from the side-channel and information  
2592 knowledge point of view. As our technique does not cover information leakage, it would  
2593 be interesting to try and combine both approaches in order to represent a more generic

2594 attacker, able to also seek knowledge.

2595 **Operational Semantic.** Given-Wilson et al. [GWL20] propose a formalization of fault  
2596 injection on a program using Turing machines. The attacker model considered can  
2597 inject active faults without consideration of the injection mean and is very generic in  
2598 supported faulty behaviors. This work is orthogonal to ours, however, to our knowledge,  
2599 no algorithm has been built for this formalization. Also, Fournet et al. [FR08] propose a  
2600 type system for program-level non-interference, taking into account an active adversary  
2601 modeled as adversarial components able to perform any action at certain steps of the  
2602 program.

## 2603 4.7 Conclusion

2604 **Summary.** In this Chapter, we started by describing how we modeled an advanced  
2605 attacker, an entity with attack capabilities, a limitation on the number that can be  
2606 used and an attack goal expressed as a program location to reach. We then showed  
2607 how to extend reachability to include adversarial transitions, building the new concept  
2608 of adversarial reachability. Finally, we extended the symbolic execution algorithm, tra-  
2609 ditionally verifying reachability, to include attacker actions as fault injections, building  
2610 adversarial symbolic execution (ASE). ASE is correct and k-complete for adversarial  
2611 reachability. The main limitation of existing SWIFI techniques is the path explosion  
2612 experienced. To prevent it, we proposed a new forkless fault encoding, that has the  
2613 drawback of increasing query complexity. We designed two optimizations dedicated to  
2614 reducing this added complexity.

2615 **Extension to Published Work.** Compared to the ESOP 2023 [DBP23] paper, we added  
2616 the instruction skip fault model and adapted our optimizations for it. We also proposed  
2617 several discussions, in particular regarding the expressivity of our adversarial transition  
2618 system.

2619 **Coming Next.** In the next chapter, we present our implementation of ASE into a  
2620 tool, BINSEC/ASE, to verify in practice the vulnerability of a program to an attacker  
2621 model. Its practical interest and effectiveness are investigated in Chapter 6.



2622 Chapter **5**

2623 The BINSEC/ASE Prototype

2624 **Contents**

---

2625	<b>5.1 Overview</b>	<b>69</b>
2626		
2627	<b>5.2 Background: the BINSEC Tool</b>	<b>69</b>
2628	5.2.1 BINSEC Presentation	70
2629	5.2.2 General Work-Flow	70
2630	5.2.2.1 Parameters	70
2631	5.2.2.2 Intermediate Representation: DBA	71
2632	5.2.2.3 Analysis Workflow	71
2633	5.2.2.4 SMT Solvers	71
2634	5.2.2.5 Output	73
2635	5.2.3 Summary	73
2636	<b>5.3 BINSEC/ASE Implementation</b>	<b>73</b>
2637	5.3.1 BINSEC/ASE Overview	73
2638	5.3.2 ASE Implementation	74
2639	5.3.2.1 Attacker Model Parameters and Goal	74
2640	5.3.2.2 Global Variables	74
2641	5.3.2.3 Assignments	74
2642	5.3.2.4 Conditional Jumps	75
2643	5.3.2.5 Exploration Directives	75
2644	5.3.2.6 BINSEC/ASE Output	76
2645	5.3.2.7 Statistics	76
2646	5.3.2.8 Implementation Details	76
2647	5.3.3 Forkless Fault Models	77
2648	5.3.3.1 Data Faults	77
2649	5.3.3.2 Test Inversion Faults	79
2650	5.3.3.3 Instruction Skip Faults	80
2651	5.3.4 Early Detection of Fault Saturation (EDS)	81
2652	5.3.5 Injection On Demand (IOD)	82

---

2653	5.3.6	Sub-fault Simplification . . . . .	82
2654	5.3.7	Forking Fault Models . . . . .	83
2655	5.3.7.1	Forking Data Faults . . . . .	83
2656	5.3.7.2	Forking Test Inversion Faults . . . . .	84
2657	5.3.7.3	Forking Instruction Skip Faults . . . . .	84
2658	5.3.8	Conclusion . . . . .	84
2659	<b>5.4</b>	<b>User Guide: a Methodology to Analyse a New Program . . . . .</b>	<b>86</b>
2660	5.4.1	Running Example . . . . .	86
2661	5.4.2	Analysis Goal . . . . .	88
2662	5.4.2.1	Security Properties . . . . .	88
2663	5.4.2.2	Attacker Model . . . . .	89
2664	5.4.3	Configuration . . . . .	93
2665	5.4.4	Reading BINSEC/ASE Output . . . . .	93
2666	5.4.5	Analysis Process . . . . .	96
2667	5.4.6	Summary . . . . .	96
2668	<b>5.5</b>	<b>Developer Guide: a Methodology to Add a New Fault Model . . . . .</b>	<b>97</b>
2669	5.5.1	Defining the New Fault Model . . . . .	97
2670	5.5.1.1	State-of-the-art Attack . . . . .	97
2671	5.5.1.2	Fault Model . . . . .	97
2672	5.5.2	Implementation . . . . .	98
2673	5.5.3	Dedicated Metrics . . . . .	99
2674	5.5.4	Testing . . . . .	99
2675	5.5.5	Summary . . . . .	99
2676	<b>5.6</b>	<b>Discussion . . . . .</b>	<b>100</b>
2677	5.6.1	BINSEC/ASE Limitations . . . . .	100
2678	5.6.1.1	Building on Top of an Existing Tool . . . . .	100
2679	5.6.1.2	Supported Fault Models . . . . .	100
2680	5.6.2	Faults on Intermediate Representation . . . . .	101
2681	5.6.3	Permanent VS Transient Faults . . . . .	101

---

2685 In this chapter, we present BINSEC/ASE, our software-implemented fault injection  
2686 tool, based on Adversarial Symbolic Execution (described in Chapter 4). After a  
2687 brief introduction explaining the goal and interest of our approach (Section 5.1), we  
2688 describe the BINSEC tool<sup>1</sup> (Section 5.2) and how we modified it to build BINSEC/ASE  
2689 (Section 5.3). Then, we provide a user guide detailing a methodology to add a new  
2690 benchmark (Section 5.4), as well as a developer guide showcasing the methodology to  
2691 add a new fault model through the instruction skip example (Section 5.5). At the end  
2692 of this chapter, we discuss the tool’s particularities and limitations (Section 5.6). The  
2693 evaluation of BINSEC/ASE is detailed in Chapter 6.

2694 The BINSEC/ASE tool can be found as an artifact on GitHub<sup>2</sup> and Zenodo<sup>3</sup>.

<sup>1</sup>BINSEC version of September 1<sup>st</sup> 2021

<sup>2</sup>[https://github.com/binsec/esop2023\\_artefact](https://github.com/binsec/esop2023_artefact)

<sup>3</sup><https://zenodo.org/record/7507112>

## 5.1 Overview

Adversarial Reachability aims at reasoning about the impact of an advanced attacker, able to inject multiple faults, on a program’s security properties. We designed an algorithm to answer the Adversarial Reachability problem, Adversarial Symbolic Execution. We implement it inside a tool, BINSEC/ASE, answering Adversarial Reachability from a bug-finding point of view (under-approximation).

This tool’s aim is to be integrated into the security evaluation of programs. Once a risk assessment has identified threat models and security properties (not in the scope of this thesis), BINSEC/ASE can help the security expert determine if the program is vulnerable or not, and provide attack paths allowing the attacker to violate the security property. This can be used to identify critical code sections to try and mount full attacks on the program. Another use is countermeasure evaluation. Adding protections also adds attack surface, so protecting a program against one attack may open the way for a new attack for the same attacker model, hence an iterative process using a tool such as BINSEC/ASE.

In the following, we motivate the main high-level design choices made for BINSEC/ASE design.

- **ISA Abstraction Level:** we implemented Adversarial Symbolic Execution (ASE) at the ISA level. ISA stands for Instruction Set Architecture, also called assembly, it defines an interface between hardware and software. An analysis at ISA level makes the tool independent from source-level languages and constructs hindering analysis precision such as register dynamics. We also wanted to be independent of the specific micro-architecture of the system running the program, which requires heavy modeling for each target system, hindering scalability and for which specifications are often closed-sourced. A program may be launched on multiple target architectures, requiring multiple costly analyses when considering the micro-architecture. ISA represents directly the code executed by the processor, instructions and operands. Many fault models characterizing physical fault injection in the literature are at the ISA level;
- **Building Inside BINSEC:** BINSEC is a symbolic execution engine working at the ISA level, designed and developed at CEA LIST, which makes help to understand BINSEC’s inner workings and to debug more easily available, which saves a lot of time and effort. The BINSEC tool consists of about 60k loc of Ocaml, the symbolic execution engine contains around 3.2k loc. We modified only the SE engine, implementing ASE in 5.8k loc;
- **ISA Support:** we currently support Intel x86-32 bits and ARM 32 bits architectures. Support can be extended to other architectures as they become available in BINSEC.

## 5.2 Background: the BINSEC Tool

In this section, we present the BINSEC tool, upon which we built BINSEC/ASE. Details of our modifications can be found in the next Section (5.3). BINSEC is an open-source symbolic execution engine for binary programs developed at the CEA by the Binsec team<sup>4</sup>. We use BINSEC version 0.4.0 in this work, forked September 1<sup>st</sup> 2021.

<sup>4</sup><https://binsec.github.io/>

### 5.2.1 BINSEC Presentation

We implement BINSEC/ASE on top of the BINSEC symbolic engine [DBT<sup>+</sup>16, DB15, BHL<sup>+</sup>11]. It has already been used in a number of significant case studies [BDM17, RBB<sup>+</sup>19, RBB<sup>+</sup>21, DBR20, DBR21], and it is notably able to achieve bounded verification (k-completeness) and to reasonably deal with symbolic pointers [FDBL18].

### 5.2.2 General Work-Flow

We detail in this section the part of BINSEC’s general workflow that is relevant to the work presented in this thesis, from user inputs, internal path exploration and communications with the SMT solver, to the provided output. Figure 5.1 provide an overview of this workflow.

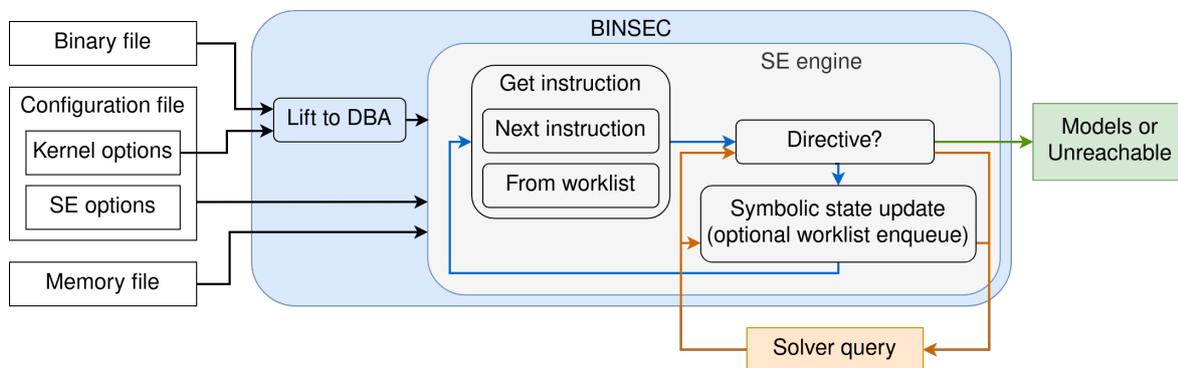


Figure 5.1: Overview of BINSEC workflow for symbolic execution

#### 5.2.2.1 Parameters

BINSEC can be parameterized using command line options or by providing it with a configuration file. The user defines a number of exploration parameters. We list here the main ones we use, divided into two categories, the kernel options and the symbolic exploration options.

**Kernel options.** They define the core elements required for BINSEC to run:

- the path to the binary *file* to analyse,
- the *ISA* for which the binary has been compiled,
- the *entrypoint* from which the symbolic analysis starts. It can be an address or more commonly the symbol of a function like `<main>`.

**Symbolic execution options.** They describe SE parameters and goals.

- Exploration objectives are defined in terms of *directives*. They can be of various forms such as *reach* an address once or multiple times, with or without a condition, *cut* the exploration when reaching an address, with or without a condition, or force constraints at a location with an *assume*;
- BINSEC performs an analysis bounded by the maximal number of instructions analyzed per path, set by the user;
- As the analysis starts at the defined entrypoint, the user can specify some variable initializations in a *memory* file. In particular, it is used to provide an initial address to the stack pointer, easing the analysis without loss of generality.

### 2769 5.2.2.2 Intermediate Representation: DBA

2770 To provide an analysis independent of the specific ISA used, BINSEC starts by reading  
 2771 the binary and lifting the ISA instructions to an intermediate representation called  
 2772 DBA (short for Dynamic Bitvector Automata). This is a language specific to BINSEC.  
 2773 It abstracts each ISA instruction into a block of DBA instructions detailing registers,  
 2774 memory and flag updates with arithmetic expressions, as well as the next instruction  
 2775 to execute. There are four main kinds of DBA instructions: assignments, conditional  
 2776 jumps, static jumps and dynamic jumps.

2777 Table 5.1 illustrates the DBA syntax from a very simple C example, with the  
 2778 associated ISA instructions (Intel x86). For readability, the updated expressions of  
 2779 flags are not displayed, except when relevant. The ISA instructions don't reflect the  
 2780 full expressivity of the *assume* C statement, just the conditional branching.

### 2781 5.2.2.3 Analysis Workflow

2782 The symbolic execution engine will start its analysis at the *entrypoint* defined by the  
 2783 user with the initialization contained in the provided *memory file*. Then instructions  
 2784 are decoded one by one in their execution order.

- 2785 – In the case of assignments, updates are made to the symbolic state, which records  
 2786 the current value, concrete or symbolic, of each variable;
- 2787 – For a conditional jump, the condition is added to the path predicate constraints  
 2788 sent to the SMT solver. The analysis will fork if both branches are possible,  
 2789 selecting one to keep exploring and storing the other in a worklist to continue the  
 2790 exploration later;
- 2791 – In the case of a static jump, the analysis will decode the instruction at the newly  
 2792 provided address and continue from there;
- 2793 – Lastly, for dynamic jumps, also known as indirect jumps (the target address has  
 2794 been computed at runtime and is stored in a register), the analysis asks the  
 2795 SMT solver for  $n$  possible values (3 by default but it can be changed through  
 2796 configuration) for the jump target address and forks the analysis, continuing one  
 2797 and storing all the others in the worklist. In practice, only one jump address  
 2798 usually makes sense.

2799 BINSEC uses a depth-first approach with its worklist of paths to analyze. The  
 2800 exploration of a path ends when it meets a directive's criteria or when it exceeds the  
 2801 maximal instruction depth.

### 2802 5.2.2.4 SMT Solvers

2803 BINSEC uses the QF\_ABV theory, standing for Quantifier Free, Arrays and Bit-  
 2804 Vectors, to discharge created formulas. BINSEC has a native binding with the SMT  
 2805 solver Bitwuzla<sup>5</sup> [NP20], winner of SMT-COMP 2022<sup>6</sup> for its category. It also inter-  
 2806 faces with other solvers, such as z3<sup>7</sup> [DMB08] or Boolector<sup>8</sup> [BB09], through the use of  
 2807 the smtlib, a universal API for SMT solvers. A connection through the smtlib is slower  
 2808 than a native binding, internally optimized for the solver.

<sup>5</sup><https://bitwuzla.github.io/>

<sup>6</sup><https://smt-comp.github.io/2022/results/qf-bitvec-single-query>

<sup>7</sup><https://github.com/Z3Prover/z3>

<sup>8</sup><https://boolector.github.io/>

Table 5.1: Example of C and Intel x86 instructions translated to DBA

C code	Intel x86	DBA
<code>int x = 9;</code>	<code>mov [ebp + 0xffffffff4], 0x9</code>	<pre> 0: @[ (ebp&lt;32&gt; + -12&lt;32&gt;),4] := 9&lt;32&gt;; 1: <b>goto</b> (08049d1c, 0) </pre>
<code>x = x + 1;</code>	<code>add [ebp + 0xffffffff4], 0x1</code>	<pre> 0: res32&lt;32&gt; := (@[(ebp&lt;32&gt; + -12&lt;32&gt;),4] + 1&lt;32&gt;); 1: OF&lt;1&gt; := [...]; 2: SF&lt;1&gt; := [...]; 3: ZF&lt;1&gt; := [...]; 4: AF&lt;1&gt; := [...]; 5: PF&lt;1&gt; := [...]; 6: CF&lt;1&gt; := [...]; 7: @[ (ebp&lt;32&gt; + -12&lt;32&gt;),4] := res32&lt;32&gt;; 8: <b>goto</b> (08049d20, 0) </pre>
<code>assert(x == 8);</code>	<code>cmp [ebp + 0xffffffff4], 0x8</code>	<pre> 0: res32&lt;32&gt; := (@[(ebp&lt;32&gt; + -12&lt;32&gt;),4] - 8&lt;32&gt;); 1: OF&lt;1&gt; := [...]; 2: SF&lt;1&gt; := [...]; 3: ZF&lt;1&gt; := (0&lt;32&gt; = res32&lt;32&gt;); 4: AF&lt;1&gt; := [...]; 5: PF&lt;1&gt; := [...]; 6: CF&lt;1&gt; := (@[(ebp&lt;32&gt; + -12&lt;32&gt;),4] &lt;u 8&lt;32&gt;); 7: <b>goto</b> (08049d24, 0) </pre>
	<code>jz 0x8049d44</code>	<pre> 0: <b>if</b> ZF&lt;1&gt; <b>goto</b> (08049d44, 0) <b>else goto</b> 1 1: <b>goto</b> (08049d26, 0) </pre>

### 2809 5.2.2.5 Output

2810 Among the usages BINSEC has, we are interested in standard reachability in this thesis.  
 2811 When the SE engine encounters a code address associated with a reach directive, it  
 2812 creates the associated formulas for the SMT solver to assess satisfiability. The solver  
 2813 will either answer that the path does not validate the directive and BINSEC forwards  
 2814 the information that the code location is unreachable from that path, either the solver  
 2815 provides a model for symbolic variables such that the formula is true and BINSEC  
 2816 forwards this model for symbolic inputs to the user.

### 2817 5.2.3 Summary

2818 BINSEC is a symbolic execution engine for binary programs with good properties:

- 2819 – It can assess the standard reachability of code locations;
- 2820 – It is correct assuming correct logical encoding of path predicates and correctness  
 2821 of the underlying solver (if the solver answers SAT with a model, then the formula  
 2822 is indeed satisfiable and the model is a real one);
- 2823 – It achieves bounded verification (k-completeness), assuming all dynamic jumps  
 2824 can be enumerated, which is usually the case in practice.

## 2825 5.3 BINSEC/ASE Implementation

2826 The work described in the rest of this chapter has been conducted during this thesis.  
 2827 BINSEC symbolic engine core is the only part we changed in order to implement BIN-  
 2828 SEC/ASE. We start by describing the main modifications made to BINSEC, before  
 2829 detailing the fault models' implementation and the implementation of our optimiza-  
 2830 tions.

### 2831 5.3.1 BINSEC/ASE Overview

2832 From a user point of view, using BINSEC/ASE is very similar to using the BINSEC  
 2833 tool. The main difference is the attacker model the user provides in the configuration.  
 2834 BINSEC/ASE then assess the adversarial reachability of the attacker's goal location  
 2835 and outputs possible attack paths if they exist. Internally, several changes are needed  
 2836 in the symbolic execution engine to account for the attacker model's impact on the  
 2837 program. Figure 5.2 presents an overview of the BINSEC/ASE workflow, detailed in  
 2838 the following sections.

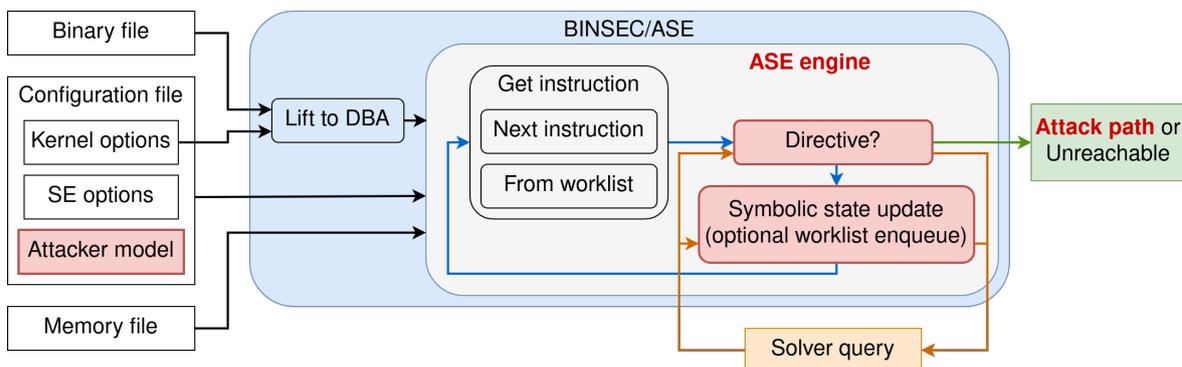


Figure 5.2: Overview of BINSEC/ASE workflow

## 2839 **5.3.2 ASE Implementation**

2840 The main modifications to the SE algorithm are described in chapter 4. We detail in  
2841 this section practical details regarding their implementations.

### 2842 **5.3.2.1 Attacker Model Parameters and Goal**

2843 The user configures the attacker model to consider for the analysis. It takes the form  
2844 of new parameters inside the configuration file, in particular, to specify the attacker’s  
2845 capability and the maximum number of faults that are allowed in one attack path.  
2846 More details on the configuration of the attacker model are available in the user guide  
2847 provided in Section 5.4.2.2. The attacker goal is expressed with a dedicated *reach*  
2848 directive.

### 2849 **5.3.2.2 Global Variables**

2850 We introduce new global variables in the symbolic state to keep track of injected faults.

2851 **Fault Counter.** We use a symbolic variable, called *FAULT\_NUMBER* to keep track  
2852 of how many faults are active, computed by summing the individual fault activation  
2853 variables. It is initialized at 0 at the initialization of the analysis.

2854 **Fault Budget.** We recover the configuration value given by the user setting the max-  
2855 imum number of faults allowed to the attacker model and store it inside a symbolic  
2856 variable called *MAX\_FAULTS* (with a concrete fixed value) at the analysis initializa-  
2857 tion. Having the fault budget as a symbolic variable makes it easy to compare it to  
2858 our symbolic fault counter in queries.

2859 **Activation Constraint.** Without some additional constraints, a fault could be declared  
2860 active but have no effect. For instance, a value could be corrupted to a new but  
2861 identical value, or a reset fault could be performed on a variable already at 0. To count  
2862 only active faults, truly modifying a variable’s value, we generate and aggregate a  
2863 set of constraints in a symbolic variable called *ACTIVATION\_CONSTRAINTS*. This  
2864 variable is a boolean value initialized at true. For each fault in the path predicate,  
2865 *ACTIVATION\_CONSTRAINTS* is updated by the computation of a logical *and* of its  
2866 current value and the new constraint with respect to the fault model used. A constraint  
2867 ensuring *ACTIVATION\_CONSTRAINTS* is true is added to each query sent to the  
2868 SMT solver. To sum up, we add activation constraints to remove ineffective faults for  
2869 a more accurate fault count and possibly help the solver eliminate them.

### 2870 **5.3.2.3 Assignments**

2871 In the program analysis process, when an assignment instruction is encountered, the  
2872 symbolic state update is augmented.

2873 **Behavior Wrapper.** We created a wrapper around the assignment behavior, which is in  
2874 charge of deciding which behavior should be applied for this assignment. In particular,  
2875 should the normal assignment behavior happen, or a behavior injecting a fault?

2876 **Fault Filter.** Each assignment isn’t a possible fault location, there are several conditions  
2877 to satisfy.

- 2878 1. The user has to have selected an attacker model able to perform faults, i.e. the  
2879 selected attacker capability isn’t *None* and the fault budget is strictly greater  
2880 than 0;

- 2881 2. The instruction must be in the target ranges of the injection set by the user in  
 2882 the configuration;
- 2883 3. The type of the variable on the left-hand side of the assignment cannot be a  
 2884 temporary variable (for DBA internal use with no architectural reality). Flags  
 2885 are faulted or not according to the attacker model;
- 2886 4. As we have no efficient algorithm for faults on addresses, we set a threshold above  
 2887 which a variable is considered an address and is not faulted. This threshold is  
 2888 set by default at `0x05000000` and can be changed by the user;
- 2889 5. There are some variables the user may not want to fault, given in the configura-  
 2890 tion. If the assigned variable is in this *blacklist*, it is not faulted. We mainly use  
 2891 this feature to avoid faulting the stack pointer.

2892 **Fault Selection.** Once the instruction qualifies for fault injection, a match-case selects  
 2893 the assignment function implementing the desired behavior. If it isn't a target for fault  
 2894 injection, the normal assignment function (from BINSEC) is selected. Those functions  
 2895 update the symbolic state according to the instruction and fault model.

#### 2896 5.3.2.4 Conditional Jumps

2897 Conditional jumps can be a place of fault injection and bear some fault injection ma-  
 2898 chinery.

2899 **Behavior Wrapper.** We also created a wrapper around the conditional jump behavior,  
 2900 which is in charge of deciding which behavior should be applied for this conditional  
 2901 jump.

2902 **Number of Fault Check.** A question was: when to make sure the number of active  
 2903 faults doesn't exceed the fault budget, for Forkless faults in particular? Checking the  
 2904 number of faults at each fault location is costly and doesn't bring more information,  
 2905 since all assignments done after the last jump don't bring new constraints. Checking  
 2906 the number of faults at the end of the path leaves room to explore unfeasible paths,  
 2907 that is to say, paths that require more faults than the user gave to the attacker model.  
 2908 We opted for the third option, checking the number of faults at each conditional jump.  
 2909 It takes advantage of the already existing queries at conditional jump locations, adds  
 2910 information due to the activation constraint and allows to prune infeasible paths. To  
 2911 implement this, we augment the normal conditional jump queries (one for the then  
 2912 branch and one for the else branch) with a boolean condition stating that the current  
 2913 value of the fault counter should be lesser or equal to the fault budget. We coupled  
 2914 this number of faults check with checking the validity of the activation constraints.

2915 **Fault Filter and Selection.** BINSEC/ASE implements the test inversion fault model,  
 2916 which injects faults into conditional jumps. A filter is first applied, containing only  
 2917 elements 1,2 and 5 of the enumeration of the assignment fault filter (Section 5.3.2.3).  
 2918 Then, if the instruction is eligible for fault injection, the wrapper selects the faulty  
 2919 behavior, otherwise, it selects the normal behavior, always passing the augmented  
 2920 query.

#### 2921 5.3.2.5 Exploration Directives

2922 The goal of the exploration is to assess the adversarial reachability of a code location.  
 2923 Our modifications to BINSEC extend its original reachability assessment to adversarial  
 2924 reachability.

2925 On top of the fault injections, a guard is added to the reachability query, when  
2926 assessing the feasibility of a path reaching a code location specified by a directive. It  
2927 consists of a combination of the fault check (the current value of the fault counter needs  
2928 to be lesser or equal to the fault budget) and the activation constraints check (ensuring  
2929 ineffective faults are not activated).

### 2930 5.3.2.6 BINSEC/ASE Output

2931 BINSEC/ASE output is similar to BINSEC's. If there exists a model for symbolic  
2932 variables such that the attacker's goal is reached, it will be displayed. While in BINSEC  
2933 symbolic variables are mostly input of the program, with BINSEC/ASE, the model also  
2934 includes an activation value for each possible fault location, and a symbolic corruption  
2935 value depending on the fault model. For instance, there will be none for reset faults, it  
2936 indicates which bit is flipped for a bit-flip fault model, or an exact value in the case of  
2937 arbitrary data faults. More details are discussed in Section 5.3.3. A model for program  
2938 input and a fault sequence defines an adversarial path.

### 2939 5.3.2.7 Statistics

2940 We extended BINSEC's computed statistics with, for instance, the number of injection  
2941 points in total and per path, statistics on queries such as the queries trivially true  
2942 or false, how many were created and how many were sent. We also implemented the  
2943 optional computation of the number of *ite* operators per query, giving a sense of the  
2944 query complexity. This is optional as it is costly to compute, requiring to traverse each  
2945 query. Those statistics help with the development effort, in particular, to ensure the  
2946 correctness of the analysis and assist the debugging process. They also allow comparing  
2947 techniques, encoding, optimizations, etc. More information on their usage is available  
2948 in the user guide (Section 5.4.4).

### 2949 5.3.2.8 Implementation Details

2950 BINSEC/ASE features other implementation details. They are not purely derived from  
2951 Adversarial Symbolic Execution algorithms but are tweaks to improve performance.

2952 **Symbolic default value.** When an evaluation is performed on a model, to see if it  
2953 can be trivially resolved or if the solver needs to be used, and that model contains  
2954 free, unconstrained variables, BINSEC attributes them an arbitrary value. It helps  
2955 for performance inside BINSEC. However, since most of our symbolic variables relate  
2956 to faults, we noticed it improves performance to give them the default value of 0. It  
2957 means faults are not activated by default.

2958 **Simplification rules.** The operator we favor for forkless encodings, the *ite* operator  
2959 is rarely used in DBA natively. For this reason, very few algorithmic simplification  
2960 rules for symbolic terms dedicated to the *ite* operator were implemented in BINSEC.  
2961 Simplification rules are used when symbolic terms are created and before they are  
2962 added to the symbolic state, path predicate or queried upon. This saves solver effort,  
2963 which is the main bottleneck in standard symbolic execution. To accommodate our use  
2964 of *ites*, we implemented in BINSEC/ASE additional rules that have been backported  
2965 to the BINSEC tool.

- 2966 – When a unary operator, *unop*, *not* for instance, is applied to an *ite*, it can be  
2967 forwarded to the then and else members of the *ite*;

- 2968 – When a binary operator, *binop*, *plus* or *minus* for instance, is applied to an *ite*  
 2969 term and a constant, the binary operation can be forwarded inside the then and  
 2970 else members of the *ite*;  
 2971 – When a test operation, *compop*, equal or lower than for example, is applied to an  
 2972 *ite*, the expression can be simplified if we know the results of the comparison for  
 2973 each internal member *t* and *e* of the *ite*.

2974 Formal rules are presented in Table 5.2.

Original expression	Expression after simplification
$unop\ (ite\ c\ ?\ t : e)$	$ite\ c\ ?\ (unop\ t) : (unop\ e)$
$binop\ (ite\ c\ ?\ t : e)\ cst$	$ite\ c\ ?\ (binop\ t\ cst) : (binop\ e\ cst)$
$binop\ cst\ (ite\ c\ ?\ t : e)$	$ite\ c\ ?\ (binop\ cst\ t) : (binop\ cst\ e)$

Original expression	Condition	Simplification
$compop\ (ite\ c\ ?\ t : e)\ cst$	$(compop\ t\ cst) \ \&\&\ (compop\ e\ cst)$	1
$compop\ (ite\ c\ ?\ t : e)\ cst$	$\neg(compop\ t\ cst) \ \&\&\ (compop\ e\ cst)$	$\neg c$
$compop\ (ite\ c\ ?\ t : e)\ cst$	$(compop\ t\ cst) \ \&\&\ \neg(compop\ e\ cst)$	c
$compop\ (ite\ c\ ?\ t : e)\ cst$	$\neg(compop\ t\ cst) \ \&\&\ \neg(compop\ e\ cst)$	0

Table 5.2: Simplification rules added to BINSEC/ASE

### 2975 5.3.3 Forkless Fault Models

2976 We detail in this section how our forkless fault models are implemented. To date,  
 2977 we have implemented arbitrary data faults, reset faults, bit-flips, test inversions and  
 2978 instruction skips. Those are generic and commonly used fault models to represent  
 2979 faults from hardware fault injection attacks and can be applied to other fault injection  
 2980 means.

#### 2981 5.3.3.1 Data Faults

2982 Once this fault behavior has been selected for an assignment, it is possible to perform  
 2983 the number of faults check and the activation constraints check, though it is not the  
 2984 recommended way. If the checks don't pass, the normal –non faulted behavior is  
 2985 executed.

2986 Otherwise, two new symbolic variables are created, one boolean value representing  
 2987 the activation of the fault, and a variable the size of the left-hand side variable repre-  
 2988 senting the fault effect. As a convenience, we add a suffix to those variables' names with  
 2989 the address of the current instruction. It helps to analyze the results and locate the  
 2990 faults in the ISA of the program. The fault counter is incremented with the symbolic  
 2991 value of the activation variable.

2992 All elements required to generate the activation constraint of the fault according to  
 2993 the fault model are set. We update the symbolic state with this additional constraint.

2994 Finally, the new right-hand side expression is computed according to the fault model  
 2995 and returned to the wrapping function which performs the final symbolic state update.

2996 A summary of this process is shown in Algorithm 5.1 for arbitrary data faults. Note  
 2997 that  $:=$  denotes an assignment of code variables, while  $\triangleq$  denotes the mapping of a  
 2998 variable in the symbolic state.

---

**Algorithm 5.1:** Implementation sketch of a forkless arbitrary data fault (with *ite* operator) assignment process

---

**Input:** path predicate  $\Phi$ , assignment instruction  $x := y$   
**Data:** fault counter  $NUMBER\_FAULT$  and activation constraint  $ACTIVATION\_CONSTRAINT$  stored in  $\Phi$   
**Output:** Updated  $\Phi$

```

1 Function eval_assign_normal( $\Phi, x, y$ ) is
2 |   return  $\Phi \wedge (x \triangleq expr')$ 
3 end

4 Function eval_assign_arbitrary_data_forkless( $\Phi, x, y$ ) is
5 |    $addr := GetAddress()$ 
6 |   Create in  $\Phi$  a symbolic value named  $b_{addr}$  of size 1
7 |   Create in  $\Phi$  a symbolic value named  $non\_det_{addr}$  of size  $sizeof(y)$ 
8 |    $\Phi := \Phi \wedge (NUMBER\_FAULTS \triangleq NUMBER\_FAULTS + b_{addr})$ 
9 |    $new\_constraint := (non\_det_{addr} \neq y) \parallel \neg b_{addr}$ 
10 |   $\Phi := \Phi \wedge (ACTIVATION\_CONSTRAINT \triangleq$ 
    |   $ACTIVATION\_CONSTRAINT \& new\_constraint)$ 
11 |   $y' := ite\ b_{addr} ? non\_det_{addr} : y$ 
12 |  return  $\Phi, y'$ 
13 end

14 Function eval_assign_handler( $\Phi, x, y$ ) is
15 |   if IsFaultLocationFilter() then
16 |     |   case ArbitraryDataFaultForkless do
17 |       |    $\Phi', y' := eval\_assign\_arbitrary\_data\_forkless(\Phi, x, y)$ 
18 |       |   return eval_assign_normal( $\Phi', x, y'$ )
19 |     |   end
20 |     |   ...
21 |   else
22 |     |   return eval_assign_normal( $\Phi, x, y$ )
23 |   end
24 end

```

---

2999 **Arbitrary data faults.** We implemented four different arithmetic encodings for arbitrary  
3000 data faults, we show in table 6.13 those encodings and their corresponding activation  
3001 constraints. We use  $b_{addr}$  to represent the activation variable and  $non\_det_{addr}$  to  
3002 represent the fault effect. Note that *ites* are written with a C-like syntax  $ite\ c ? t : e$   
3003 where  $c$  is the boolean condition that if true, the expression results in the expression  
3004  $t$  and if false, the *ite* expression results in the expression  $e$ ;  $\oplus$  is used to denote a  
3005 bitwise xor operator. We omit the extension operations from boolean to variable size  
3006 value for readability. Forkless arbitrary data fault with the *ite* operator is illustrated  
3007 in Algorithm 5.1.

3008 **Reset Faults.** Similar to arbitrary data faults, we implement a reset fault model. The  
3009 differences lie in the fault effect, which does not require a symbolic variable to represent  
3010 it but is simply the value 0 or an expression resulting in 0 if  $b_{addr}$  is active, and the  
3011 activation constraint that is adapted to exclude fault location assigning the value 0.  
3012 We also implemented four different forkless encodings with various operators. There

Table 5.3: Arbitrary data encodings and their associated activation constraint

Fault model	Fault expression	Activation constraint
None	$x := y$	
Inlined if-then-else	$x := \text{ite } b_{addr} ? non\_det_{addr} : y$	$(non\_det_{addr} \neq y) \parallel \neg b_{addr}$
Multiplication	$x := y + b_{addr} \times non\_det_{addr}$	$(non\_det_{addr} \neq 0) \parallel \neg b_{addr}$
Bitwise and	$x := y + (-b_{addr}) \& non\_det_{addr}$	$(non\_det_{addr} \neq 0) \parallel \neg b_{addr}$
Bitwise xor	$x := y \oplus (-b_{addr}) \& non\_det_{addr}$	$(non\_det_{addr} \neq 0) \parallel \neg b_{addr}$

3013 are details in Table 5.4.

Table 5.4: Reset encodings and their associated activation constraint

Fault model	Fault expression	Activation constraint
None	$x := y$	
Inlined if-then-else	$x := \text{ite } b_{addr} ? 0x0000 : y$	$(y \neq 0x0000) \parallel \neg b_{addr}$
Subtraction	$x := y - b_{addr} \times y$	$(y \neq 0x0000) \parallel \neg b_{addr}$
Bitwise and	$x := y - (-b_{addr}) \& y$	$(y \neq 0x0000) \parallel \neg b_{addr}$
Bitwise xor	$x := y \oplus (-b_{addr}) \& y$	$(y \neq 0x0000) \parallel \neg b_{addr}$

3014 Here, the activation constraint is always the same: the original right-hand side  
3015 expression should not be 0.

3016 **Bit-flip Faults.** The implementation of the bit-flip fault model is again very similar to  
3017 the arbitrary data fault process. Two symbolic variables are created for the activation  
3018 and the effect, their combination is different to result in a bit-flip, as shown in Table  
3019 5.5.

Table 5.5: Bit-flip encoding and its activation constraint

Fault model	Fault expression	Activation constraint
None	$x := y$	
Inlined if-then-else	$x := \text{ite } b_{addr} ?$ $y \oplus (1 \ll non\_det_{addr}) : y$	$(non\_det_{addr} < sizeof(y))$ $\parallel \neg b_{addr}$

3020 We only implemented the *ite* encoding for bit-flip but others can be imagined. The  
3021 activation constraint limits the value of  $non\_det_{addr}$  with respect to the size of the  
3022 right-hand side expression to avoid overflows. This encoding of bit-flips allows only  
3023 one bit-flip per assignment, even in a multi-fault context. This represents real classes  
3024 of attacks, such as the Rowhammer attack. We believe an extended bit-flip encoding  
3025 can be devised using masks to encode the possibility of multiple bit-flips in one fault  
3026 location.

### 3027 5.3.3.2 Test Inversion Faults

3028 The implementation of the test inversion fault model differs from the implementation  
3029 of other data faults. The assignment handler takes the normal –non-faulted assignment  
3030 function. The faulty behavior is instantiated in the conditional jump handler. When  
3031 the test inversion behavior is selected, it computes the new, faulted condition to give  
3032 to the normal conditional jump function. This is illustrated in Algorithm 5.2.

**Algorithm 5.2:** Implementation sketch of a forkless test inversion process

---

```

Input: path predicate  $\Phi$ , conditional jump instruction if cdt lt else le
Data: worklist WL, fault counter NUMBER_FAULT and fault budget
        MAX_FAULT stored in  $\Phi$ 
Output: Updated  $\Phi$  and WL

1 Function eval_assign_normal( $\Phi$ , cdt, lt, le) is
2   | if  $\Phi \wedge cdt \wedge (NUMBER\_FAULTS \leq MAX\_FAULT)$  is satisfiable then
3   |   | Add ( $\Phi \wedge cdt$ , lt) to WL
4   | end
5   | /* Idem for else branch ( $\neg cdt$ )                                     */
6 end

6 Function eval_test_inversion_forkless( $\Phi$ , cdt) is
7   | addr := GetAddress()
8   | Create in  $\Phi$  a symbolic value named baddr of size 1
9   |  $\Phi := \Phi \wedge (NUMBER\_FAULTS \triangleq NUMBER\_FAULTS + b_{addr})$ 
10  | cdt' := ite baddr ?  $\neg cdt$  : cdt
11  | return  $\Phi$ , cdt'
12 end

13 Function eval_condition_jump_handler( $\Phi$ , cdt, lt, le) is
14  | if IsFaultLocationFilter() then
15  |   | case TestInversionForkless do
16  |   |   |  $\Phi'$ , cdt' := eval_test_inversion_forkless( $\Phi$ , cdt)
17  |   |   | return eval_conditional_jump_normal( $\Phi'$ , cdt', lt, le)
18  |   | end
19  |   | ...
20  | else
21  |   | return eval_conditional_jump_normal( $\Phi$ , cdt, lt, le)
22  | end
23 end

```

---

3033 There is no need for an activation constraint in this fault model, as the bitwise  
3034 negation of a boolean value is always its opposite. We used the *ite* operator but other  
3035 forkless encodings are possible.

3036 **5.3.3.3 Instruction Skip Faults**

3037 The last fault model we implemented is the instruction skip fault model. The idea is  
3038 to be able to create a nop (no-operation) with a forkless encoding. As each kind of  
3039 DBA instruction can be the target of a skip, the remaining ones (for static jumps and  
3040 dynamic jumps) are refactored with a wrapper and a fault filter like the conditional  
3041 jump handler. A summary of the fault effect for each DBA instruction is presented in  
3042 Table 5.6, where *addr* represents a distant code address, *next* is the next instruction in  
3043 the memory layout and *expr* corresponds to an expression computing one or multiple  
3044 addresses. No activation constraint is required. Except for assignments, the other  
3045 faulted instructions are likely to open new paths.

3046 **Faulting a DBA Block.** The instruction skip fault model applies to an entire DBA

Table 5.6: Instruction skip encoding for each dba instruction

DBA instruction	Fault model	Fault expression
Assignment	None	$x := y$
	Forkless Skip	$x := \text{ite } b_{addr} ? x : y$
Conditional jump	None	if $cdt$ goto $addr$ else goto $next$
	Forkless Skip	if $(cdt \ \& \ \neg b_{addr})$ goto $addr$ else goto $next$
Static jump	None	goto $addr$
	Forkless Skip	if $b_{addr}$ goto $next$ else goto $addr$
Dynamic jump	None	goto $expr$
	Forkless Skip	goto $(\text{ite } b_{addr} ? next : expr)$

3047 block, not just one DBA instruction, to correspond to one ISA instruction being  
 3048 skipped. This is relevant only for assignment instructions that need more than one  
 3049 instruction for flags and register updates.

3050 **Assignment.** The Forkless encoding of the instruction skip fault model acts for assign-  
 3051 ments as a data fault, with the effect corresponding to the old value of the right-end-side  
 3052 variable. This is performed for all variables involved in the computation, may it be  
 3053 flags, register or memory operations, with the same activation variable. The difficulty  
 3054 is to know if one instruction in the DBA block is blacklisted (if it updates a variable  
 3055 that should always remain without faults). If so, the fault needs to be canceled for  
 3056 all previous instructions in the block as they were optimistically injected with faults.  
 3057 In practice, the variable containing the activation for the whole bloc is assigned the  
 3058 value 0 in the path predicate. An assignment DBA block ends with a static jump to  
 3059 the next instruction, we reset the cancellation flag and the shared activation variable  
 3060 at this stage.

3061 **Conditional Jump.** Here, the condition is augmented with the possibility of a fault  
 3062 before it is given to the normal conditional jump process.

3063 **Static Jump.** This instruction is in effect replaced by a conditional jump, with the  
 3064 fault activation variable as the condition.

3065 **Dynamic Jump.** The expression that results in the target address of the jump is faulted  
 3066 as a data expression would be.

### 3067 5.3.4 Early Detection of Fault Saturation (EDS)

3068 Our first optimization, EDS aims at stopping injection as soon as possible. The imple-  
 3069 mentation follows the algorithm described in Section 4.4.1. Note that this optimization  
 3070 has been designed and implemented for data faults only. Test inversion contains very  
 3071 few fault locations in comparison. We believe it can straightforwardly be adapted for  
 3072 instruction skips.

3073 A new behavior is added to the conditional jump handler, where the satisfiability  
 3074 query is first asked for  $nb_f < max_f$  and then for  $nb_f == max_f$  to detect the fault  
 3075 saturation. We created an internal flag value to register the activation of EDS and not  
 3076 inject fault when the flag is true.

3077 We encountered a challenge implementing EDS. Imagine some faults  $F$  are injected,  
 3078 then a conditional jump where no faults is needed to take a certain branch  $A$ . Then the  
 3079 constraint  $nb_f < max_f$  is recorded in the path predicate. Imagine the path encounters

3080 another conditional jump, where one of its branches,  $B$ , needs a fault in  $F$ . In the  
3081 context of a single-fault attacker model, taking the branch  $A$  constrains  $F$  to be inactive  
3082 since the sum of the faults before  $A$  needs to be 0, hence preventing the exploration of  
3083 branch  $B$ , even if it is possible for the attacker model.

3084 To solve this issue, we modified the implementation of the function performing  
3085 the solver call. An additional argument was added, containing the constraints that  
3086 shouldn't be added to the path predicate but that are still part of the query. We call  
3087 this parameter *such\_that*: can the path predicate and the conditional condition be true  
3088 *such that* this parameter is also true.

### 3089 5.3.5 Injection On Demand (IOD)

3090 Our second optimization, IOD, described in Section 4.4.2, aims at starting injection  
3091 as late as possible, and only adding new faults when necessary. Note that again, this  
3092 optimization has been designed with data faults. Test inversion contains very few fault  
3093 locations in comparison. We show later how it was adapted for instruction skips, and  
3094 focus the first part of the description on its application to data faults.

3095 Contrary to EDS, IOD required heavy refactoring. The injection on demand re-  
3096 quires to compute and keeping in parallel two path predicates. We name *dual* the  
3097 combined symbolic states. All functions have to be updated with this new type. Ei-  
3098 ther they need only one symbolic state and the caller function has to be modified to  
3099 give the relevant symbolic state, or the function receives a dual and its logic must be  
3100 adapted to take the dual into account.

3101 The initialization instantiates an extra symbolic state. Path predicate update func-  
3102 tions need to perform the same operations twice, once on the normal predicate and  
3103 once on the faulted predicate. This includes assignments, conditional jumps, static  
3104 jumps and dynamic jumps, but also all functions interacting with the SMT solver  
3105 and the directive handlers. A new internal variable was created to keep track of the  
3106 under-approximation counter. Its value is checked before injecting faults.

3107 The assignment process is duplicated, the normal path predicate is updated with  
3108 the normal assignment behavior, and the faulted path predicate is updated with the  
3109 faulty behavior presented above.

3110 **IOD for Instruction Skip.** All types of instructions need to be taken into account for  
3111 this fault model.

- 3112 – We keep the same mechanism for assignments and conditional jump, though  
3113 adding the possibility of skipping the latter;
- 3114 – We never merge for static jumps, as previous instruction skips do not influence  
3115 the static jump address;
- 3116 – We systematically merge for dynamic jumps. This is a choice different than  
3117 for data faults where paths are not merged there. As we do not have an effi-  
3118 cient algorithm for faults on addresses, this choice for data faults limits faults on  
3119 dynamically computed addresses. However, the forkless instruction skip doesn't  
3120 have limitations on skipping address assignments or on faults regarding addresses  
3121 in general, all instructions can be skipped.

### 3122 5.3.6 Sub-fault Simplification

3123 To improve performance further by reducing the number of injection points, we devised  
3124 a new optimization, sub-fault simplification. The idea is to detect faults subsumed by

3125 another fault, especially in the case of arbitrary data faults. The dominated fault can  
 3126 be nullified without loss of generality.

$\begin{aligned} x &:= 8 \\ x &:= x + 1 \end{aligned}$
--

(a) Original statement

$\begin{aligned} x &:= \text{ite } \text{fault\_here\_1} ? \text{fault\_value\_1} : 8 \\ x &:= \text{ite } \text{fault\_here\_2} ? \text{fault\_value\_2} : \\ &\quad (\text{fault\_here\_1} ? \text{fault\_value\_1} : 8) + 1 \end{aligned}$
---

(b) Forkless transformation for arbitrary data fault

Figure 5.3: Illustration of sub-fault simplification

3127 An example of this pattern is provided in Figure 5.3. The same left-hand-side  
 3128 variable is assigned multiple times, only the last one could be faulted with an arbitrary  
 3129 data fault to reduce fault injection points.

3130 We implemented this concept as a query preprocessing, adding the expression nul-  
 3131 lifying fault activation variables in queries. We designed two variants:

- 3132 – one where the dominated faults are detected on a per-query basis (*Free*),
- 3133 – and another using a memoization technique in the hope of improving performance  
 3134 (*Mem*).

3135 However, a dominated fault in one query is not necessarily dominated in another,  
 3136 hence there are some soundness issues with this technique. We detect unsound cases  
 3137 and propose two recovery mechanisms:

- 3138 – *Abort* where the SFS simplification is not added to the query,
- 3139 – and *Lazy* still trying the query and sending it again without SFS if the query is  
 3140 UNSAT.

3141 Unfortunately, experiments showed (see Section 6.3.2) that this optimization was  
 3142 costly in execution time due to query traversal. It was then abandoned.

### 3143 5.3.7 Forking Fault Models

3144 In order to compare our novel forkless encoding technique against the state-of-the-art  
 3145 techniques in the fairest way, we implemented forking encodings alongside FASE. Note  
 3146 that mutant generation scales worst. This allows us to restrict the comparison to the  
 3147 encodings only, the rest of the analysis workflow being the same inside BINSEC.

3148 Forking faults appear as other faulting options inside the instruction handler. Pre-  
 3149 viously described optimizations do not apply to forking faults as their placements are  
 3150 concrete.

#### 3151 5.3.7.1 Forking Data Faults

3152 This behavior differs from forkless data faults as the path is split between a path  
 3153 with faults for this instruction and one without. This is similar to a normal conditional  
 3154 jump. Both paths are stored in BINSEC’s worklist  $WL$  and one is dequeued to continue  
 3155 the exploration inside the wrapper. Algorithm 5.3 illustrates the general process for  
 3156 arbitrary data fault.

3157       The forking encodings split the path in two, in one path the activation variable  $b_{addr}$   
3158 is true, the assignment is faulted and the fault counter is incremented; in the second  
3159 path, the assignment proceeds with the normal behavior. We perform the number of  
3160 faults check at each fault location for forking faults because the number of faults is  
3161 concrete and it allows pruning unfeasible branches early on.

3162       The arbitrary data fault model simply replaces the right-hand side of the assignment  
3163 with a nondeterministic value. The reset fault model is identical except it replaces the  
3164 right-hand side with zeros, and bit-flips flip a bit with an encoding similar to the forkless  
3165 one.

### 3166 **5.3.7.2 Forking Test Inversion Faults**

3167       The implementation of the forking test inversion fault model follows the same principles  
3168 as the forking data fault models, except it takes place in the conditional jump handler.  
3169 A verification of the number of faults already performed is done, then the path is split  
3170 in two, one will execute the normal conditional jump and the other one, a conditional  
3171 jump with the negation of the condition.

### 3172 **5.3.7.3 Forking Instruction Skip Faults**

3173       The instruction skips fault model also has its forking counterpart implemented, still  
3174 following the same principles. Here, when the fault on an assignment block has to be  
3175 canceled because it updates a black-listed variable, the exploration is simply stopped  
3176 in the faulted branch and continues in the branch without the fault.

## 3177 **5.3.8 Conclusion**

3178       We saw in this section how BINSEC was used to build an Adversarial Symbolic Exe-  
3179 cution tool, BINSEC/ASE, able to assess the adversarial reachability of a location in a  
3180 binary program. Most of BINSEC is kept as is, the main changes happen at the path  
3181 predicate computation.

3182       We will now take a step back from technical details and explore a methodology to  
3183 analyze a new program as a user guide and a methodology to add a new fault model  
3184 as a developer guide.

---

**Algorithm 5.3:** Implementation sketch of a forking arbitrary data fault assignment process

---

**Input:** path predicate  $\Phi$ , assignment instruction  $x := y$   
**Data:** fault counter  $NUMBER\_FAULT$  and fault budget  $MAX\_FAULT$  stored in  $\Phi$   
**Output:** Updated  $\Phi$

```

1 Function eval_assign_normal( $\Phi, x, y$ ) is
2 |   return  $\Phi \wedge (x \triangleq expr')$ 
3 end

4 Function eval_assign_arbitrary_data_forking( $\Phi, x, y$ ) is
5 |   if  $\Phi \wedge (NUMBER\_FAULTS \leq MAX\_FAULT)$  is satisfiable then
6 |      $addr := GetAddress()$ 
7 |     Create in  $\Phi$  a symbolic value named  $b_{addr}$  of size 1
8 |     Create in  $\Phi$  a symbolic value named  $non\_det_{addr}$  of size  $sizeof(y)$ 
9 |     Split in  $\Phi$  the values for variable  $b_{addr}$ 
10 |    case  $0b1$  do
11 |       $\Phi' := \Phi \wedge (b_{addr} \triangleq 0b1)$ 
12 |       $\Phi' := \Phi' \wedge (NUMBER\_FAULTS \triangleq NUMBER\_FAULTS + 1)$ 
13 |       $new\_constraint := (non\_det_{addr} \neq y) \parallel \neg b_{addr}$ 
14 |       $\Phi' := \Phi' \wedge (ACTIVATION\_CONSTRAINT \triangleq$ 
15 |         $ACTIVATION\_CONSTRAINT + new\_constraint)$ 
16 |       $y' := non\_det_{addr}$ 
17 |      return eval_assign_normal( $\Phi', x, y'$ )
18 |    end
19 |    case  $0b0$  do
20 |       $\Phi' := \Phi \wedge (b_{addr} \triangleq 0b0)$ 
21 |      return eval_assign_normal( $\Phi', x, y$ )
22 |    end
23 |   else
24 |     eval_assign_normal( $\Phi, x, y$ )
25 |   end
26 return  $\Phi', y'$ 
end

```

---

## 5.4 User Guide: a Methodology to Analyse a New Program

The aim of this section is to provide a brief guide to the interested user so they can use BINSEC/ASE on their own programs, from the setup to the evaluation phase. In particular, this section does not aim to be exhaustive but to give an idea of the engineering effort required to perform an analysis and list the main steps. An overview is presented in Figure 5.4.

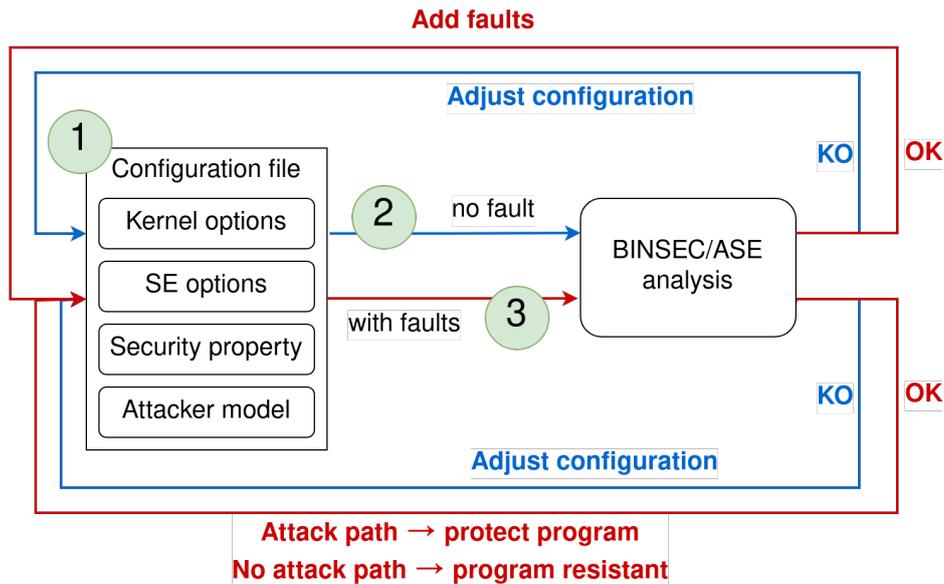


Figure 5.4: User guide workflow

### 5.4.1 Running Example

To illustrate the usage of BINSEC/ASE, let's consider again the basic version of the VerifyPIN program that we saw in Chapter 2. It is presented in Figure 5.5, with some PIN initialization (lines 7 to 17). We assume the attacker does not know the correct PIN, so they try one at random, "0 0 0 0" for instance. The attacker goal is to be authenticated, i.e. to pass the assertion line 41.

The program works as follows. The *main* function line 38 starts by initializing the PINs, and then it calls the *verifyPIN* function. This function resets the *g\_authenticated* variable to false by default, then calls the *byteArrayCompare* function to compute the PIN check. If *byteArrayCompare* returns 1, then the program authenticates the user. Inside *byteArrayCompare*, each digit is checked against the card PIN digit and if one is detected to be different, the function returns 0.

Note that this program is very naive and, in particular, not constant time. Other versions are proposed in FISSC [DPP<sup>+</sup>16] that are. Here, we only consider an attacker able to inject faults into a program, and not perform side-channel analysis.

We compile toward Intel x86-32 bits. As we use a previous version of BINSEC (2021), support for Intel x86-64 bits is not available. It becomes available in newer BINSEC versions. Note that our technique is architecture independent and considering 32-bit programs still makes sense in an embedded system context.

This C program is compiled with *gcc* using the following command line:

```
1 #define PIN_SIZE 4
2
3 bool g_authenticated = 0;
4 int g_userPin[PIN_SIZE];
5 int g_cardPin[PIN_SIZE];
6
7 void initialize() {
8     int i;
9     // card PIN = 1 2 3 4
10    for (i = 0; i < PIN_SIZE; ++i) {
11        g_cardPin[i] = i+1;
12    }
13    // user PIN = 0 0 0 0
14    for (i = 0; i < PIN_SIZE; ++i) {
15        g_userPin[i] = 0;
16    }
17 }
18
19 bool byteArrayCompare(int* a1, int* a2, int size) {
20     int i;
21     for (i = 0; i < size; i++) {
22         if (a1[i] != a2[i]) {
23             return 0;
24         }
25     }
26     return 1;
27 }
28
29 void verifyPIN() {
30     g_authenticated = 0;
31     if (byteArrayCompare(g_userPin, g_cardPin,
32         PIN_SIZE) == 1) {
33         g_authenticated = 1; // Authentication();
34     }
35     return;
36 }
37
38 void main() {
39     initialize();
40     verifyPIN();
41     assert(g_authenticated == 1);
42 }
```

Figure 5.5: Running example, inspired by VerifyPIN [DPP+16]

```
3212 $ gcc -m32 -static -O0 -g verifyPIN_0.c
```

3215 The flag *m32* indicates to compile into a 32 bits binary, *static* means all depen-  
3216 dencies are statically linked which is required by the BINSEC tool, the optimization  
3217 level *O0* and the inclusion of debug symbols with the *-g* flag allow for easy mapping  
3218 between source code and binary program.

## 3219 5.4.2 Analysis Goal

3220 An analysis with BINSEC/ASE starts by defining what type of security property the  
3221 user wants to study and what type of attacker the program faces.

3222 Depending on the analysis goal, the exploration can be configured through a ded-  
3223 icated option to stop at the first attack found if any, or to explore exhaustively all  
3224 adversarial paths in the program. If the program is resistant to the considered attacker  
3225 model, the analysis being *k*-complete, it explores all paths up to a bounded depth to  
3226 ensure no bounded attack paths are missed. It is sometimes possible to increase the  
3227 depth bound to get the completeness of the analysis.

### 3228 5.4.2.1 Security Properties

3229 Security properties are application-specific and finding relevant ones is an active area  
3230 of research. It is not in the scope of this thesis. However, we provide the following few  
3231 pointers.

3232 Many different properties can be important to ensure in a program like authentica-  
3233 tion, data integrity or correct program output in general. When verifying an authen-  
3234 tication program, an interesting property is to make sure the password given is indeed  
3235 the reference password. For cryptographic programs, there may be mathematical for-  
3236 mulas that should remain true to avoid leaking the secret key like the BelCore attack  
3237 in CRT-RSA. The attacker's goal is to violate that security property.

3238 BINSEC/ASE supports properties that can be expressed with existing program  
3239 variables, or even ghost variables that the user updates through stubs. The security  
3240 oracle can also be written in DBA as part of the configuration file, then it can be  
3241 expressed with all DBA variables involved, and in particular it grants access to flags,  
3242 registers and memory locations. We can even imagine verifying trace properties. In the  
3243 examples we provide, the property is expressed with program variables as an 'assert'  
3244 in the program.

3245 Once the desired security property is established, there remains the question of the  
3246 location in the program where it should be true. It may be at the end of the program to  
3247 verify a program's output or in the middle, checking for memory access out of bounds  
3248 for instance. To increase the efficiency of the analysis, we suggest identifying locations  
3249 in the code from where the security property cannot be violated anymore to prune the  
3250 exploration. It can correspond to some program returns or calls to library functions  
3251 handling error cases.

3252 **Running example.** In the VerifyPIN program, we consider that the attacker goal is to  
3253 be authenticated despite not having the correct PIN. At the program level, it corre-  
3254 sponds to reaching the true branch of the assertion in the binary (address 0x08049e3a)  
3255 corresponding to line 41 in C. We place a *cut* at the false branch of the assertion,  
3256 at address 0x08049e1e, since the attacker cannot succeed past that point. Figure 5.6  
3257 shows the assembly code for the *main* function.

```

08049deb <main>:
8049deb:      endbr32
8049def:      lea    0x4(%esp),%ecx
8049df3:      and    $0xffffffff0,%esp
8049df6:      pushl  -0x4(%ecx)
8049df9:      push  %ebp
8049dfa:      mov    %esp,%ebp
8049dfc:      push  %ebx
8049dfd:      push  %ecx
8049dfe:      call  8049bc0 <__x86.get_pc_thunk.bx>
8049e03:      add    $0x9b1fd,%ebx
8049e09:      call  8049ce5 <initialize>
8049e0e:      call  8049da7 <verifyPIN>
8049e13:      movzbl 0x12dc(%ebx),%eax
8049e1a:      cmp    $0x1,%al
8049e1c:      je    8049e3a <main+0x4f>
8049e1e:      lea   -0x30fd4(%ebx),%eax
8049e24:      push  %eax
8049e25:      push  $0x2b
8049e27:      lea   -0x30ff8(%ebx),%eax
8049e2d:      push  %eax
8049e2e:      lea   -0x30fea(%ebx),%eax
8049e34:      push  %eax
8049e35:      call  804aff0 <__assert_fail>
8049e3a:      nop
8049e3b:      lea   -0x8(%ebp),%esp
8049e3e:      pop   %ecx
8049e3f:      pop   %ebx
8049e40:      pop   %ebp
8049e41:      lea   -0x4(%ecx),%esp
8049e44:      ret

```

Figure 5.6: Running example: disassembly of the main function

### 3258 5.4.2.2 Attacker Model

3259 Once the security property to verify has been defined, the relevant attacker model has  
3260 to be specified. Note that a very powerful attacker model can break almost any security  
3261 property, hence the need to choose the right strength level, after a risk analysis for the  
3262 application.

3263 **Attacker Capability and Attacker Budget.** Those two parameters define the strength  
3264 of the attacker model and can give an idea of the injection means at their disposal.  
3265 They should be selected according to the risk analysis previously performed.

3266 **Target Locations.** This corresponds to the set of instruction addresses that can be  
3267 faulted by the attacker. It is expressed as a list of intervals. For large programs, we  
3268 recommend focusing the target locations on the critical sections of code to get faster  
3269 results. In a single fault scenario, the analysis can be split with a different section  
3270 targeted each time. If the security oracle has been added to the code, we recommend

3271 excluding that part from the target locations.

3272 This can also help refine the attacker model. If an attacker model can only fault  
3273 some specific instructions, target locations can be enumerated exhaustively. For in-  
3274 stance, skipping only function calls can be modeled by listing as targets only the func-  
3275 tion call with an instruction skip fault model. Faulting data only on loads and stores,  
3276 at the memory interface, can be represented in this way too. In those two examples, it  
3277 is possible to design specific fault models, we refer the interested reader to Section 5.5.

3278 **Other Fault Modifiers.** To prevent faults on variables deemed secure, not interesting,  
3279 or with no efficient algorithm to handle them, a target blacklist can be used. We  
3280 mainly use it for the *esp* and *gs* registers in Intel x86 binaries. The *esp* register is the  
3281 stack pointer and we do not currently have an efficient algorithm to handle faulting it.  
3282 The *gs* register is used as a base pointer to point to a data region in some programs.  
3283 Faulting *gs* results in a combinatorial explosion as the base pointer for data can now be  
3284 taken anywhere in a data fault model. In general, we do not have an efficient algorithm  
3285 to handle faults on memory addresses, those can be filtered out with a configuration  
3286 option.

3287 By default, BINSEC/ASE doesn't fault flags as it would largely increase the number  
3288 of fault locations and new symbolic variables. Faults on 'normal' variables are most of  
3289 the time expressive enough with respect to the attacker model. However, it is possible  
3290 to activate faults on flags with a configuration option.

3291 **Running example.** Let's consider an attacker model able to perform arbitrary data  
3292 fault for instance. Since the VerifyPIN program does not have any protection, we  
3293 expect it to be vulnerable to an attacker able to inject one fault. For target loca-  
3294 tions, we exclude the *initialize* and *main* functions. This leaves the *verifyPIN* and  
3295 *byteArrayCompare* functions (assembly code in Figures 5.7 and 5.8), where we also  
3296 exclude the beginning which put the return address onto the stack, as well as the  
3297 end of the functions handling the return process. It leaves two ranges of addresses  
3298 to fault: 0x08049dba to 0x08049de6 in *verifyPIN* and 0x08049d5e to 0x08049da0 in  
3299 *byteArrayCompare*.

```
08049da7 <verifyPIN >:
8049da7:      endbr32
8049dab:      push    %ebp
8049dac:      mov     %esp,%ebp
8049dae:      push    %ebx
8049daf:      call   8049bc0 <__x86.get_pc_thunk.bx>
8049db4:      add     $0x9b24c,%ebx
8049dba:      movb   $0x0,0x12dc(%ebx)
8049dc1:      push    $0x4
8049dc3:      mov    $0x80e6c94,%eax
8049dc9:      push    %eax
8049dca:      mov    $0x80e6c84,%eax
8049dd0:      push    %eax
8049dd1:      call   8049d4a <byteArrayCompare>
8049dd6:      add     $0xc,%esp
8049dd9:      cmp    $0x1,%al
8049ddb:      jne    8049de5 <verifyPIN+0x3e>
8049ddd:      movb   $0x1,0x12dc(%ebx)
8049de4:      nop
8049de5:      nop
8049de6:      mov    -0x4(%ebp),%ebx
8049de9:      leave
8049dea:      ret
```

Figure 5.7: Running example: disassembly of the verifyPIN function

```
08049d4a <byteArrayCompare>:
8049d4a:    endbr32
8049d4e:    push   %ebp
8049d4f:    mov    %esp,%ebp
8049d51:    sub    $0x10,%esp
8049d54:    call   8049e45 <__x86.get_pc_thunk.ax>
8049d59:    add    $0x9b2a7,%eax
8049d5e:    movl   $0x0,-0x4(%ebp)
8049d65:    jmp    8049d98 <byteArrayCompare+0x4e>
8049d67:    mov    -0x4(%ebp),%eax
8049d6a:    lea   0x0(,%eax,4),%edx
8049d71:    mov    0x8(%ebp),%eax
8049d74:    add   %edx,%eax
8049d76:    mov   (%eax),%edx
8049d78:    mov   -0x4(%ebp),%eax
8049d7b:    lea   0x0(,%eax,4),%ecx
8049d82:    mov   0xc(%ebp),%eax
8049d85:    add   %ecx,%eax
8049d87:    mov   (%eax),%eax
8049d89:    cmp   %eax,%edx
8049d8b:    je    8049d94 <byteArrayCompare+0x4a>
8049d8d:    mov   $0x0,%eax
8049d92:    jmp   8049da5 <byteArrayCompare+0x5b>
8049d94:    addl  $0x1,-0x4(%ebp)
8049d98:    mov   -0x4(%ebp),%eax
8049d9b:    cmp   0x10(%ebp),%eax
8049d9e:    jl    8049d67 <byteArrayCompare+0x1d>
8049da0:    mov   $0x1,%eax
8049da5:    leave
8049da6:    ret
```

Figure 5.8: Running example: disassembly of the byteArrayCompare function

### 5.4.3 Configuration

For detailed configuration options, we refer the interested reader to the online documentation and examples provided as part of the BINSEC/ASE artifact<sup>9</sup>.

**Running Example.** Figure 5.9 shows what a configuration file for the VerifyPIN program would look like, with the described attacker model. The memory file *mem* only contains a starting address for the stack pointer. An appropriate value for the maximal exploration depth *depth* is discussed in Section 5.4.5.

```
esp <32> := 0xFFFFFFFF00;
```

```
[kernel]
isa = x86
entrypoint = main
file = verifyPIN_0.x

[sse]
enabled = true
depth = 200
memory = mem

fault-model = ArbitraryData
max-faults = 1
goal-address = 0x08049e3a
assert-fail-address = 0x08049e1e
target-addresses = 0x08049dba,0x08049de6,0x08049d5e,0x08049da0
injection-method = on-demand
target-blacklist = esp
where-check-nb-faults = branch
```

Figure 5.9: Running example: BINSEC/ASE configuration file

### 5.4.4 Reading BINSEC/ASE Output

We base this section on the output of BINSEC/ASE for the VerifyPIN running example.

**Run Statistic.** The output ends with some analysis statistics (see Figure 5.10). For instance, We can see on line 2 that the analysis explored 5 different paths, analyzed 206 instructions in total (line 5) and the longest path is 163 instructions deep (line 6).

To determine if the analysis was complete (as opposed to k-complete), the following statistics should be checked:

- The number of ‘cut early paths’ line 8 should be 0;
- The number of SMT queries the solver answered ‘unknown’ for should be 0 (line 20).

The ‘SMT queries’ section provides a number of information about queries and reliance on the SMT solver. It gives an idea about the query complexity handled by the analysis.

<sup>9</sup>[https://github.com/binsec/esop2023\\_artefact](https://github.com/binsec/esop2023_artefact)

- 3323 – Line 11, we see that 61 queries were generated but only 13 of them needed to be  
3324 sent to the solver (line 12). The difference has been arithmetically simplified into  
3325 true or false (lines 18 and 19 for assume queries);
- 3326 – The ‘Average solving time’ (line 13) gives in seconds the average time the solver  
3327 took to answer per query. Here, our tool wasn’t precise enough to record the  
3328 small value, hence the display of a 0.

3329 Among the ‘Fault injection’ statistics, we have a reminder of the selected attacker  
3330 model options (lines 23, 25, 26 and 27) and chosen optimizations (lines 24 and 28).  
3331 The total number of injections performed is indicated in line 29. In the end, resides  
3332 overall analysis results:

- 3333 – ‘Models found’ (line 33) indicates the number of adversarial paths produced by  
3334 BINSEC/ASE;
- 3335 – ‘Assert cut’ (line 34) relates to the number of paths where the attacker goal was  
3336 not adversarially reachable.

3337 **Attack Path.** We explore in Figure 5.11 one of the produced attack paths to under-  
3338 stand how to read that part of the BINSEC/ASE output. Line 1 indicates that the  
3339 goal address 0x08049e3a has been reached and the solver proposed the model follow-  
3340 ing. Only 3 different faults were injected in that path and the fault at the address  
3341 0x08049d5e is the activated one (line 4). It corresponds to the initialization of the  
3342 loop counter  $i$ , equivalent to line 20 in Figure 5.5. The solver proposed to inject the  
3343 value 0x7ffffffc (line 9) instead of the original value 0. Since 0x7ffffffc is strictly greater  
3344 than PIN\_SIZE, the program does not enter the loop and returns 1 as if all digits were  
3345 correct. Hence, injecting this fault allows the attacker to reach its goal.

```
1 [sse:info] Exploration
2     Paths 5
3     Paths continuing after max reached 0
4     Secondary queries (EDS only) 0
5     Analysed instructions 206
6     Max depth instruction 163
7     Max depth branchments 16
8     Cut early paths 0
9
10 [sse:info] SMT queries
11     Number of queries 61
12     Number of queries sent to the solver 13
13     Average solving time 0.000000
14     Total assume queries 49
15     Assume queries sent to solveur 13
16     Total enumerate queries 12
17     Enumerate queries sent to solveur 0
18     Trivial true assume 18
19     Trivial false assume 18
20     Unknown solver 0
21
22 [sse:info] Fault injection
23     Fault model ArbitraryData
24     Injection method on-demand
25     Maximum number of faults 1
26     Where number of faults checked: branch
27     Where non_det constrains checked: branch
28     At branch optim: false
29     Injection locations 15
30     Min Injection locations per path 0
31     Max Injection locations per path 15
32     Avg Injection locations per path 11
33     Models found 3
34     Assert cut 2
```

Figure 5.10: Running example: BINSEC/ASE statistics output

```

1 [sse:result] Model 08049e3a
2   —— Model ——
3   # Variables
4   b_8049d5e : #b1
5   b_8049d98 : #b0
6   b_8049dba : #b0
7   ebp : —
8   ebx : —
9   non_det_8049d5e : #x7fffffff c
10  non_det_8049d98 : #x7fffffff d
11  non_det_8049dba : #x00
12
13  —— empty memory ——

```

Figure 5.11: Running example: BINSEC/ASE attack path

### 3346 5.4.5 Analysis Process

3347 We describe here the recommended process when adding a new benchmark, after the  
3348 security property and attacker model definition step has been performed. This is an  
3349 iterative process to ensure correct configuration and early results.

3350 As configuring a BINSEC/ASE analysis may be a trial-and-error process, we recom-  
3351 mend starting with a non-faulted run. This helps to see if there is any warning or error  
3352 to resolve. If an attack is found, it should be investigated. Either there is an existing  
3353 bug in the program or the configuration is not adapted. Other statistics can be read  
3354 and examined, like the number of explored paths or solver statistics. As a convenience,  
3355 we use the maximal depth of a non-faulted run, rounded up to the hundred as a bound  
3356 for faulted runs, this metric can be recovered at this step. The default maximal depth  
3357 setting may need to be increased to adapt to the program.

3358 The next step is to add the attacker model to the configuration. The analysis is  
3359 likely to take longer and explore more paths. We recommend starting from a rather  
3360 weak attacker model, able to perform only one fault, then analyzing the BINSEC/ASE  
3361 run, before incrementally strengthening the attacker model until the desired one. This  
3362 incremental process enables the detection of any abnormal behavior, warnings or errors  
3363 before the analysis gets too complex. If an attack is found for a weaker attacker model  
3364 than the desired one, the program is vulnerable. The attack can be analyzed and  
3365 possibly protected against.

### 3366 5.4.6 Summary

3367 We presented in this section a methodology to set up a new program for a BINSEC/ASE  
3368 analysis. Here is a summary of the different steps.

- 3369 1. Express the security properties of the program;
- 3370 2. Decide on an attacker model relevant to the program;
- 3371 3. Write the configuration file;
- 3372 4. Run BINSEC/ASE without faults to verify the configuration;
- 3373 5. Incrementally increase the strength of the attacker model until the desired one.

3374 Now that the program analysis is set up as intended, the results can be integrated

3375 into the bigger picture, for instance, a security evaluation by a security practitioner.

- 3376 – If the program was vulnerable when it was not supposed to be, a strengthening
- 3377 process can begin, adding countermeasures and starting the analysis process pre-
- 3378 sented above, iterating until the analysis results meet the security requirements;
- 3379 – If the program proves resistant to the considered attacker model, it can be in-
- 3380 teresting to consider other properties the program should have, and start the
- 3381 analysis process again with this new attacker model.

## 3382 5.5 Developer Guide: a Methodology to Add a New Fault 3383 Model

3384 The objective of this section is to present a methodology to add a new fault model  
3385 inside BINSEC/ASE. We acknowledge that there exist many different fault models.  
3386 We implemented some in this thesis as proof of concepts, but others may be more  
3387 relevant to a specific application. This guide is meant to provide a starting point for  
3388 interested developers and does not aim to be exhaustive.

3389 **Running Example.** In this section, the instruction skip fault model is the running  
3390 example. It was the last to be implemented and required modifications in many parts  
3391 of the symbolic engine.

### 3392 5.5.1 Defining the New Fault Model

3393 The first step to adding a new fault model is to design it on paper. We discuss it in  
3394 this section.

#### 3395 5.5.1.1 State-of-the-art Attack

3396 A fault model is usually derived from a capability given to an attacker by a real attack  
3397 that a security practitioner wishes to simulate. To represent this attack, a developer  
3398 needs to understand it in terms of:

- 3399 – what type of capability it gives to an attacker (what action on what object);
- 3400 – when it can be applied in the execution;
- 3401 – in which part of the system the fault is introduced;
- 3402 – does multiple faults make sense and if so how.

3403 **Running Example.** Instruction skip is a common model used to describe the effects of  
3404 hardware fault injection on programs. In practice, instructions can be modified so they  
3405 don't have side effects, or the prefetch buffer may be tempered with. A skip can occur  
3406 on any ISA instruction. Multiple separated faults are possible in the state-of-the-art,  
3407 as well as skipping multiple consecutive instructions.

#### 3408 5.5.1.2 Fault Model

3409 The details of the attacker capability to model have to be expressed in terms of their  
3410 impact on the path predicate computation. That is to say, how they modify the  
3411 path predicate initialization, the adversarial reachability query and the behavior of the  
3412 four main DBA instructions: assigns, static, conditional and dynamic jumps. At this  
3413 step, the capability is often over-approximated for simplicity. It is possible to imagine  
3414 complex fault models with interactions between different instructions. It can be the

3415 case when considering precisely the micro-architecture [TAC<sup>+</sup>22]. It would require  
3416 keeping track of a trigger flag in the symbolic state.

3417 **Running Example.** We modeled instruction skip as not executing an instruction en-  
3418 tirely, not updating the memory, the registers or the flags. Forkless instruction skip  
3419 has the following impact on the path predicate computation (see Section 5.3.3.3 for  
3420 more details).

- 3421 – *Assign*: a data fault is injected, with the corrupted value being the old value of  
3422 the variable.
- 3423 – *Static jump*: this instruction is transformed into a conditional jump, with the  
3424 fault activation as the condition, the next instruction directly below in addresses  
3425 as the target for the true branch, and the original target for the else branch.
- 3426 – *Conditional jump*: the condition is augmented with an activation variable, that if  
3427 true, the analysis goes to the ‘then’ branch which corresponds to the next address  
3428 in the memory layout.
- 3429 – *Dynamic jump*: the expression providing the possible jump targets is augmented  
3430 similarly to a data fault, adding the next instruction address as a jump target  
3431 possibility.

3432 Here, we see an example of a fault model modifying all the DBA instructions, it is  
3433 not necessarily the case. For instance, a new kind of data fault would only modify the  
3434 assignment process.

3435 We choose to implement multiple distinct faults to represent an attacker with a  
3436 multiple instruction skip capability, with the existing machinery (number of faults  
3437 check in particular). We believe implementing a fixed number of consecutive faults can  
3438 be done by storing fault activation variables and a timer in the path predicate, adding  
3439 with a logical *and* the stored activation variables to the current instruction’s activation  
3440 and decreasing the timer at each time until their timer reaches 0.

## 3441 5.5.2 Implementation

3442 Once the plan for which instruction handler needs to be modified and how, it can be  
3443 implemented inside BINSEC/ASE. A new fault model may need other adjustments,  
3444 such as:

- 3445 – the initialization process,
- 3446 – the directives handlers,
- 3447 – the fault location filtering process already implemented may work or may need  
3448 to be modified,
- 3449 – it may also require additional internal variables stored in the path predicate, like  
3450 counters, we have some for injected faults, or boolean flags, indicating a fault  
3451 saturation for instance,
- 3452 – queries sent to the SMT solver may be impacted in their form (adding new  
3453 elements) and in which part of the query should be stored as a new constraint in  
3454 the path predicate or not,
- 3455 – the new fault model has to be added to the configuration option to allow a user  
3456 to select it.

3457 We provide Ocaml interfaces in BINSEC/ASE for the elements listed above in Table  
3458 5.7.

3459 For comparison, we implemented a forking version of all our forkless fault models.  
3460 If the developer wishes to do the same, a new fault model needs to be designed, as the  
3461 impact of the forking fault model on the path predicate is different.

Table 5.7: Ocaml interfaces

Elements to modify	BINSEC/ASE Ocaml interface
Initialization process	src/ee/sse.ml:do_sse
Directive handlers	src/ee/sse.ml:[handle_assumptions, handle_reach], src/ee/sse.ml:[handle_enumerate, handle_cut]
Fault filter	src/ee/sse.ml:is_fault_location
Path predicate	src/ee/sse_types.[ml, mli]:Path_state
SMT queries	src/ee/senv.[ml, mli]:[assume, enumerate]
Options	src/ee/sse_options.[ml, mli]
Exploration metrics	src/ee/sse.ml:Stats
Solver metrics	src/ee/senv.ml:Query_stats

### 3462 5.5.3 Dedicated Metrics

3463 We recommend implementing dedicated metrics to monitor the new fault model. It  
 3464 can be the number of times it is triggered, a notion of where it is triggered, etc. This  
 3465 helps the debugging process and provides more details in the BINSEC/ASE output  
 3466 for further program or attack study. We distinguish two types of metrics, exploration  
 3467 metrics and solver metrics. Interfaces for them are listed in Table 5.7.

3468 **Running Example.** For instance, the instruction skip fault model comes with counters  
 3469 for how many of each kind of DBA instruction has been faulted.

### 3470 5.5.4 Testing

3471 Implementing a new fault model is not trivial and requires testing to ensure the correct-  
 3472 ness of the implementation. We recommend writing a few minimal testing programs,  
 3473 each showcasing a different aspect of the fault model, for one and two faults. Then,  
 3474 configure them, run the BINSEC/ASE analysis and analyze manually the results. This  
 3475 means examining the number of paths explored and ensuring it is correct, seeing where  
 3476 queries have been created and that it is reasonable, verifying the total number of injec-  
 3477 tion points, etc. And study the dedicated metrics added to be sure to understand what  
 3478 happens. BINSEC comes with different levels of debug logs which can be very helpful  
 3479 in understanding the unrolling of the analysis. If the implementation of the new fault  
 3480 model was somewhat ‘invasive’ and could potentially disturb other fault models, we  
 3481 recommend performing non-regression tests.

### 3482 5.5.5 Summary

3483 We presented in this section a methodology to add a new fault model. Here is a  
 3484 summary of the different steps.

- 3485 1. Define what the fault model should represent (a type of attack for instance);
- 3486 2. Define the fault model in terms of how it transforms DBA instructions;
- 3487 3. Implement those transformations and other alterations to the symbolic engine;
- 3488 4. Implement dedicated metrics;
- 3489 5. Create testing programs and evaluate them to ensure the correctness of the im-  
 3490 plementation.

3491 When all of those steps are done, the new fault model is ready to be used in a real  
 3492 security scenario.

## 3493 5.6 Discussion

3494 We discuss in this section considerations on what can or cannot be represented with  
3495 BINSEC/ASE, in particular limits to fault model support.

### 3496 5.6.1 BINSEC/ASE Limitations

3497 We discuss here the main limitations of the BINSEC/ASE prototype.

#### 3498 5.6.1.1 Building on Top of an Existing Tool

3499 The prototype presented in this chapter is based on BINSEC version 0.4.0, forked  
3500 September 1<sup>st</sup> 2021. BINSEC underwent a major refactoring shortly after, which made  
3501 rebasing impossible, only a full re-implementation could have worked. This was not  
3502 done for time reasons. Hence BINSEC/ASE does not benefit from recent developments  
3503 in configuration syntax readability, architecture support (Intel x86-64 bits for instance)  
3504 and other continued improvements. We leave as future work re-implementing the FASE  
3505 engine in the current BINSEC version.

#### 3506 5.6.1.2 Supported Fault Models

3507 **Faults on Addresses.** We do not have an efficient algorithm for faults on addresses. As  
3508 many values are possible, the solver can lose time considering them and the possible  
3509 effects of their propagation in the following instructions. Another issue is that the  
3510 solver can select an address in an uninitialized part of the memory and fill it with  
3511 exactly the convenient content, which is not very realistic. Hence, we chose to avoid  
3512 faulting values ‘looking like’ address, with a configurable threshold. In practice, it also  
3513 limits our ability to study faults related to stack manipulations.

3514 **Faults on Instruction Op-Code.** They can be either permanent faults when the code  
3515 memory is corrupted, or transient faults when an element of the micro-architecture  
3516 like the instruction buffer is attacked. The case of permanent VS transient faults  
3517 is discussed in Section 5.6.3. Changing one instruction into another in the general  
3518 sense is a hard problem in symbolic execution as it relates to self-modifying code,  
3519 inducing a state explosion to consider all possible options for each instruction. We  
3520 do not have an efficient forkless algorithm for faulting instruction op-codes. However,  
3521 some fault models can be considered as specific cases of instruction corruption like  
3522 the test inversion fault model where a conditional jump (jump-if-equal for instance)  
3523 is transformed into its opposite (jump-if-not-equal in that example). Instruction skips  
3524 can also be seen as a specific case of instruction corruption, changing an instruction to  
3525 a NOP.

3526 **Combination of Fault Models.** BINSEC/ASE does not currently support the possibility  
3527 to combine fault models in one multi-fault adversarial path. It would make sense in  
3528 a security scenario considering an advanced attacker able to leverage different attack  
3529 vectors in one attack. We have not imagined a more efficient algorithm than a naive  
3530 combination of fault models into one big encoding. This naive technique could be  
3531 implemented without much effort and we leave as future work to design an efficient  
3532 combined injection technique.

3533 **Generic Support for Fault Models.** We acknowledge that we implemented only a  
3534 limited number of fault models having specific effects in this prototype. The aim was

3535 to showcase the ability of our tool to support various fault models that can represent  
3536 existing security scenarios. The interested user can extend BINSEC/ASE with new  
3537 fault models more suitable for their analysis context, based on our developer guide  
3538 (Section 5.5).

### 3539 **5.6.2 Faults on Intermediate Representation**

3540 Working on intermediate representation allows to abstract from specific ISAs and to  
3541 have a generic analysis. The semantics are preserved, however, it changes the syntax.  
3542 This becomes meaningful when considering faults modifying the binary encoding of  
3543 instructions, for instance flipping a bit in the op-code or switching one register for  
3544 another. Those faults require micro-architectural details to be known to the analysis,  
3545 at least through a mapping of ISA op-codes to DBA instructions. This is another  
3546 difficulty hindering the support of generic faults on instructions by BINSEC/ASE. Only  
3547 restrictions that can be expressed with DBA and that can be reasoned about without  
3548 micro-architectural information (instruction skip, test inversion, operand modification)  
3549 have been implemented. We also miss micro-architectural considerations such as fetch  
3550 and prefetch of instructions, without which we cannot model faults happening in those  
3551 components.

### 3552 **5.6.3 Permanent VS Transient Faults**

3553 In this work, we only consider transient faults, i.e. a corruption that lasts until the value  
3554 is overwritten. We believe permanent faults require micro-architectural information to  
3555 precise exactly which bit of which physical component is affected, and then propagate  
3556 it to the ISA level. Moreover, memory is often encrypted and permanent damage in  
3557 it can be hard to reason with. Some restrictions on permanent faults can likely be  
3558 implemented, such as permanent faults on a register, which is available at DBA level.



3559 Chapter **6**

3560 Experimental Evaluation

3561 **Contents**

---

3562	<b>6.1 Evaluation Overview . . . . .</b>	<b>104</b>
3563		
3564	6.1.1 Research Questions . . . . .	104
3565	6.1.2 Experimental Setting . . . . .	105
3566	6.1.2.1 Benchmark . . . . .	105
3567	6.1.2.2 Attacker Model . . . . .	106
3568	6.1.2.3 Competitor . . . . .	106
3569	6.1.2.4 Experimental Setup . . . . .	106
3570	6.1.3 Artifact Availability . . . . .	107
3571	<b>6.2 Correctness and K-completeness (RQ1) . . . . .</b>	<b>107</b>
3572	<b>6.3 FASE Evaluation for Arbitrary Data Faults (RQ2) . . . . .</b>	<b>108</b>
3573	6.3.1 Scalability (RQ2.1) . . . . .	108
3574	6.3.2 Impact of Optimizations (RQ2.2) . . . . .	110
3575	6.3.3 Comparison of the Different Forkless Encodings (RQ2.3) . . . . .	116
3576	<b>6.4 FASE Evaluation of Other Fault Models (RQ3) . . . . .</b>	<b>118</b>
3577	6.4.1 FASE Evaluation of Reset Faults (RQ3.1) . . . . .	119
3578	6.4.2 FASE Evaluation of Bit-Flip Faults (RQ3.2) . . . . .	120
3579	6.4.3 FASE Evaluation of Test Inversion Faults (RQ3.3) . . . . .	121
3580	6.4.4 FASE Evaluation of Instruction Skip Faults (RQ3.4) . . . . .	122
3581	6.4.5 Summary . . . . .	123
3582	<b>6.5 Forkless Faults in Instrumentation (RQ4) . . . . .</b>	<b>124</b>
3583	6.5.1 Experimental Settings . . . . .	125
3584	6.5.2 Scalability . . . . .	125
3585	6.5.3 Conclusion . . . . .	126
3586	<b>6.6 Security Scenarios . . . . .</b>	<b>127</b>
3587	6.6.1 CRT-RSA . . . . .	128
3588	6.6.2 Secret-keeping Machine . . . . .	129
3589	6.6.3 SecSwift Countermeasure . . . . .	130

3590	6.6.4	Neural Network . . . . .	131
3591	6.6.5	Security Scenarios Feedback . . . . .	132
3592	<b>6.7</b>	<b>Case Study: WooKey Bootloader . . . . .</b>	<b>132</b>
3593	6.7.1	Presentation of WooKey . . . . .	132
3594	6.7.2	Security Scenario and Goal of our Study . . . . .	133
3595	6.7.3	Analyze Key Parts of Wookey . . . . .	133
3596	6.7.4	Analyze a Security Patch of WooKey . . . . .	136
3597	6.7.5	Propose a New Patch and Evaluate It . . . . .	136
3598	6.7.6	Other Attacks on WooKey . . . . .	137
3599	6.7.6.1	Attack Vectors Combination . . . . .	137
3600	6.7.6.2	Faulty Redundant Test . . . . .	137
3601	6.7.7	Case Study Conclusion . . . . .	138

---

3602  
3603

3605 This chapter is dedicated to the experiment evaluation of the approach presented  
3606 in this thesis through the BINSEC/ASE prototype described in Chapter 5. After  
3607 presenting our experimental setting (Section 6.1), we confirm the correctness and k-  
3608 completeness of our BINSEC/ASE prototype (Section 6.2) and systematically evaluate  
3609 its performance against our implementation of the forking technique representing the  
3610 state-of-the-art approach and compared to our various optimizations for arbitrary data  
3611 faults (Section 6.3) and for the other faults models (Section 6.4). We also investigate  
3612 the advantages of using a forkless fault encoding in instrumentation (Section 6.5).  
3613 Secondly, we explore various security scenarios (Section 6.6) and a case study (Section  
3614 6.7) to showcase the interest and feasibility of our technique.

## 3615 6.1 Evaluation Overview

3616 In this section, we introduce the research questions driving this experimental evaluation  
3617 of our technique, FASE (Forkless Adversarial Symbolic Execution), implemented in  
3618 the BINSEC/ASE tool. We also detail the experimental settings and present the  
3619 benchmarks used for the evaluation.

### 3620 6.1.1 Research Questions

3621 We aim to answer the following research questions during this evaluation.

3622 **RQ 1: Correctness and K-completeness.** FASE algorithm has been shown to be correct  
3623 and k-complete for adversarial reachability (see Section 4.3.3.3). We start by checking  
3624 that our FASE implementation preserves those properties in Section 6.2.

3625 **RQ 2: FASE Evaluation for Arbitrary Data Faults.** In Section 6.3, we evaluate the  
3626 performance of our technique compared to the forking technique implemented alongside  
3627 FASE in BINSEC/ASE. We also compare the performance of our optimizations and  
3628 forkless encoding variations through the following research questions.

- 3629 – **RQ 2.1: Scalability.** We verify that FASE is able to scale in number of faults  
3630 without path explosion, better than the forking technique (Section 6.3.1);

- 3631 – **RQ 2.2: Impact of Optimizations.** We investigate the impact of our optimizations  
 3632 on the performance of FASE and wonder which is the fastest optimization (Section  
 3633 6.3.2);
- 3634 – **RQ 2.3: Forkless Encodings.** We investigate the impact of the operators used  
 3635 for the forkless encoding and wonder which one does result in the fastest analysis  
 3636 (Section 6.3.3).

3637 **RQ 3: FASE Evaluation for Other Fault Models.** We evaluate the impact of the other  
 3638 implemented fault models: reset (RQ3.1 Section 6.4.1), bit-flip (RQ3.2 Section 6.4.2),  
 3639 test inversion (RQ3.3 Section 6.4.3) and instruction skip (RQ3.4 Section 6.4.4). For  
 3640 each, we verify that we can scale in number of faults without path explosion, compared  
 3641 to the forking technique.

3642 **RQ 4: Forkless Faults in Instrumentation.** In Section 6.5, we evaluate the performance  
 3643 of forkless encodings compared with forking ones in source-level instrumentation. The  
 3644 instrumentation process is described in Appendix C. We check that we can scale in  
 3645 number of faults without path explosion, compared to the forking technique in instru-  
 3646 mentation.

## 3647 6.1.2 Experimental Setting

3648 We now present the benchmark used for this performance evaluation and detail our  
 3649 experimental settings.

### 3650 6.1.2.1 Benchmark

3651 **Performance Benchmarks.** We used here a standard set of programs from the SWIFI  
 3652 literature on physical fault injections and high-security devices. This constitutes a  
 3653 benchmark of 12 programs, characterized in Table 6.1.

- 3654 – Eight versions of VerifyPIN from the FISSC [DPP<sup>+</sup>16] benchmark suite, ded-  
 3655 icated to the evaluation of physical fault attack analyses. VerifyPIN is a toy  
 3656 authentication program. There are one unprotected and 7 different protected  
 3657 versions, some vulnerable, some resistant to one test inversion fault. An oracle is  
 3658 provided by FISSC, checking if the user PIN truly corresponds to the reference  
 3659 PIN;
- 3660 – Two manually unrolled versions of the unprotected VerifyPIN, with a PIN size  
 3661 of 4 and 16, to add diversity in the benchmarks with programs without loops;
- 3662 – Two versions of the npo2 program from Le *et al.* [LHGD18], together with their  
 3663 oracles. Npo2 is a program computing an integer’s upper power of two. The  
 3664 attacker’s goal is to perform a silent data corruption, i.e. change the end result  
 3665 without triggering countermeasures. One version is vulnerable to one arbitrary  
 3666 data fault, the second is resistant due to extra arithmetic checks.

3667 In the remainder of this chapter, this set of programs is denoted as the *performance*  
 3668 *benchmarks*.

3669 **Compilation.** The benchmarks are written in C and have been compiled with gcc for  
 3670 the Intel x86-32 architecture. This was the first ISA supported by BINSEC/ASE and  
 3671 we expect other ISAs to give the same results as only the lifting of the ISA to DBA  
 3672 changes, not the adversarial symbolic execution itself. We use the flag “-O0” to preserve  
 3673 countermeasures, and subsequently to improve assembly readability. For BINSEC com-  
 3674 patibility, we use the “-static” flag to include the necessary library functions directly in

Table 6.1: Benchmarks characteristics and statistics of a standard SE analysis

Program group (#)	C loc	x86 loc	BINSEC analysis - no fault			Time
			#instructions (explored)	#paths	max #branch in a path	
Sections 6.2 to 6.5						
VerifyPINs (8)	80-140	160-215	192-269	1	17-34	< 0.1s
VerifyPIN unrolled (2)	40-85	140-430	142-442	5-17	5-17	< 0.1s
npo2 (2)	50	200-220	607-653	3	31-33	< 0.1s

3675 the binary.

### 3676 6.1.2.2 Attacker Model

3677 The attacker model chosen in this evaluation can perform a varying number of faults.  
 3678 Its goal is expressed as a security oracle directly written in C for each benchmark,  
 3679 the computation of which is not faulted. Specific details are provided for the different  
 3680 experiments when relevant. In particular, the exact fault models used are indicated.  
 3681 Except when explicitly stated otherwise, we do not fault addresses, nor fault directly  
 3682 the stack pointer or the program counter, and we do not fault the status flags.

### 3683 6.1.2.3 Competitor

3684 We chose to represent the state-of-the-art work with the forking technique, which scales  
 3685 better than mutant generation techniques for multiple faults. The forking technique  
 3686 split the explored path into two at each possible fault location, creating one containing  
 3687 the fault, and keeping one without the fault at that particular location. In order to  
 3688 only compare encodings we implemented forking faults encodings inside BINSEC/ASE,  
 3689 along with our new forkless ones, as described in Chapter 5. Note that the forking  
 3690 technique induces path splits in the analysis exploration at fault locations, it does not  
 3691 explicitly fork processes. We compare our approach to others on a more general level  
 3692 in Section 4.5.

### 3693 6.1.2.4 Experimental Setup

3694 **Machine Used.** We ran our experiments on a cloud machine with a processor Intel  
 3695 Dual Xeon 4214R with 48 CPU cores and 384GB of RAM. Experiments ran in parallel  
 3696 on the 48 cores, each run using only one core.

3697 **BINSEC settings.** We limit the maximal depth of an analysis to the depth necessary  
 3698 to perform an exhaustive non-faulty analysis, rounded to the upper hundred. We may  
 3699 not always be complete, which is not necessary for a performance evaluation. We  
 3700 exhaustively explore all the possible adversarial paths up to this bound in order to  
 3701 have comparable results. We set the global analysis timeout for 1 day.

### 3702 6.1.3 Artifact Availability

3703 Our benchmark infrastructure, case studies and the executable of BINSEC/ASE have  
 3704 been made available through artifacts for reproducibility purposes on GitHub<sup>1</sup> and  
 3705 Zenodo<sup>2</sup>. The BINSEC/ASE source code will be open-sourced.

## 3706 6.2 Correctness and K-completeness (RQ1)

3707 We start with a basic sanity check, in order to check whether our implementation  
 3708 indeed is correct and k-complete.

3709 **RQ1.** Is our tool correct and k-complete?

3710 **Goal.** The goal of this first research question is to check that our tool works as expected  
 3711 on several codes with known ground truth. In particular, we show that we are able  
 3712 to find attacks on vulnerable programs and provide a complete analysis showing the  
 3713 absence of vulnerabilities in secure programs.

3714 **Protocol.** The performance benchmarks are used in this experiment. We consider an  
 3715 attacker model able to perform either 1 or 2 arbitrary data faults, or 1 or 2 test inversion  
 3716 faults. To study the correctness and completeness of BINSEC/ASE, we proceed as  
 3717 follows.

- 3718 1. We check that indeed, with no fault allowed, no attack is found in any of the  
 3719 benchmarks. We make sure that the analysis is complete for this setting;
- 3720 2. For each program of the benchmark, we get the ground truth provided by their  
 3721 authors, setup the same attacker model, and observe BINSEC/ASE output.

Table 6.2: FASE vulnerability results compared to benchmark ground truth (RQ1)

AD: arbitrary data, TI: test inversion

V: vulnerable, S: secure, -: unspecified in the reference paper

Program	Ground truth / FASE result			
	1 TI	2 TI	1 AD	2 AD
np02 c1 (insecure)	- / S	- / S	V / V	V / V
np02 c2 (secure)	- / S	- / S	S / S	S / S
VerifyPIN 0	V / V	V / V	- / V	- / V
VerifyPIN 1	V / V	V / V	- / V	- / V
VerifyPIN 2	V / V	V / V	- / V	- / V
VerifyPIN 3	V / V	V / V	- / V	- / V
VerifyPIN 4	V / V	V / V	- / V	- / V
VerifyPIN 5	S / S	V / V	- / V	- / V
VerifyPIN 6	S / S	V / V	- / V	- / V
VerifyPIN 7	S / S	V / V	- / V	- / V
VerifyPIN unrolled size 4	S / S	S / S	- / V	- / V
VerifyPIN unrolled size 16	S / S	S / S	- / V	- / V

3722 **Results.** Results are summarized in Table 6.2. Vulnerability results are the same for  
 3723 FASE and its optimizations, as well as for the forking technique. The exploration of our

<sup>1</sup>[https://github.com/binsec/esop2023\\_artefact](https://github.com/binsec/esop2023_artefact)

<sup>2</sup><https://zenodo.org/record/7507112>

3724 analysis was either complete or at least an adversarial path was found, demonstrating  
 3725 the vulnerability of the program. As expected, the insecure npo2 program is vulnerable  
 3726 to a single arbitrary data fault while the secure version is not, nor is it vulnerable to  
 3727 two arbitrary data faults. According to their authors, the VerifyPIN versions 0 to 4 are  
 3728 vulnerable to one test inversion, while VerifyPIN 5 to 7 are resistant to it. We indeed  
 3729 reproduce these results. When allowing two faults, all VerifyPINs become vulnerable.  
 3730 When using one arbitrary data fault against the VerifyPINs, all versions are found  
 3731 vulnerable. We manually check that indeed the identified attack paths make sense.  
 3732 Our manually unrolled versions of VerifyPINs do not contain conditional branching  
 3733 instructions in the targeted function, making them resistant to test inversion. We  
 3734 check that this is the case, while they are still vulnerable to a single arbitrary data  
 3735 fault.

**Conclusion RQ1.** Our implementation of FASE preserves its correctness and k-completeness properties on our benchmark. In particular, we can show a program’s vulnerability to fault injection attacks and prove a program secure when it is.

## 3736 6.3 FASE Evaluation for Arbitrary Data Faults (RQ2)

3737 Arbitrary data faults are the most complex fault model BINSEC/ASE implements with  
 3738 regard to the number of added symbolic variables. One determines the activation of  
 3739 a fault and the second the corrupted value, and that for each possible fault location,  
 3740 encompassing all assignments. Hence, characterizing FASE for arbitrary data faults is a  
 3741 form of worst-case evaluation, and also where the efficiency gain is the most important.

3742 We only consider arbitrary data faults in this section. First, we compare scaling  
 3743 capabilities of FASE against the forking technique in terms of number of faults. Then,  
 3744 we want to show which optimization results in the fastest analysis and what their  
 3745 impact on the analysis is. We also evaluate the impact of the forkless encoding operator  
 3746 used and show which is the fastest. We consider optimizations and encodings to be  
 3747 independent and evaluate them separately. As a reminder, the analyses are set to  
 3748 explore exhaustively all program paths within the 24-hour time limit.

### 3749 6.3.1 Scalability (RQ2.1)

3750 First, we show the general performance improvement of FASE over the forking tech-  
 3751 nique, especially as more faults are considered. We selected our best performing opti-  
 3752 mization, FASE-IOD, as shown in Section 6.3.2, to represent FASE.

3753 **RQ2.1.** Can we scale in number of faults without path explosion compared to the  
 3754 forking technique (for arbitrary data faults)?

3755 **Goal.** The goal of this experiment is to check whether our technique is able to prevent  
 3756 path explosion, especially when the number of faults increases, and that it translates  
 3757 to improved analysis time.

3758 **Protocol.** The performance benchmark is used in this experiment. We consider an  
 3759 attacker model able to perform 1 to 10 arbitrary data faults.

3760 **Results.** We start by looking at the analysis time to compare FASE and the forking  
 3761 technique, as analysis time conditions the usability of a tool. Both techniques have  
 3762 the same exploration strategy, analyze all possible paths, and can be compared fairly.

3763 Analysis time results are presented in Table 6.3 (p.109) and illustrated in Figure 6.1  
 3764 (p.109). First, we see that the forking technique starts to timeout for half of the  
 3765 benchmarks for an attacker model able to perform 3 faults, and timeout for all for 8  
 3766 faults and above. FASE never timeouts in this experiment. FASE is on average x19  
 3767 times faster for 1 fault and x403 times faster for 2 faults. Values for 3 faults and above  
 for the forking technique are for incomplete runs, hence not comparable.

Table 6.3: FASE-IOD and forking analysis time comparison for arbitrary data faults (RQ2.1)

		Analysis time (s)						
		1f	2f	3f	4f	6f	8f	10f
		average value						
FASE-IOD		2.08	8.88	33.4	100	467	909	1.04k
Forking*		39.5	3.58k	8.17k	9.15k	35.6k	86.4k	86.4k
		timeouts (24h) over 12 benchmarks in total						
FASE-IOD		0	0	0	0	0	0	0
Forking		0	0	6	9	11	12	12

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts)

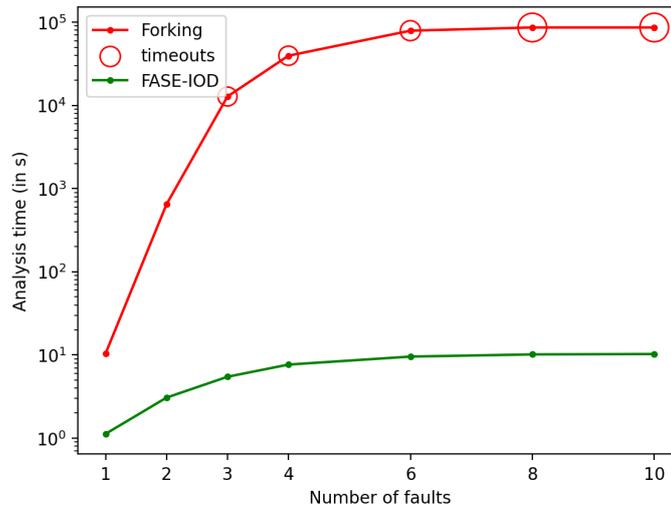


Figure 6.1: FASE-IOD and forking analysis time comparison for arbitrary data faults (RQ2.1)

3768 Then, we compare both techniques on the number of explored paths. FASE aims  
 3769 to reduce path explosion, we check whether it is experimentally true. The number of  
 3770 explored paths is presented in Table 6.4 (p.110) and illustrated in Figure 6.2 (p.110).  
 3771 The forking technique shows an exponential increase in the number of paths explored  
 3772 as the number of faults increases, while FASE increase in explored paths is much more  
 3773 mitigated, still opening new paths as they become feasible due to the introduced faults.  
 3774 On average, the forking technique explores x17 times more paths for 1 fault and x267  
 3775 times more for 2 faults.  
 3776

3777 Lastly, we notice that values from around 4 faults to 10 tend to plateau. This is  
 3778 especially visible in Figures 6.1 (p.109) and 6.2 (p.110). This is due not to our technique  
 3779 but to the fact that the programs in our benchmark are not overly complex, and 4 to

Table 6.4: FASE-IOD and forking number of explored paths comparison for arbitrary data faults (RQ2.1)

		Number of explored paths						
		1f	2f	3f	4f	6f	8f	10f
		average value						
FASE-IOD		22.3	108	381	1.05k	4.64k	11.1k	15.2k
Forking*		369	28.8k	170k	215k	1.41M	2.17M	2.2M

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts)

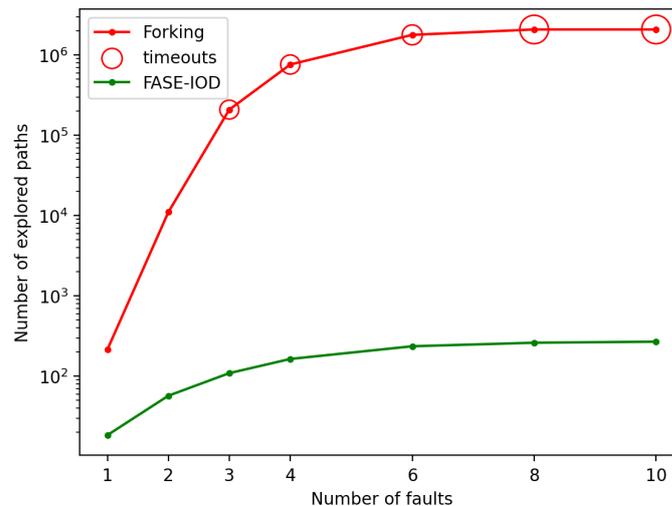


Figure 6.2: FASE-IOD and forking number of explored paths comparison for arbitrary data faults (RQ2.1)

3780 6 faults are sufficient to open all possible paths in the control-flow graph. Hence the  
 3781 number of explored paths plateaus. Interestingly, the analysis time also plateaus, as  
 3782 adding unnecessary faults does not open new paths to explore and relax the constraints  
 3783 on fault placing, not making the solver’s tasks harder.

**Conclusion RQ2.1.** For arbitrary data faults, FASE is able to scale in the number of faults considered while the forking technique struggles. A direct correlation between the path explosion experienced by the forking technique and its slower analysis time can be observed, compared to FASE which mitigates path explosion and obtains faster analysis times.

### 3784 6.3.2 Impact of Optimizations (RQ2.2)

3785 We now evaluate our different optimizations for FASE. Two angles were considered to  
 3786 optimize FASE analysis:

- 3787 – We have two optimizations changing the computation of the path predicates  
 3788 while keeping the same attacker power, EDS<sup>3</sup> and IOD<sup>4</sup>, plus their combination  
 3789 EDS+IOD;

<sup>3</sup>Early Detection of fault Saturation

<sup>4</sup>Injection On Demand

3790 – In addition, we also implemented optimizations aimed at adding query simplifi-  
 3791 cations between their computation and them being sent to the solver, sub-fault  
 3792 simplification (SFS), described in Section 5.3.6. As a reminder, SFS aims to  
 3793 detect and nullify faults subsumed by another one. Different variants were im-  
 3794 plemented.

3795 **RQ2.2.** What is the impact of our optimizations?

3796 **Goal.** The goal of this experiment is to evaluate our optimizations and how they  
 3797 contribute to reducing query complexity.

3798 **Protocol.** The performance benchmark is used in this experiment. We consider first an  
 3799 attacker model able to perform 1 to 10 arbitrary data faults for our first optimizations  
 3800 (EDS, IOD and EDS+IOD) evaluation. Then, for SFS, we consider an attacker able  
 3801 to inject 1 to 4 arbitrary data faults.

3802 **Results - EDS, IOD, EDS+IOD.** We start with our first set of optimizations and look at  
 3803 whether they result in a faster analysis, with all parameters equal otherwise. Analysis  
 3804 time data is presented in Table 6.5 (p.111) and illustrated in Figure 6.3 (p.112). Overall,  
 3805 as expected, FASE without optimizations is the slowest, followed by FASE-EDS (x1.8  
 3806 times faster than FASE for 1 fault), then FASE-EDS+IOD (x3.3 times faster than  
 3807 FASE for 1 fault) and the fastest being FASE-IOD (x3.7 times faster than FASE for 1  
 fault).

Table 6.5: FASE optimizations, analysis time for arbitrary data faults (RQ2.2)

	Analysis time (s)						
	1f	2f	3f	4f	6f	8f	10f
	average value						
FASE	7.73	27.6	92.6	261	1.06k	2.2k	2.56k
FASE-EDS	4.27	17.6	64.0	201	958	2.22k	2.7k
FASE-IOD	2.08	8.88	33.4	100	467	909	1.04k
FASE-EDS+IOD	2.36	10.5	38.5	111	502	936	1.01k
Forking*	39.5	3.58k	8.17k	9.15k	35.6k	86.4k	86.4k
	timeouts (24h) over 12 benchmarks in total						
FASE	0	0	0	0	0	0	0
FASE-EDS	0	0	0	0	0	0	0
FASE-IOD	0	0	0	0	0	0	0
FASE-EDS+IOD	0	0	0	0	0	0	0
Forking	0	0	6	9	11	12	12

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts)

3808 We continue our experimental evaluation with a sanity check, looking at the number  
 3809 of explored paths, presented in Table 6.6 (p.112). FASE, FASE-EDS, FASE-IOD and  
 3810 FASE-EDS+IOD explore the same number of paths, which is expected since they  
 3811 encode the same attacker model.

3812 To provide more intuition on how optimizations impact the analysis and in par-  
 3813 ticular contribute to reducing query complexity, we consider other, internal, metrics.  
 3814 First, we consider the ratio of queries created over queries sent to the solver, presented  
 3815 in Table 6.7 (p.113), to see if some queries have been made so simple that they can  
 3816 be arithmetically reduced to true or false by the analysis, without the help of the  
 3817

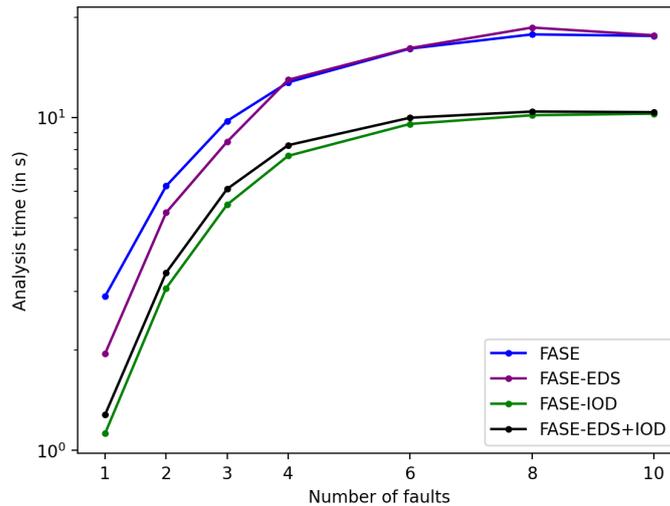


Figure 6.3: FASE optimizations, analysis time for arbitrary data faults (RQ2.2)

Table 6.6: FASE optimizations, number of explored paths for arbitrary data faults (RQ2.2)

	Number of explored paths						
	1f	2f	3f	4f	6f	8f	10f
	average value						
FASE	22.3	108	381	1.05k	4.64k	11.1k	15.2k
FASE-EDS	22.3	108	381	1.05k	4.64k	11.1k	15.2k
FASE-IOD	22.3	108	381	1.05k	4.64k	11.1k	15.2k
FASE-EDS+IOD	22.3	108	381	1.05k	4.64k	11.1k	15.2k
Forking*	369	28.8k	170k	215k	1.41M	2.17M	2.2M

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts)

3818 solver. At first glance, we notice the optimizations generate slightly more queries than  
 3819 FASE, as expected from the related algorithms. For instance, on average for 1 fault  
 3820 compared to FASE, FASE-EDS and FASE-IOD generate x1.04 times more queries and  
 3821 FASE-EDS+IOD generates x1.09 times more queries. Despite this slight increase, fewer  
 3822 queries are sent to the solver with the optimizations. For instance, for 1 fault, 44%  
 3823 of queries are simplified without a solver call for FASE, 58% for FASE-EDS, 65% for  
 3824 FASE-IOD and 62% for FASE-EDS+IOD. This shows that generating more queries  
 3825 can, somewhat counterintuitively, reduce overall query complexity, as up to 2/3 are  
 3826 resolved before needing to be sent to the solver. For the sake of comparison, 51% of  
 3827 queries are simplified without a solver call for the forking technique.

3828 The second internal metric measuring query complexity is the average solving time  
 3829 per query, presented in Table 6.8 (p.114) and illustrated in Figure 6.4 (p.114). Measures  
 3830 for average solving time per query tend to vary somewhat, which can have a significant  
 3831 impact due to the small values. To mitigate this, we perform 100 runs for each con-  
 3832 figuration (program, number of faults, technique) and only use the average values per  
 3833 program. FASE encodings introduce many new symbolic variables and are expected to  
 3834 increase the average solving time per query, which indeed is more than twice the value  
 3835 of the forking technique (x2.7 times averaging values for 1 and 2 faults). As expected,  
 3836 our optimizations reduce the average solving time per query. On average for all num-

Table 6.7: FASE optimizations, number of queries created and sent to the solver for arbitrary data faults (RQ2.2)

	Number of queries created / Number of queries sent to the solver									
	1f	2f	3f	4f	6f	8f	10f			
FASE	295 / 164	1.24k / 679	3.99k / 2.11k	10.1k / 5.22k	37.1k / 18.7k	69.7k / 34.2k	83.1k / 39.8k			
FASE-EDS	308 / 130	1.32k / 584	4.32k / 1.98k	11.1k / 5.29k	41.8k / 20.9k	79.6k / 39.9k	94.0k / 45.4k			
FASE-IOD	308 / 109	1.38k / 558	4.62k / 1.98k	12.1k / 5.46k	46.5k / 22.9k	89.0k / 44.9k	106k / 53.0k			
FASE-EDS+IOD	321 / 121	1.51k / 650	5.02k / 2.25k	13.1k / 6.14k	50.1k / 25.2k	96.6k / 49.3k	116k / 58.0k			
Forking*	3.54k / 1.74k	204k / 128k	1.03M / 533k	1.5M / 557k	6.99M / 4.44M	7.53M / 4.76M	7.02M / 4.65M			

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts)

3837 bers of faults, FASE-EDS is x1.04 faster than FASE in solving time, and FASE-IOD  
 3838 and FASE-EDS+IOD are x2 times faster than FASE.

Table 6.8: FASE optimizations, average solving time per query for arbitrary data faults (RQ2.2)

Average solving time per query (s)							
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE	0.050	0.041	0.037	0.037	0.038	0.041	0.038
FASE-EDS	0.039	0.045	0.037	0.037	0.036	0.041	0.037
FASE-IOD	0.020	0.020	0.021	0.021	0.019	0.020	0.019
FASE-EDS+IOD	0.020	0.019	0.021	0.021	0.020	0.020	0.019
Forking*	0.015	0.018	0.012	0.012	0.008	0.019	0.020

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts)

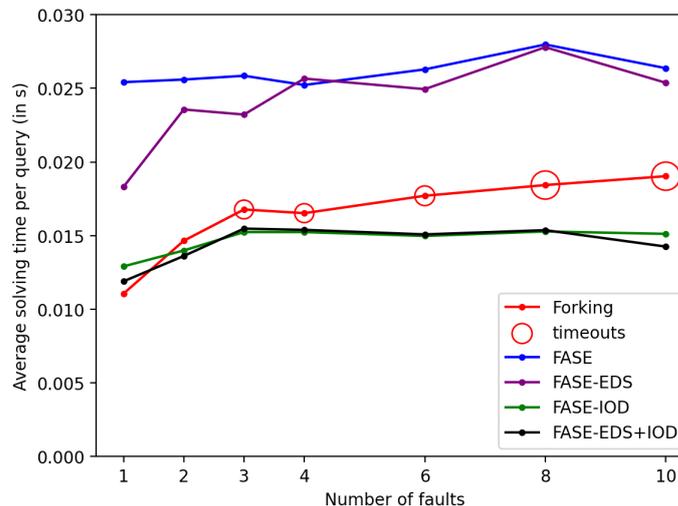


Figure 6.4: FASE optimizations, average solving time per query for arbitrary data faults (RQ2.2)

3838

3839 Our last internal metric to measure query complexity is the average number of ite  
 3840 operators per query, presented in Table 6.9 (p.115). For this metric, we add the values  
 3841 for non-faulted analysis for comparison. For a non-faulted run, there is almost no ite  
 3842 operator in queries, which makes an average of 0. FASE queries contain on average x7.4  
 3843 times more ite operators per query for 1 fault than the forking technique. This value  
 3844 continues to increase as we consider more faults, but only slightly, as faulted expressions  
 3845 are already computed in the path predicate, they are simply less constrained in number  
 3846 of activated faults. The slight increase is due to some new paths opening thanks to  
 3847 faults. FASE's optimizations follow the same trend as FASE. On average for all fault  
 3848 numbers, FASE-EDS only spear 6% of ite operators compared to FASE, while FASE-  
 3849 IOD spears 39% of ite operators and FASE-EDS+IOD 38%. This shows that our  
 3850 optimizations can indeed reduce injection points in query. Interestingly, the forking  
 3851 technique also induces more ite operators in queries as the number of faults considered  
 3852 increases, likely due to the new possibility offered by faults in conditional jumps that

cannot be arithmetically simplified anymore, in particular for the arbitrary data faults in this experiment.

Table 6.9: FASE optimizations average number of ite operator per query for arbitrary data faults (RQ2.2)

Average number of ite operator per query								
	0f	1f	2f	3f	4f	6f	8f	10f
average value								
FASE	0	2.87k	2.94k	2.99k	3.02k	3.05k	3.06k	3.07k
FASE-EDS	0	1.64k	3.11k	2.95k	3.0k	3.05k	3.06k	3.07k
FASE-IOD	0	1.53k	1.76k	1.83k	1.88k	1.92k	1.94k	1.94k
FASE-EDS+IOD	0	1.53k	1.94k	1.86k	1.9k	1.93k	1.94k	1.94k
Forking*	0	368	570	647	754	1.24k	1.32k	1.37k

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts)

Finally, we remark that a plateau appears in the different metrics as 4 to 6 faults are sufficient to open all possible paths (detailed in Section 6.3.1), except for the average time per query, which does not directly depend on the number of paths explored.

**Results - SFS.** We start by looking at the analysis time, presented in Table 6.10 (p.115), to check whether SFS optimizations result in a faster analysis. Unfortunately, SFS variants yield significantly slower analysis times, which start to equal those of FASE for large numbers of faults.

Table 6.10: Analysis time for sub-fault simplification for arbitrary data faults (RQ2.2)

Analysis time (s)				
	1f	2f	3f	4f
average value				
FASE	5.41	23.0	79.9	223
FASE-Free-Lazy	8.23	31.3	94.7	255
FASE-Free-Abort	6.77	25.9	85.1	225
FASE-Mem-Lazy	17.2	44.2	116	215
FASE-Mem-Abort	16.0	40.3	106	217

To get more intuitions on the impact of SFS optimizations, we consider the average solving time per query. As expected, SFS allows for slower solving time, as it computes query simplifications. SFS reduces the average solving time per query, from x1.1 to x1.4 times faster. The memoization of part of the queries does not yield faster processing than traversing queries each time (FASE-Free). Aborting SFS mechanism proves more efficient than risking extra queries with the lazy approach.

SFS only adds a constraint on queries stating that the selected faults are null to help the solver. BINSEC/ASE does not implement arithmetical simplification rules complex enough to be at query scale, hence SFS does not reduce the number of ite operators in queries as shown in Table 6.12.

Table 6.11: Average solving time per query for sub-fault simplification for arbitrary data faults (RQ2.2)

Average solving time per query (s)				
	1f	2f	3f	4f
average value				
FASE	0.030	0.025	0.026	0.027
FASE-Free-Lazy	0.026	0.020	0.020	0.020
FASE-Free-Abort	0.028	0.023	0.024	0.025
FASE-Mem-Lazy	0.026	0.020	0.019	0.019
FASE-Mem-Abort	0.028	0.024	0.024	0.024

Table 6.12: Average number of ite per query for sub-fault simplification for arbitrary data faults (RQ2.2)

Average number of ite per query				
	1f	2f	3f	4f
average value				
FASE	2.87k	2.94k	2.99k	3.02k
FASE-Free-Lazy	2.87k	2.94k	2.99k	3.02k
FASE-Free-Abort	2.87k	2.94k	2.99k	3.02k
FASE-Mem-Lazy	2.87k	2.94k	2.99k	3.02k
FASE-Mem-Abort	2.87k	2.94k	2.99k	3.02k

**Conclusion RQ2.2.** Our different optimizations reduce query complexity compared to FASE. While EDS, IOD and EDS+IOD result in significantly faster analysis time overall, SFS slows down the analysis, its computation inside BINSEC is likely too expensive compared to the gain obtained in query complexity. FASE-IOD is our best-performing optimization. It is followed by FASE-EDS+IOD where the combination of optimizations does not produce better results, maybe due to the redundancy between the under-approximation counter of IOD and the saturation logic of EDS. Then, FASE-EDS only slightly improves upon FASE. Last, SFS while simplifying queries results in worst performance.

### 3872 6.3.3 Comparison of the Different Forkless Encodings (RQ2.3)

3873 We now consider different forkless encodings with various operators as described in  
 3874 Section 5.3.3.1 (reminded Table 6.13) for the arbitrary data fault model. We compare  
 3875 four forkless arbitrary data encodings, FASE-Ite, FASE-Mul (with a multiplication  
 3876 operator), FASE-And (with a logical and operator) and FASE-Xor (with a logical xor  
 3877 operator). We used FASE-Ite until now, as it will be shown to be the fastest encoding.  
 3878 Here, FASE is used without optimizations.

3879 **RQ2.3.** What is the impact of the different encodings on the performance of the solver?

3880 **Goal.** Our goal is to explore the impact of different operators in the forkless encoding  
 3881 and their affinity with the solver used by BINSEC/ASE.

3882 **Protocol.** The performance benchmarks are used in this experiment. We consider an  
 3883 attacker model able to perform 1 to 10 arbitrary data faults.

Table 6.13: Arbitrary data encodings and their associated activation constraint

Fault model	Fault expression	Activation constraint
None	$x := y$	
Inlined if-then-else	$x := \text{ite } b_{addr} ? non\_det_{addr} : y$	$(non\_det_{addr} \neq y) \parallel \neg b_{addr}$
Multiplication	$x := y + b_{addr} \times non\_det_{addr}$	$(non\_det_{addr} \neq 0) \parallel \neg b_{addr}$
Bitwise and	$x := y + (-b_{addr}) \& non\_det_{addr}$	$(non\_det_{addr} \neq 0) \parallel \neg b_{addr}$
Bitwise xor	$x := y \oplus (-b_{addr}) \& non\_det_{addr}$	$(non\_det_{addr} \neq 0) \parallel \neg b_{addr}$

3884 **Results.** We start by looking at the impact of encodings on the general analysis time,  
3885 presented in Table 6.14 (p.117). For 1 fault, FASE-Ite is x1.2 times faster than FASE-  
3886 Xor, and x1.5 times faster than FASE-Mul and FASE-And. FASE-Xor starts to per-  
3887 form better than FASE-Ite for 8 faults and more. FASE-Mul and FASE-And start to  
experiment timeouts at 8 and 6 faults respectively.

Table 6.14: Encodings, analysis time for arbitrary data faults (RQ2.3)

	Analysis time (s)						
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE-Ite	8.59	27.7	96.9	271	1.13k	2.18k	2.63k
FASE-Xor	10.4	31.0	105	287	1.16k	1.77k	2.21k
FASE-Mul*	12.5	46.4	213	884	6.75k	7.32k	7.32k
FASE-And*	12.5	52.7	308	1.74k	7.31k	7.31k	7.31k
timeouts (24h) over 12 benchmarks in total							
FASE-Ite	0	0	0	0	0	0	0
FASE-Xor	0	0	0	0	0	0	0
FASE-Mul	0	0	0	0	0	1	1
FASE-And	0	0	0	0	1	1	1

FASE-Mul\*: values for 8 faults and above are computed for incomplete runs due to timeouts,  
FASE-And\*: values for 6 faults and above are computed for incomplete runs due to timeouts

3888

3889 Then, we evaluate the encodings on internal metrics, evaluating query complexity.  
3890 The ratio of queries created on the number of queries sent, presented in Table 6.15  
3891 (p.118), to reflect query complexity before solver call. The number of queries created  
3892 and sent ratio is the same for the 4 encodings up to 4 faults, around 50%. FASE-  
3893 Mul and FASE-And start to experience timeouts for 6 faults and above. This shows  
3894 that arithmetic simplification rules implemented inside BINSEC to simplify expressions  
3895 before queries are sent to the solver impact our four operators in the same way.

3896 Our last query complexity metric is the average solving time per query, presented  
3897 in Table 6.16 (p.118), to reflect the query complexity in solver calls. FASE-Ite is the  
3898 fastest, followed by FASE-Xor x1.3 times slower on average for all number of faults,  
3899 FASE-Mul is x2.4 times slower and FASE-And x3 times.

Table 6.15: Encodings, number of queries created and sent to the solver for arbitrary data faults (RQ2.3)

Number of queries created / Number of queries sent to the solver				
	1f	2f	3f	4f
average value				
FASE-Ite	295 /164	1.24k /679	3.99k /2.11k	10.1k /5.22k
FASE-Xor	295 /164	1.24k /679	3.99k /2.11k	10.1k /5.22k
FASE-Mul	295 /164	1.24k /679	3.99k /2.11k	10.1k /5.22k
FASE-And	295 /164	1.24k /679	3.99k /2.11k	10.1k /5.22k
		6f	8f	10f
FASE-Ite	37.1k /18.7k	69.7k /34.2k	83.1k /39.8k	
FASE-Xor	37.1k /18.7k	69.7k /34.2k	83.1k /39.8k	
FASE-Mul*	37.1k /18.7k	17.3k /9.22k	19.6k /10.2k	
FASE-And*	10.9k /5.79k	11.9k /6.16k	16.9k /8.67k	

FASE-Mul\*: values for 8 faults and above are computed for incomplete runs due to timeouts,  
FASE-And\*: values for 6 faults and above are computed for incomplete runs due to timeouts

Table 6.16: Encodings, average solving time per query for arbitrary data faults (RQ2.3)

Average solving time per query (s)							
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE-Ite	0.056	0.042	0.048	0.042	0.041	0.043	0.042
FASE-Xor	0.094	0.049	0.054	0.052	0.049	0.060	0.047
FASE-Mul*	0.089	0.070	0.107	0.084	0.101	0.165	0.138
FASE-And*	0.088	0.091	0.092	0.128	0.207	0.201	0.149

FASE-Mul\*: values for 8 faults and above are computed for incomplete runs due to timeouts,  
FASE-And\*: values for 6 faults and above are computed for incomplete runs due to timeouts

**Conclusion.** FASE-Ite is the best-performing encoding for arbitrary data faults with the solver used by BINSEC/ASE, up to 6 faults, then FASE-Xor becomes better. Variations in overall analysis time when changing the operator of the forkless encoding are significant. Arithmetic simplification rules implemented in BINSEC to simplify queries before sending them to the solver are also important and can spear solver calls if efficient. Here around 50% of queries are speared.

3900 In this experiment, we only considered the solver natively bound with BINSEC/ASE,  
3901 Bitwuzla. However, different solvers may favor different operators. We investigate this  
3902 question in Appendix B.5, comparing Bitwuzla to two other common solvers. While  
3903 Boolector also favors the 'ite' operator, z3 favors the xor encoding.

## 3904 6.4 FASE Evaluation of Other Fault Models (RQ3)

3905 After BINSEC/ASE evaluation for arbitrary data faults, we evaluate it for the other  
3906 implemented fault models: reset, bit-flip, test inversion and instruction skip.

3907 We aim to assess the extensibility of arbitrary data fault results. We expect the  
 3908 same performance trends for analysis time and the number of explored paths for other  
 3909 fault models, in particular for data faults since they are particular cases of arbitrary  
 3910 data faults.

3911 **RQ3.** For each fault model, can we scale in number of faults without path explosion  
 3912 compared to the forking technique?

3913 **Protocol.** We used the performance benchmarks in this experiment. We consider an  
 3914 attacker model able to perform 1 to 10 faults of the selected fault model. Our analysis  
 3915 explores exhaustively all possible paths until the time limit to obtain a fair comparison.

3916 For data faults (reset and bit-flip) and for instruction skip, we only consider FASE-  
 3917 IOD as it is implemented for those fault models and this optimization was shown to be  
 3918 the best-performing one for arbitrary data faults. The test inversion fault model has  
 3919 no optimization implemented.

### 3920 6.4.1 FASE Evaluation of Reset Faults (RQ3.1)

3921 First, we consider the overall analysis time, presented in Table 6.17 (p.119) . FASE-  
 3922 IOD shows better performance than the forking technique for reset faults, on average  
 3923 we are x2.5 times faster for 1 fault, x16 times faster for 2 faults, x90 times faster for  
 3924 3 faults and x560 times faster for 4 faults. The forking technique starts to timeout for  
 3925 6 faults instead of 3 for arbitrary data faults. This is likely due to the simpler fault  
 3926 model, as the effect of the fault is hardcoded to zero, instead of being a new symbolic  
 value.

Table 6.17: Analysis time for reset faults (RQ3.1)

	Analysis time (s)						
	1f	2f	3f	4f	6f	8f	10f
	average value						
FASE-IOD	0.363	1.15	2.81	5.04	9.78	13.3	15.2
Forking	0.922	18.8	255	2.82k	3.69k	27.8k	11.3k
	timeouts (24h) over 12 benchmarks in total						
FASE-IOD	0	0	0	0	0	0	0
Forking*	0	0	0	0	9	9	11

(Forking\*: values for 6 faults and above are computed for incomplete runs due to timeouts)

3927  
 3928 Then, the number of explored paths results are presented in Table 6.18 (p.120). We  
 3929 see that the forking technique explodes in number of paths explored while FASE-IOD  
 3930 does not. On average, we explore x5.4 times fewer paths for 1 fault, x48 for 2 faults  
 3931 and x400 for 3 faults.

**Conclusion RQ3.1.** We confirmed for the reset fault model the general trends of improved performance for analysis time and number of explored paths observed so far for arbitrary data faults.

Table 6.18: Number of explored paths for reset faults (RQ3.1)

Number of explored paths							
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE-IOD	11.0	30.2	66.1	113	202	265	300
Forking*	59.2	1.45k	26.3k	388k	851k	8.48M	2.0M

(Forking\*: values for 6 faults and above are computed for incomplete runs due to timeouts)

### 3932 6.4.2 FASE Evaluation of Bit-Flip Faults (RQ3.2)

3933 Analysis time results are presented in Table 6.19 (p.120). The forking technique seems  
 3934 to favor this fault model compared to arbitrary data and reset fault models. The same  
 3935 analysis time trends from other data faults can be seen, except that the forking line  
 3936 drops below FASE-IOD for 1 fault. The forking technique starts to experience timeouts  
 3937 for 4 faults. The forking technique is x1.9 times faster for 1 fault compared to FASE-  
 3938 IOD but becomes x3.9 times slower than FASE-IOD for 2 faults. A possible explanation  
 3939 for the forking technique being faster for one fault is the chosen forkless encoding, which  
 3940 results in harder-to-solve queries than the arbitrary data faults (in average solving time  
 3941 per query), that it is supposed to be a particular case of. Furthermore, as described  
 3942 in Section 5.3.7, the forking technique is implemented such that a fork in the analysis  
 3943 happens between a path with a bit-flip fault in an assignment and a path without. A  
 3944 path is not created for each possible bit the fault could flip. This implementation may  
 be considered an optimization of the forking technique for bit-flips.

Table 6.19: Analysis time for bit-flip faults (RQ3.2)

Analysis time (s)							
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE-IOD	1.25	2.89	4.69	6.35	8.76	9.27	9.15
Forking*	0.645	15.1	225	952	2.13k	13.4k	9.47k
timeouts (24h) over 12 benchmarks in total							
FASE-IOD	0	0	0	0	0	0	0
Forking	0	0	0	2	7	8	10

(Forking\*: values for 4 faults and above are computed for incomplete runs due to timeouts)

3945 The number of explored paths are presented in Table 6.20 (p.121). We can see that  
 3946 the number of explored paths explodes for the forking technique, from 68.9 on average  
 3947 for 1 fault to 1.44k for 2. FASE-IOD is able to mitigate the path explosion experienced  
 3948 by the forking technique, which explores on average x4.8, x38 and x311 times more  
 3949 paths for 1, 2 and 3 faults respectively.  
 3950

**Conclusion RQ3.2.** FASE-IOD outperforms the forking technique in scalability in number of faults considered. However, the forking technique results in a faster analysis time for 1 fault and then becomes much slower as more faults are considered. This could indicate that our forkless encoding of bit-flips is not the optimal one in terms of performance. We leave designing a faster one for future work.

Table 6.20: Number of explored paths for bit-flip faults (RQ3.2)

Number of explored paths							
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE-IOD	14.5	37.2	60.4	80.8	110	119	122
Forking	68.9	1.44k	18.8k	92.3k	207k	999k	734k

(Forking\*: values for 4 faults and above are computed for incomplete runs due to timeouts)

### 3951 6.4.3 FASE Evaluation of Test Inversion Faults (RQ3.3)

3952 As a reminder, the test inversion fault model has no optimization implemented.

3953 Analysis time results are presented in Table 6.21 (p.121). The forking technique  
 3954 results in faster analysis time for all numbers of faults, except for 10 faults. It is on  
 3955 average, over all number of faults, x1.4 times faster. Neither FASE nor the forking  
 technique experiment timeouts for this fault model.

Table 6.21: Analysis time for test inversion faults (RQ3.3)

Analysis time (s)							
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE	0.130	0.262	0.501	0.829	1.69	2.30	2.59
Forking	0.096	0.161	0.300	0.522	1.24	2.17	2.82
timeouts (24h) over 12 benchmarks in total							
FASE	0	0	0	0	0	0	0
Forking	0	0	0	0	0	0	0

3956 The number of explored paths is presented in Table 6.22 (p.121). The forking  
 3957 technique explores always more paths as the analysis considers more faults, but so does  
 3958 FASE, as by definition a test inversion fault opens a new path when activated. FASE  
 3959 explores slightly fewer paths on average. The gap in explored paths is significantly lesser  
 3960 than for the other presented fault models. There is not enough of a gain to compensate  
 3961 for the added query complexity of the forkless encoding. Furthermore, there are only  
 3962 few possible fault locations for test inversion in the performance benchmark, hence  
 3963 comparatively few new paths are created by the forking technique compared to data  
 3964 faults, which tips to our disadvantage the explored paths VS query complexity trade-  
 3965 off. However, as the gap in explored paths increases with faults, we can imagine that  
 3966 FASE, which is faster for 10 faults, remains faster for more than 10 faults with an  
 3967 increasing gain.  
 3968

Table 6.22: Number of explored paths for test inversion faults (RQ3.3)

Number of explored paths							
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE	12.2	33.2	69.7	122	249	343	379
Forking	15.3	42.6	92.4	171	433	769	1.02k

**Conclusion RQ3.3.** FASE supports the test inversion fault model, however, it is not the fastest technique for a small number of faults. FASE is more interesting than the forking technique for a large number of faults (10 or more). This is likely because FASE tends to perform better than the forking technique when there is a big gap in the number of paths explored, which is the case for test inversion faults only for a large number of faults. We emit the hypothesis that faults at control-flow points would favor forking faults.

#### 6.4.4 FASE Evaluation of Instruction Skip Faults (RQ3.4)

As a reminder, the only optimization implemented for instruction skips is IOD.

Analysis time results are presented in Table 6.23 (p.122). First, we can see that there is 1 program for which FASE-IOD cannot finish in time, the unrolled version of VerifyPIN with a PIN size of 16. It contains a long sequence of 175 assignments, that seem to overload the analysis when faulted in a forkless manner. For comparison, the basic VerifyPIN program never exceeds 10 consecutive assignments. We believe that in a program with regular branching instructions, constraints are regularly added, hence there is actually not that much freedom in variable assignments, contrary to a program with a long assignment sequence adding no constraints. The forkless instruction skip appears not well suited for programs with long assignment sequences.

The forking technique passes this program, likely because it spreads the difficulty of placing the fault by exploring different paths. The forkless technique does not experience any other timeout for the rest of the programs for any number of faults. The forking technique, however, starts to experience timeouts from 3 faults. With our incomplete results, FASE-IOD is slower for 1 fault but becomes faster for 2 faults. While exact numbers do not have meaning, since we count the full 24h in the case of a timeout, FASE-IOD is indeed faster for 2 faults onward.

Table 6.23: Analysis time for instruction skip faults (RQ3.4)

		Analysis time (s)						
		1f	2f	3f	4f	6f	8f	10f
		average value						
FASE-IOD*		95.6	150	329	515	822	842	962
Forking*		9.6	857	3.52k	16.3k	32.1k	86.4k	86.4k
		timeouts (24h) over 12 benchmarks in total						
FASE-IOD		1	1	1	1	1	1	1
Forking		0	0	3	6	10	12	12

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts  
FASE\*: all value from incomplete runs)

The number of explored paths results are presented in Table 6.24 (p.123). Despite the incomplete run, which is an unrolled program with very few paths to explore for FASE-IOD, we can see that FASE-IOD mitigated the path explosion faced by the forking technique for instruction skip faults.

Table 6.24: Number of explored paths for instruction skip faults (RQ3.4)

		Number of explored paths						
		1f	2f	3f	4f	6f	8f	10f
		average value						
FASE-IOD*		19.1	87.7	283	734	2.67k	4.7k	5.78k
Forking*		164	11.0k	211k	3.44M	7.6M	28.1M	29.2M

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts  
 FASE\*: all value from incomplete runs)

**Conclusion.** To conclude, instruction skip is a harder fault model for FASE-IOD, as we experience a timeout. However, FASE-IOD clearly out-performs the forking technique when considering multiple faults (3 faults or more).

### 3991 6.4.5 Summary

3992 Table 6.25 summarizes fault model data for comparison on analysis time, number of  
 3993 explored paths and average solving time per query.

**Conclusion RQ3.** FASE implements various fault models, with unequal results compared to the forking technique.

- We always outperform the forking technique for arbitrary data faults;
- We also always beat the forking technique for reset faults;
- It is only for 2 or more faults that FASE shows better results than the forking technique for bit-flip faults;
- We struggle more for instruction skip faults as we experience a timeout, but we clearly perform better than the forking technique for 3 or more faults.;
- FASE does not shine for test inversion faults as the forking technique experiences a restricted path explosion. We only outperform the forking technique for large numbers of faults (10 or more).

Overall, FASE tends to perform better than the forking technique for multi-fault analyses. It would seem like the forkless technique favors faults at data assignments, while the forking technique favors faults at control-flow points.

3994 We refer the interested reader to Appendix B for more experimental data on fault  
 3995 models other than arbitrary data faults. In particular, we investigate the impact of  
 3996 the optimizations implemented for each fault model and comment on the impact of the  
 3997 chosen encoding operator for reset faults.

Table 6.25: Summary table comparing FASE and the forking technique (RQ3)

		1f	2f	3f	4f	6f	8f	10f
Analysis time (s)								
FASE-IOD	AD	2.08	8.88	33.4	100	467	909	1.04k
Forking		39.5	3.58k	8.17k*	9.15k*	35.6k*	86.4k*	86.4k*
FASE-IOD	RS	0.363	1.15	2.81	5.04	9.78	13.3	15.2
Forking		0.922	18.8	255	2.82k	3.69k*	27.8k*	11.3k*
FASE-IOD	BF	1.25	2.89	4.69	6.35	8.76	9.27	9.15
Forking		0.645	15.1	225	952*	2.13k*	13.4k*	9.47k*
FASE	TI	0.130	0.262	0.501	0.829	1.69	2.30	2.59
Forking		0.096	0.161	0.300	0.522	1.24	2.17	2.82
FASE-IOD	IS	95.6*	150*	329*	515*	822*	842*	962*
Forking		9.6	857	3.52k*	16.3k*	32.1k*	86.4k*	86.4k*
Explored paths								
FASE-IOD	AD	22.3	108	381	1.05k	4.64k	11.1k	15.2k
Forking		369	28.8k	170k*	215k*	1.41M*	2.17M*	2.2M*
FASE-IOD	RS	11.0	30.2	66.1	113	202	265	300
Forking		59.2	1.45k	26.3k	388k	851k*	8.48M*	2.0M*
FASE-IOD	BF	14.5	37.2	60.4	80.8	110	119	122
Forking		68.9	1.44k	18.8k	92.3k*	207k*	999k*	734k*
FASE	TI	12.2	33.2	69.7	122	249	343	379
Forking		15.3	42.6	92.4	171	433	769	1.02k
FASE-IOD	IS	19.1*	87.7*	283*	734*	2.67k*	4.7k*	5.78k*
Forking		164	11.0k	211k*	3.44M*	7.6M*	28.1M*	29.2M*
Average solving time per query (s)								
FASE-IOD	AD	0.020	0.020	0.021	0.021	0.019	0.020	0.019
Forking		0.015	0.018	0.012*	0.012*	0.008*	0.019*	0.020*
FASE-IOD	RS	0.005	0.006	0.007	0.006	0.006	0.007	0.006
Forking		0.003	0.003	0.003	0.003	0.001*	0.001*	0.002*
FASE-IOD	BF	0.035	0.024	0.024	0.024	0.026	0.026	0.027
Forking		0.003	0.004	0.003	0.002*	0.003*	0.004*	0.004*
FASE	TI	0.002	0.002	0.002	0.002	0.002	0.002	0.002
Forking		0.002	0.002	0.002	0.002	0.002	0.002	0.002
FASE-IOD	IS	0.017*	0.025*	0.023*	0.020*	0.023*	0.021*	0.023*
Forking		0.010	0.012	0.006*	0.005*	0.002*	0.003*	0.003*

XXX\*: value computed for incomplete run due to timeouts

AD: arbitrary data, RS: reset, BF: bit-flip, TI: test inversion, IS: instruction skip

## 6.5 Forkless Faults in Instrumentation (RQ4)

3998

3999

4000

4001

4002

4003

Instrumentation is the process of adding data or metadata to a program, that an analyzer will pick up and reason upon. It is possible to inject faults through instrumentation in order to reuse existing analyzers, but it allows for less tuning of the analysis. For instance, the Lazart tool [PMPD14] takes C code as input and uses instrumentation with the Klee<sup>5</sup> symbolic execution engine. The goal of this section is to show that the

<sup>5</sup><https://klee.github.io/>

4004 forkless technique can also benefit instrumentation.

### 4005 **6.5.1 Experimental Settings**

4006 We start by briefly describing the experimental settings relevant to this section.

4007 **Benchmarks.** We still use the performance benchmarks (see Section 6.1.2.1).

4008 **Instrumentation Process.** Each source code, which is written in C, is instrumented  
4009 with forkless and forking fault encodings. We can no longer use the BINSEC framework,  
4010 easily at least, since we need a way to tell the analyzer which variables are symbolic  
4011 and to specify constraints on them. We use the Klee symbolic engine instead and take  
4012 advantage of its C-level API to instrument our C codes. We use Clang<sup>6</sup> to compile the  
4013 instrumented C to bytecode, which is given as input to Klee for analysis.

4014 To make sure the analysis would not fork for the forkless encoding, we did not use  
4015 the ite operator, but the encoding using the bitwise and operator. For more details,  
4016 we refer the interested reader to Appendix C.

4017 **Attacker Model.** For this section, we consider an attacker model able to perform 1 to  
4018 10 arbitrary data faults.

4019 **Exploration Strategy.** We set Klee to explore all paths for both the forkless and forking  
4020 encoding, in order to be able to compare results.

4021 **Other Settings.** We limit this experimental evaluation with a 1h timeout.

### 4022 **6.5.2 Scalability**

4023 **RQ4.** Can we scale in number of faults without path explosion, compared to the forking  
4024 encoding when used in instrumentation?

4025 **Goal.** The goal of this experiment is to check whether forkless faults can also benefit  
4026 program instrumentation.

4027 **Results.** We start by comparing the analysis time in order to see the interest of using  
4028 the forkless encoding in instrumentation compared to the forking encoding. Analysis  
4029 time results are presented in Table 6.26 (p.126) and illustrated in Figure 6.5 (p.126).  
4030 Interestingly, the forking encoding is slightly faster (x1.07 times) for 1 fault than the  
4031 forkless encoding. We believe results may vary depending on the operator used for  
4032 the forkless encoding, as they did with BINSEC/ASE, we leave for future work to  
4033 determine which forkless operator Klee favors and if one results in better performance  
4034 than the forking encoding. There are also fewer injection points and paths to explore in  
4035 C compared to binary code, which may tip the paths explored / query complexity trade-  
4036 off in our disfavor. We see that, for a 1h timeout, the forkless technique also experiences  
4037 a timeout when considering an attacker able to inject 3 faults. This happens for the  
4038 VerifyPIN\_7 program, the most complex of the VerifyPINs. The forking encoding also  
4039 experiences a timeout for this program and 3 more when considering 3 faults. Overall,  
4040 the forking encoding always experiences more timeouts than the forkless encoding.  
4041 Except for 1 fault, the forkless encoding shows improved performance in analysis time  
4042 (x6.4 times faster for 2 faults) and mitigates the path explosion experiences by the  
4043 forking encoding.

---

<sup>6</sup><https://clang.llvm.org/>

Table 6.26: Analysis time for arbitrary data faults in instrumentation with Klee (RQ4)

		Analysis time (s)						
		1f	2f	3f	4f	6f	8f	10f
		average value						
Forkless*		14.6	92.9	350	505	820	1.1k	1.2k
Forking*		13.7	592	1.4k	1.93k	2.71k	2.8k	2.84k
		timeouts (1h) over 12 benchmarks in total						
Forkless		0	0	1	1	2	3	4
Forking		0	0	4	5	8	10	10

(Forking\*, Forkless\*: values for 3 faults and above are computed for incomplete runs due to timeouts)

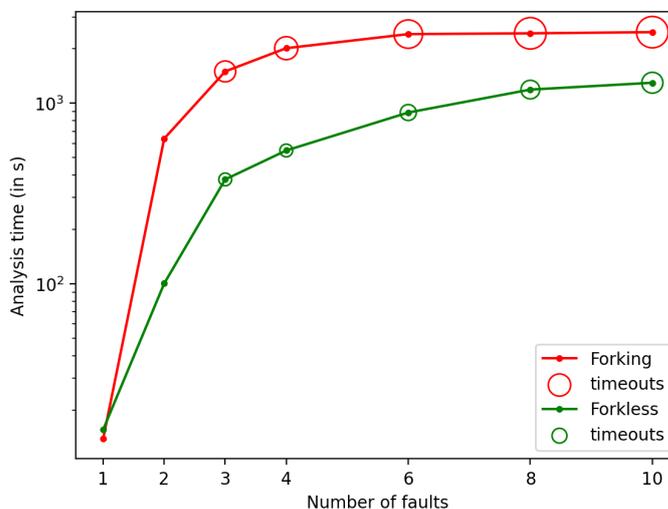


Figure 6.5: Forkless and forking analysis time for arbitrary data faults in instrumentation with Klee (RQ4)

4044 We also consider the number of explored paths to compare instrumentation tech-  
 4045 niques, presented in Table 6.27 (p.127) and illustrated in Figure 6.6 (p.127). We can  
 4046 see that, like inside BINSEC, a binary symbolic execution engine, a forking encoding  
 4047 suffers a path explosion in instrumented C code when the number of faults increases.  
 4048 The forking encoding explores x2.3 times more paths when considering 1 fault and x8.7  
 4049 times when considering 2.

4050 We note that the same plateau effect as in the BINSEC/ASE evaluation appears,  
 4051 maybe for fewer faults even. This phenomenon is due in part to incomplete runs, and  
 4052 in part to the benchmarks used that are not so complex and few faults already open  
 4053 most paths of the CFG.

### 4054 6.5.3 Conclusion

Table 6.27: Number of explored paths for arbitrary data faults in instrumentation with Klee (RQ4)

Number of explored paths							
	1f	2f	3f	4f	6f	8f	10f
average value							
Forkless*	17.7	84.3	154	384	420	498	527
Forking*	41.2	736	3.77k	15.0k	16.1k	12.5k	23.8k

(Forking\*, Forkless\*: values for 3 faults and above are computed for incomplete runs due to timeouts)

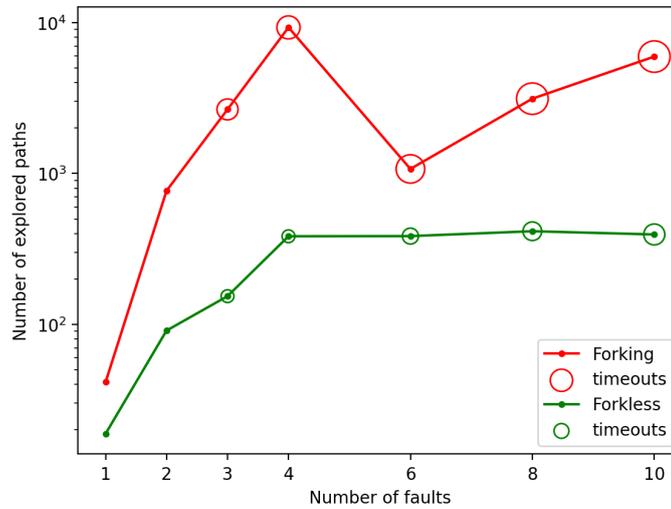


Figure 6.6: FASE and forking number of explored paths for arbitrary data faults in instrumentation with Klee (RQ4)

**Conclusion RQ4.** The forkless encoding enables improved performance and mitigates the path explosion experienced by the forking encoding when considering a multi-fault analysis, also when they are used for instrumentation, demonstrating the interest of the forkless encoding. Still, our other optimizations cannot be used.

## 4055 6.6 Security Scenarios

4056 In this section, we explore different security scenarios. This showcases the usability  
 4057 of our tool on various real-world programs, with different attacker models and attack  
 4058 goals. Programs' characteristics are presented in Table 6.28 (p.128).

4059 In particular, we aim to show that:

- 4060 – We can replay known attacks with BINSEC/ASE;
- 4061 – Our technique is expressive enough to handle real-world scenarios;
- 4062 – We are able to scale beyond our performance benchmarks to real-world programs.

Table 6.28: Benchmarks characteristics and statistics of a standard SE analysis

Program group (#)	C loc	x86 or ARM loc	BINSEC analysis - no fault			Time
			#instr (explored)	#paths	max #branch in a path	
Section 6.6						
CRT-RSA (3)	125-170	400-600	108k-29M	1	5k-1.3M	0.4s-1m27
Secret keeping machine (2)	100-200	240-360	1k-1.3k	1	130-150	< 0.1s
VerifyPIN_0 with SecSwift (1)	80	430	430	1	22	< 0.1s
Neural network (1)	100	275	774	1	109	0.1s

### 6.6.1 CRT-RSA

RSA is an asymmetric encryption algorithm. A way to implement this algorithm, yielding faster execution, is known as CRT-RSA. We study this algorithm through three implementations with various protections.

**Attacker Model.** We consider an attacker able to perform one reset fault in order to conduct a BellCoRe attack [BDL97]. If the output of the CRT-RSA is mathematically linked to the secret key, then, with a number of tries using the algorithm, it is possible to reconstruct the secret key.

**The Programs.** We base this analysis on the work of Puys et al. [PRBL14]. They describe three versions of CRT-RSA: unprotected, Shamir version and Aumuller version. Only the last one is shown to resist the BellCoRe attack.

**Results.** We were able to automatically reproduce the attack with 1 reset fault on the unprotected version of CRT-RSA in under 2 seconds. The whole program took just under 3 hours to fully analyze, yielding nine different adversarial paths executing a BellCoRe attack.

Our analysis experienced a timeout (24h) for the remaining two implementations without yielding attack paths. This is not due to a path explosion but to a query complexity explosion. We compare analysis metrics for 1h of analysis and 12 hours. The number of paths explored remained the same, 2 for both, and only x1.8 to x3.5 more instructions were analyzed. However, the average solving time per query went from 0.5 seconds to 1.7 seconds for both, which is a big increase, and significantly higher than the 0.005 seconds in the performance benchmark. No more queries were found trivially true or false by BINSEC/ASE without a solver call between 1h and 12h of analysis.

We summarize those results in Table 6.29.

Table 6.29: Summary of the CRT-RSA security scenario

Version	Ground truth	Result
CRT-RSA basic	Insecure	Insecure ✓
CRT-RSA Shamir	Insecure	Time-out without finding attacks ✗
CRT-RSA Aumuller	Secure	Time-out without finding attacks ✓

4088 **Conclusion.** We were able to find attacks on the basic CRT-RSA version and found  
4089 none in Aumuller’s version. Our analysis was not able to provide a conclusive result  
4090 for Shamir’s version.

4091 Cryptography is a weak spot of symbolic execution, as it relies on difficult-to-solve  
4092 problems to avoid reverse-engineering, which is related to what symbolic execution  
4093 attempts. It is hence not unexpected we would encounter scalability issues for this  
4094 security scenario. We leave for future work to design optimization techniques to improve  
4095 our analysis performance of cryptographic programs.

## 4096 6.6.2 Secret-keeping Machine

4097 Dullien [Dul17] proposes two versions of a toy secret-keeping machine in an exploration  
4098 of weird states that deviate from normal program behavior and showcases the impact  
4099 of implementation on program vulnerability.

4100 **The Programs.** One version is based on linked lists and is manually shown to be  
4101 exploitable by an attacker able to perform a single bit-flip, while the array version is  
4102 proven secure against that attacker model. This bit-flip happens in the memory of the  
4103 secret-keeping machine, modeling a bit-flip from Rowhammer for instance.

4104 The program works as follows. To store a new secret, a user gives the machine a  
4105 password and a secret. To retrieve their secret, the user sends just their password to  
4106 the secret-keeping machine, which will output the associated secret.

4107 **Attacker Model.** Dullien considers an attacker model able to perform one bit-flip in  
4108 the memory (not in registers or intermediary variables). Its goal is to find the victim’s  
4109 secret without knowing the password.

4110 We start by adapting Dullien’s attacker model to BINSEC/ASE, as we do not  
4111 support the capability ‘a bit-flip in memory’. Instead, we use BINSEC/ASE bit-flip  
4112 capability which happens when a value (register or memory) is updated. Then we  
4113 manually filter the results to only consider faults in memory loads overwritten afterward  
4114 and in memory stores.

4115 **Attack Scenario.** This attack is performed in a scenario of interaction with the secret-  
4116 keeping machine:

- 4117 1. The attacker interacts with the machine, storing and retrieving secrets, to set the  
4118 machine’s memory in a specific state;
- 4119 2. The victim stores a password-secret pair;
- 4120 3. The attacker interacts again with the machine, induces a bit-flip in memory at  
4121 some point and enters a value in the machine that should output the victim’s  
4122 secret.

4123 We now present how we experimented with this security scenario.

4124 **Experimental Settings.** We reimplemented the two versions presented in the paper.  
4125 One implements the memory with a linked list, the second version uses an array. We  
4126 simulate the interaction by calling multiple times the secret-keeping machine program  
4127 with different inputs recreating the scenario proposed in the paper. We discriminate  
4128 when to inject faults and when not by duplicating the functions interacting with the  
4129 machine’s memory, which is unique and shared. Only the duplicates are faulted.

4130 For this benchmark, we activated faults on variables used as addresses.

4131 **Finding a Vulnerability.** Our first goal is to reproduce the attack found in the paper

4132 on the vulnerable version. We are happy to report that BINSEC/ASE is able to find  
4133 the attack described in the paper following the same interaction scenario.

- 4134 – In memory, data are structured by triplets of password, secret and pointer to  
4135 the next triplet. When a secret is returned to the user, the program clears the  
4136 triplet, and puts it on the free linked list, setting the pointer to the next triplet  
4137 accordingly;
- 4138 – The attack changes a pointer to the next memory triplet, such that instead of  
4139 pointing to the next password, it points to its associated secret, both belonging  
4140 to the attacker;
- 4141 – Then the attacker sends this secret, acting as a password, to the secret machine,  
4142 which will output the next pointer and clear the triplet, i.e. the secret, the pointer  
4143 and the password of the next triplet. The (victim’s) password of that next triplet  
4144 is overwritten by a computable pointer, hence now known to the attacker;
- 4145 – The attacker sends to the machine the pointer value and the machine outputs  
4146 the victim’s secret.

4147 With this restrictive scenario, it is the only attack we found. It took BINSEC/ASE  
4148 0.1s to find it and 7h30min to analyze the whole program.

4149 **Program Resistance.** Our second goal was to see if we could show that the secure  
4150 version is indeed secure. We analyze the second, array-built version with the same  
4151 attacker model. Our analysis was complete and did not report any attacks for a bit-flip  
4152 in memory.

4153 **A More Powerful Attacker Model.** Then we explore for which attacker model the  
4154 resistant version becomes vulnerable. Our analysis of the array-built program shows  
4155 that by inducing a bit-flip in registers, the attack succeeds. We find the first adversarial  
4156 path in 78 minutes and analyze the whole program in 13h20min, finding 34 adversarial  
4157 paths. In the linked-list version, pointers to the next cell are placed in memory and  
4158 are the target of the attack found. In the array version, array indices provide that  
4159 function but are stored in registers, hence the array version becomes vulnerable when  
4160 broadening the specter of the attacker’s capability.

Table 6.30: Summary of the Secret keeping machine security scenario

Version	Attacker model	Ground truth	Result
Linked-list	1 bit-flip in memory	Insecure	Insecure ✓
Array	1 bit-flip in memory	Secure	Secure ✓
Array	1 bit-flip anywhere	Insecure	Insecure ✓

4161 **Conclusion.** To conclude, BINSEC/ASE was able to replay the results from Dullien  
4162 [Dul17] as summarized in Table 6.30. They highlight the importance of the implemen-  
4163 tation decisions in program vulnerability. We also show that specifying the attacker  
4164 model against which to protect a program from is important, as a slight difference can  
4165 open new attack surface for an attacker.

### 4166 6.6.3 SecSwift Countermeasure

4167 SecSwift [dF21, CdFB21] is a control-flow integrity (CFI) countermeasure at the llvm  
4168 level. It is developed by STMicroelectronics.

4169 **Experimental Settings.** We borrow an implementation of the SecSwift ControlFlow  
4170 countermeasure from Boespflug *et al.* [BEMP20] at the llvm level. According to the

4171 authors, it associates a unique identifier to each basic block and uses a xor-based  
 4172 mechanism to ensure that the correct branch has been taken. In a nutshell, it prevents  
 4173 the execution from deviating from the CFG. It is a restricted, less powerful version of the  
 4174 real, closed-source, SecSwift countermeasure. We applied the protection to VerifyPIN  
 4175 version 0, the basic version without protections, to study the effects of SecSwift without  
 4176 interference from other protections.

4177 **Goal.** Our goal is to assess the increased protection provided by the SecSwift coun-  
 4178 termeasure and to investigate whether we can find attacks on a SecSwift-protected  
 4179 program.

4180 **Attacker Model.** We consider an attacker able to perform one arbitrary data fault  
 4181 or one test inversion. Their goal is to be authenticated without entering the correct  
 4182 password.

4183 **Results.** We were able to find nine different attack paths in this binary, all yielding  
 4184 an early loop exit, with either a single test inversion or a single arbitrary data fault.  
 4185 Those attacks take the branch not entering the loop body, or exiting it early, paths that  
 4186 belong to the CFG of the program. These attacks are not unexpected as the adversarial  
 4187 paths stay in the legitimate CFG, yet it is still interesting that our technique finds them  
 4188 automatically. It took BINSEC/ASE 22 minutes to find the first attack and 30 minutes  
 to analyze the whole program. Results are summarized in Table 6.31.

Table 6.31: Summary of the SecSwift security scenario

Version	Ground truth	Result
VerifyPIN_0 with SecSwift	Insecure	Insecure ✓

4189

4190 **Conclusion.** It is important to characterize exactly against what type of attack a  
 4191 protection is efficient, so they can be combined to provide increased resistance. An  
 4192 incomplete countermeasure may not increase a program’s resistance to an attacker.

#### 4193 6.6.4 Neural Network

4194 The CEA List from Grenoble provided us with a very small neural network, based  
 4195 on an implementation from CEA Gardanne [DHM<sup>+</sup>23]. This neural network is used  
 4196 for evaluating hardware fault injection attack efficiency on microcontrollers for neural  
 4197 networks and highlights their vulnerability. While actual neural networks for embedded  
 4198 system evaluations have tens to hundreds more neurons than this one, it has the same  
 4199 representative architecture (few layers deep, no convolution, ReLU) and allows testing  
 4200 proof of concepts. It has 3 layers of 4,5 and 3 nodes.

4201 The goal of this experiment is to see if an attacker can disturb the computation  
 4202 of the neural network so that it classifies the input incorrectly, finding adversarial  
 4203 examples.

4204 This security scenario was the motivation behind extending BINSEC/ASE support  
 4205 to ARM executables.

4206 **Attacker Model.** The attacker model chosen is able to perform 1 bit-flip with the goal  
 4207 of making the neural network classify a particular input in the desired incorrect class.

4208 **Vulnerabilities Found.** Since this program has not been designed to resist fault in-  
 4209 jection, it is expected that we would find vulnerabilities. BINSEC/ASE found 13

4210 adversarial attack paths in 18 hours, the first attack path was found in 66 seconds.  
4211 Here are a few examples of successful fault injection:

- 4212 – Corrupting the load of a weight;
- 4213 – Corrupting the multiplication between a weight and the value coming from the  
4214 previous layer;
- 4215 – Corrupting the formatting function used.

4216 Here, many corruptions result in a wrong computation giving the desired class, but  
4217 they still need to be targeted depending on the input and the desired class. Table 6.32  
summarizes those results.

Table 6.32: Summary of the Neural Network security scenario

Version	Ground truth	Result
Neural Network	Insecure	Insecure ✓

4218

4219 **Conclusion.** We have shown on this small neural network that an attacker has a wide  
4220 attack surface at their disposal to corrupt to their wishes the output of the network.  
4221 This stresses the importance of taking security into account for neural networks, espe-  
4222 cially for embedded usage and if they perform a security feature.

### 4223 6.6.5 Security Scenarios Feedback

4224 After those security scenarios, this section provides some feedback on the usability of  
4225 BINSEC/ASE.

4226 **What Was Easy With BINSEC/ASE.** There was no real difficulty in configuring BIN-  
4227 SEC/ASE once the program was implemented with its oracle. This last part was not  
4228 always easy though.

4229 **What Was Not so Easy With BINSEC/ASE.** Understanding BINSEC/ASE output is  
4230 not trivial for a complex program. In particular, BINSEC/ASE provides the address  
4231 of activated faults, but only a very loose sense of the attack path, which hinders attack  
4232 understanding.

4233 **Improvements Due to the Security Scenarios.** In order to analyze the neural network,  
4234 which was provided to us as an ARM binary, we had to extend BINSEC/ASE support  
4235 to this ISA by back-porting it from BINSEC.

## 4236 6.7 Case Study: WooKey Bootloader

4237 We confront BINSEC/ASE to a real-life security system, WooKey. After presenting  
4238 WooKey and the goals of this case study, we explore the bootloader, replaying attacks,  
4239 and evaluate a recent protection scheme, where we found an attack not mentioned  
4240 before. We propose and evaluate our own patch. Then we explore two other attacks  
4241 on the WooKey project.

### 4242 6.7.1 Presentation of WooKey

4243 First presented in 2018 by ANSSI, the French cybersecurity agency, the WooKey plat-  
4244 form [BRT<sup>+</sup>18, Woo] is “a custom STM32-based USB thumb drive with mass storage  
4245 capabilities designed for user data encryption and protection, with a full-fledged set

4246 of in-depth security defenses”. Their choice to be open source and open hardware  
 4247 makes WooKey a relevant case study: it is a real-life, complex device, security-focused  
 4248 and available for reproducibility. Wookey has been extensively analyzed, as it was the  
 4249 target of an ANSSI cybersecurity challenge for security professionals [AAE<sup>+</sup>20].

4250 We focus on the WooKey bootloader, a dual-bank system enabling hot firmware  
 4251 updates. The system is hardened, especially redundant test protections are present in  
 4252 critical sections to protect against an attacker able to perform one test inversion fault.

## 4253 6.7.2 Security Scenario and Goal of our Study

4254 We detail in this section the experimental setting and the analysis overview constituting  
 4255 this case study.

4256 **WooKey Bootloader.** We see in Table 6.33 (p.134) that WooKey bootloader size is  
 4257 orders of magnitude larger than the programs used for performance evaluation. Wookey  
 4258 is available as C code. We borrowed C stubs from Lacombe et al. [LFBP21] and we  
 4259 compile the bootloader as we did for the evaluation benchmarks (Section 6.1.2.1).

4260 **Attacker Model.** We consider the same attack goal as the ANSSI challenge did  
 4261 [AAE<sup>+</sup>20]: the attacker seeks to manipulate the bootloader logic to boot on the older  
 4262 firmware, more likely to contain security vulnerabilities. We also consider an attacker  
 4263 able to perform a single arbitrary data fault.

4264 **Analysis Overview.** We conduct the following three analyses:

- 4265 1. We automatically analyze WooKey at binary-level to check whether we are able  
 4266 to find previously known attacks [LFBP21], and/or new ones in Section 6.7.3: we  
 4267 are indeed able to find the two attacks identified by prior work [LFBP21] (A1,  
 4268 A2), *as well as an attack they did not report* (A3);
- 4269 2. We automatically analyze at binary level the patched version of Wookey proposed  
 4270 by Lacombe et al. [LFBP21] in Section 6.7.4: we found that the proposed patch  
 4271 indeed blocks the two known attacks (A1 and A2), but not the unreported attack  
 4272 (A3);
- 4273 3. We propose a definitive patch by adding a counter-measure for A3 and removing  
 4274 parts of the counter-measures that are shown to be useless here (Section 6.7.5).  
 4275 The patch is proven correct w.r.t. our attack model;
- 4276 4. We explore in Section 6.7.6 other parts of the WooKey project, replaying two  
 4277 other attacks.

## 4278 6.7.3 Analyze Key Parts of Wookey

4279 The relevant parts of the bootloader’s functions are presented in Figure 6.7. Attacks  
 4280 found are summarized in Table 6.34.

4281 **Methodology.** In this case study, since we only consider an attacker model able to  
 4282 perform one fault (the program is not meant to resist a more powerful attacker), we  
 4283 divide our analysis, faulting only one function at a time. This divide-and-conquer  
 4284 approach allows to obtain fast results.

4285 **Results.** Lacombe et al. [LFBP21] find an attack in the *loader\_exec\_req\_selectbank*  
 4286 function (A1) and another in the *loader\_exec\_req\_flashlock* function (A2). They cor-  
 4287 respond to data corruption in branching conditions.

Table 6.33: Benchmarks characteristics and statistics of a standard SE analysis

Program group (#)	C loc	x86 loc	#instructions (explored)	#paths	BINSEC analysis - no fault		Time
					max #branch in a path		
Sections 6.2 to 6.5							
Verify-PINs (8)	80-140	160-215	192-269	1		17-34	< 0.1s
Verify-PIN unrolled (2)	40-85	140-430	142-442	5-17		5-17	< 0.1s
npo2 (2)	50	200-220	607-653	3		31-33	< 0.1s
Section 6.6							
CRT-RSA (3)	125-170	400-600	108k-29M	1		5k-1.3M	0.4s - 1m27
Secret keeping machine (2)	100-200	240-360	1k-1.3k	1		130-150	< 0.1s
Verify-PIN_0 with SecSwift	80	430	430	1		22	< 0.1s
Neural network	100	275	774	1		109	0.1s
Section 6.7							
WooKey bootloader	3.2k	2350	290k	17		18k	9s

```

1 static loader_request_t loader_exec_req_selectbank (loaderstate_t
  nextstate){
2     // ...
3     if((flip_shared_vars.fw.bootable == FW_BOOTABLE &&
        flop_shared_vars.fw.bootable == FW_BOOTABLE) &&
4         !(flip_shared_vars.fw.bootable != FW_BOOTABLE ||
        flop_shared_vars.fw.bootable != FW_BOOTABLE)){
5         // ...
6     }
7     if(flop_shared_vars.fw.bootable == FW_BOOTABLE
8 (CM1 - X)    && flip_shared_vars.fw.bootable != FW_BOOTABLE){
9         if(!(flop_shared_vars.fw.bootable == FW_BOOTABLE
10 (CM2)      && flip_shared_vars.fw.bootable != FW_BOOTABLE))
11             goto err;
12         ctx.boot_flop = sectrue ;
13         // ...
14     }
15     if(flip_shared_vars.fw.bootable == FW_BOOTABLE
16 (CM3 - X)    && flop_shared_vars.fw.bootable != FW_BOOTABLE){
17         ctx.boot_flip = sectrue ;
18         // ...
19     }
20     // ...
21 }
22
23 static loader_request_t loader_exec_req_flashlock (loader_state_t
  nextstate){
24     // ...
25     if (ctx.dfu_mode == sectrue) {
26 (CMA)      if (ctx.dfu_mode != sectrue)
27 (CMA)      goto err;
28         // ...
29     }
30     else if (ctx.dfu mode == secfalse ) {
31         if (ctx.bootflip == sectrue ) {
32 (CM4)      if (ctx.bootflip != sectrue)
33 (CM4)      goto err;
34         // ...
35         ctx.next_stage = (app_entry_t) (FWLSTART);
36     }
37     // ...
38 }
39 }

```

Figure 6.7: functions of WooKey’s bootloader, with [LFBP21] fixes and **our patch**

Table 6.34: Table summarizing the effects of countermeasures

Protection scheme	A1	A2	A3
Normal Wookey	1.3	1.31	1.25
Prior patch (CM1+CM2+CM3+CM4)	✓	✓	✓
Our patch (CM2+CM4+CMA)	✗	✗	✗

Legend - ✓: attack path found by our tool / ✗: no attack found

- 4288 – For attack A1, faulting the test at line 3 when both firmware flip and flop are  
4289 bootable results in the execution carrying on to a later test (line 7), which only  
4290 checks if flop can be booted, assuming at least one firmware cannot. If both are  
4291 bootable, and flop is the older one, the attacker’s goal can be satisfied;
- 4292 – Attack A2 takes advantage of the lack of countermeasure in *loader\_exec\_req\_*  
4293 *flashlock*, which computes the pointer to the boot function of the chosen firmware.  
4294 Inducing a fault that inverts the test line 31 leads to the wrong pointer being se-  
4295 lected.

4296 We are able to find both attacks, linking faults back to their locations in the C  
4297 code with debug information. *We also find an additional, unreported attack*<sup>7</sup>, faulting  
4298 another part of the *loader\_exec\_req\_flashlock* function (A3).

- 4299 – In attack (A3), by inverting the test line 25, it is possible to select the wrong  
4300 firmware pointer to boot on, in particular, if we are not in a secure Direct  
4301 Firmware Update (DFU) mode, it is possible to boot as if we were.

4302 **Remark.** While we consider an attacker model able to perform an arbitrary data fault,  
4303 all attack paths found can equivalently be produced with a test inversion fault, the  
4304 fault model the WooKey project aims to be resistant to. This shows the scalability  
4305 of our tool to consider a more powerful attacker model. Interestingly, using arbitrary  
4306 data faults instead of test inversion does not open new attack paths here.

#### 4307 6.7.4 Analyze a Security Patch of WooKey

4308 We now evaluate the protection scheme proposed by Lacombe et al. [LFBP21] for these  
4309 attacks. It consists of four extra counter-measures, adding duplicated tests, named from  
4310 CM1 to CM4, illustrated in Figure 6.7 (p.135).

4311 Our results are summarized in Table 6.34 (p.135). We found indeed that:

- 4312 – The full protection prevents attacks A1 and A2, as claimed by the authors of the  
4313 patch;
- 4314 – Yet, our analysis shows that the protection does not prevent the unreported  
4315 attack A3.

#### 4316 6.7.5 Propose a New Patch and Evaluate It

4317 **Patching Attack A3.** We manually inspect the analysis results to understand what  
4318 happens. We have especially been able to identify the root cause of A3 and propose  
4319 a dedicated countermeasure for it (named CMA and illustrated in Figure 6.7 (p.135)).  
4320 It consists again of adding a duplicated test line 26.

4321 **Counter-measure Analysis.** We analyze each counter-measure in isolation, meaning  
4322 we implement only one per binary and evaluate each binary separately. We have been  
4323 able to understand that:

- 4324 – Counter-measures CM1 and CM3 do not block any attack path as they are redun-  
4325 dant with other tests in the code and can be safely removed, for the considered  
4326 attacker model. As protections often induce an important overhead, only adding  
4327 useful ones is an interesting area of research. From CM1, we derive the following  
4328 principle: when there are two nested tests, and no instruction belonging solely  
4329 to the outer test, doubling the outer test is unnecessary (illustrated Figure 6.8

---

<sup>7</sup>After discussion with the authors [LFBP21], it turns out that they actually found this path but did not report it in the article, as they did not consider it as a real attack w.r.t. the Wookey challenge.

- 4330 (p.137)). We have not been able to derive a general rule from the removal of  
 4331 CM3.  
 4332 – CM2 and CM4 each block one of the attacks found by Lacombe et al.

```

1  if (condition_1) {
2      if (condition_1){
3          if (condition_2) {
4              if (condition_2) {
5                  // code only here
6              }
7          }
8      }
9  }
10 }
```

Figure 6.8: CM1 double test pattern

4333 Overall, our new patch (CMA + refined former patch) is shown by our tool to protect  
 4334 against all the attacks, for an attacker able to perform one arbitrary data fault, which  
 4335 is an attacker more powerful than the one considered initially in the WooKey project.

### 4336 6.7.6 Other Attacks on WooKey

4337 We also explore other parts of the WooKey project and we were able to replay two  
 4338 other known attacks.

#### 4339 6.7.6.1 Attack Vectors Combination

4340 The iso8716 library is used in WooKey for secure communication.

4341 **Attacker Model.** We consider an attacker able to perform one arbitrary data fault.  
 4342 The attack goal (A4) is to induce a buffer overflow to be exploited in a combined attack.

4343 **Results.** Our analysis found a vulnerability to fault injection which enables a software  
 4344 buffer-overflow in function *SC\_get\_ATR* [LFBP21]. Contrary to the attack paths found  
 4345 before, this one could not be created with a test inversion fault.

#### 4346 6.7.6.2 Faulty Redundant Test

4347 Martin *et al.* [MKP22] study the correctness of implemented counter-measures. In  
 4348 particular, they were able to detect that one counter-measure doubling a test was  
 4349 incorrectly implemented.

4350 **Attacker Model.** We consider an attacker able to perform an arbitrary data fault,  
 4351 including on flags. Since we study a small function that directly starts with the double  
 4352 test of interest, this is equivalent to a test inversion fault.

4353 **Results.** We show that indeed, a redundant test preventing single test inversion faults in  
 4354 the *loader\_set\_state* function is in fact vulnerable (A5), hence incorrectly implemented  
 4355 as pointed out by Martin et al. We implement the correction they propose and show  
 4356 the correctness of their patch.

4357 **6.7.7 Case Study Conclusion**

4358 In this case study of WooKey’s bootloader, we demonstrated some uses of our tool, in  
 4359 a real-life context. It can replay known attacks, including from source-level analysis  
 4360 and find new ones. It also assists in the countermeasure design and evaluation process.  
 4361 A summary table is presented in Table 6.35.

Table 6.35: Summary of the WooKey bootloader case study

Attacker model		Version	Attack	Ground truth	Result
Capability	Attack goal				
1 AD	boot on older firmware	basic	A1, A2	<i>Insecure</i>	<i>Insecure</i>
		basic	A3	–	<i>Insecure</i>
		CM 1+2+3+4	A1, A2	<b>Secure</b>	<b>Secure</b>
		CM 1+2+3+4	A3	–	<i>Insecure</i>
		CM 2+4	A1, A2	–	<b>Secure</b>
		CM 2+4+A	A1, A2	<b>Secure</b>	<b>Secure</b>
		CM 2+4+A	A3	<b>Secure</b>	<b>Secure</b>
1 AD	trigger a buffer overflow	basic	A4	<i>Insecure</i>	<i>Insecure</i>
1 AD	bypass redundant test	basic	A5	<i>Insecure</i>	<i>Insecure</i>
		basic	A5	<b>Secure</b>	<b>Secure</b>

AD: arbitrary data fault

4362

# Chapter 7

4363

## Conclusion

4364

### Contents

4365

4366

**7.1 Conclusion** . . . . . **139**

4367

**7.2 Perspectives** . . . . . **140**

4368

4370

4371

## 7.1 Conclusion

4372

4373

4374

4375

4376

4377

4378

Program analysis techniques for safety consider a rather weak attacker model, only able to craft smart inputs. In practice, an advanced attacker has many attack vectors at their disposal, providing them with a wide range of (reusable) capabilities, threatening program security. In this work, we address program analysis in the presence of an advanced attacker. Our goal is to design a bug-finding program analysis technique that is automatic, efficient and generic, integrating a model of an advanced attacker to assess their impact on a program’s security properties.

4379

4380

4381

4382

First, we define a model of what we consider to be an advanced attacker. We formalize the impact of this advanced attacker on a program’s execution. We extend the standard notion of reachability to adversarial reachability, by including attacker actions as a new type of transition in the transition system representing the program.

4383

4384

4385

4386

4387

4388

4389

4390

4391

Then, we propose a dedicated symbolic algorithm for adversarial reachability, adversarial symbolic execution. We extend symbolic execution by integrating attacker actions as fault injections into the program analysis. A main limitation experienced by the state-of-the-art techniques is path explosion, we mitigate it by integrating a novel forkless encoding of attacker actions. Furthermore, we design two optimizations dedicated to reducing the number of fault injections in queries to alleviate the solver’s task. Early Detection of fault Saturation aims to stop fault injection as soon as possible, while Injection On Demand only adds faults when necessary to explore a branch, injecting them as late as possible. Both preserve the attacker model’s power.

4392

4393

4394

4395

4396

4397

We implement adversarial symbolic execution on top of the BINSEC framework for binary-level analysis. While adversarial symbolic execution is independent of the analysis abstraction level, what is executed on the processor, and hence faulted, is the binary program, motivating our implementation abstraction level. Our technique is shown to significantly reduce the number of paths to explore, and scales up to 10 faults on a standard SWIFI benchmark, where prior forking attempts timeout for 3 faults.

4398 We illustrate the interest and feasibility of our technique through different security  
4399 scenarios, replaying attacks of attackers with various capabilities. Also, we show that  
4400 our method scales to realistic size examples, such as the WooKey project where we have  
4401 been able to replay known fault attacks and even find a vulnerability not reported in  
4402 a recently proposed countermeasure patch.

4403 Essentially, this work is a first step in designing efficient program analysis techniques  
4404 able to take into account advanced attackers.

## 4405 7.2 Perspectives

4406 We present several directions to extend or improve the work presented in this thesis.

4407 **Extend the Supported Attacker Model.** While our formalization allows for a generic  
4408 attacker model, our implementation of adversarial symbolic execution is more restricted  
4409 in attacker capabilities. For instance, it would be interesting and would widen the scope  
4410 of this tool’s usability to implement an attacker model able to perform actions from  
4411 different fault models in one attack. Similarly, our technique would benefit from having  
4412 an efficient algorithm to deal with fault on addresses, for instance by recording seen  
4413 addresses deemed ‘interesting’, or having a preprocessing selecting a set of ‘interesting’  
4414 addresses. It would allow corrupting pointers, for instance, to point to the reference  
4415 password instead of the user input password in the password comparison of an authenti-  
4416 cation program. Efficient faults on addresses can also be used for advanced control-flow  
4417 attacks, such as return-oriented programming (ROP), when reasoning about code ad-  
4418 dresses. Currently, we support those faults with a combinatorial explosion. We are also  
4419 missing an efficient algorithm for general instruction corruption. Another angle would  
4420 be to extend the supported fault models to broaden the scope of represented attacks.  
4421 In particular, micro-architectural components could be modeled to account for more  
4422 attacks targeting micro-architectural elements such as physical fault injection on the  
4423 prefetch buffer.

4424 Such extensions would likely increase the complexity of the analysis. We believe  
4425 more optimizations can be designed to minimize the burden of fault injection on sym-  
4426 bolic execution, working towards reducing query complexity. It would also be interest-  
4427 ing to see how different fault encodings are in fact handled inside the solver to select  
4428 or design the best one. It could even be a mix of encoding.

4429 **A Hybrid Forking/Forkless Technique.** We have seen in Chapter 6 that not all pro-  
4430 grams and not all fault models respond best to a forkless fault encoding. At first glance,  
4431 forkless encodings seem best suited for data faults while forking encodings could be fa-  
4432 vored for control-flow faults. It would be interesting to characterize what types of  
4433 programs respond best to a forkless technique and which respond best to a forking  
4434 one. Going further, we believe that inside a program, some possible fault locations  
4435 could respond best to one technique or the other. Hence, a future research direction  
4436 would be to design a hybrid forkless/forking technique and heuristics to choose which  
4437 technique to use on which possible fault location. Interestingly, this problem relates  
4438 to path merging in symbolic execution, where heuristic split cases into either fork the  
4439 analysis (analog to forking fault injection) or keep a merged state (analog to forkless  
4440 fault injection). The best performance for the path merging problem is found with a  
4441 hybrid technique, motivating this future research direction.

4442 **Finding the Minimal Attacker.** The approach presented in this thesis focuses on as-

4443 sessing the vulnerability or resistance of a program to an attacker model. The problem  
4444 can also be seen from the opposite perspective, that is to say, having an analysis find a  
4445 ‘minimal’ attacker model for which a program is vulnerable. Here, an order that makes  
4446 sense between various fault models and number of faults considered would have to be  
4447 coined.



## Éléments de traduction en français

*French summary of the thesis. The remaining of this appendix will be written in French.*

Une introduction substantielle est proposée Section [A.1](#), suivie d'un résumé de

chaque chapitre Section [A.2](#), terminant par une conclusion Section [A.3](#).

### **A.1 Chapitre 1 : Introduction**

**Contexte.** Des programmes se trouvent partout dans nos vies : des smartphones dans nos poches aux ordinateurs sur lesquels nous travaillons, nos télévisions, les systèmes informatiques du supermarché du coin, mais aussi les puces de nos cartes de crédits, nos systèmes de transport, de nos hôpitaux, les systèmes de distribution d'eau, d'électricité, etc. Ces programmes peuvent contenir des fonctionnalités de sécurité, dont le mauvais fonctionnement peut avoir de graves conséquences en termes monétaires, matérielles, humaines, ou en termes de fuite de secret.

Cependant, les programmes étant écrits à la main par des développeurs, des erreurs s'y immiscent. Ces bogues peuvent ensuite être exploités par des attaquants, par exemple, utiliser un dépassement de tampon pour lire ou écrire des données, ou exécuter du code arbitraire. Des protections existent mais restent souvent incomplètes et le programme continue à être vulnérable face à un attaquant avancé, par ailleurs capable d'utiliser des vecteurs d'attaque autres que les vulnérabilités logicielles tel que l'injection de faute durant l'exécution du programme. De nouveaux vecteurs et techniques d'attaque sont régulièrement découverts.

De nombreux travaux de recherche, ces dernières décennies, se sont intéressés au problème de *l'analyse automatique de programme* afin de les rendre plus robustes. De nombreuses techniques ont été développées telles que l'exécution symbolique [[CS13](#), [GLM12](#), [BGM13](#)], l'analyse statique [[Fac](#)], l'interprétation abstraite [[CGJ+03](#)] ou l'analyse bornée de modèles [[CBRZ01](#)], détectant des bogues ou cherchant à montrer leur absence [[CCF+05](#), [KKP+15](#)], conduisant à une adoption industrielle [[BGM13](#), [Fac](#), [BCLR04](#), [KKP+15](#), [LRV+22](#)]. S'assurer qu'un programme fonctionne comme voulu, c'est-à-dire sans bogue, relève de la *sûreté* de fonctionnement. Comme les bogues sont une porte d'entrée aux attaques, les supprimer est un premier pas vers la *sécurité* logicielle.

**Problème.** Ces analyses de programmes pour la sûreté se reposent sur un modèle d'attaquant faible, seulement capable de concevoir des entrées malveillantes. En pra-

4481 tique, un attaquant avancé est capable d’exploiter des vulnérabilités complexes pour  
4482 monter des attaques micro-architecturales ou encore des attaques matérielles, physiques  
4483 ou contrôlées par logiciel tel que Rowhammer. Il peut également combiner ces différents  
4484 vecteurs d’attaque.

4485 Pour construire une analyse de programme pour la sécurité, grâce à laquelle il  
4486 serait possible de rejouer des attaques, faisant varier différents paramètres et d’évaluer  
4487 des protections logicielles, nous avons identifié le besoin de réconcilier analyse de  
4488 programme et modèle d’attaquant avancé. Cependant, prendre en compte les nom-  
4489 breuses actions possibles d’un attaquant augment significativement l’espace d’état du  
4490 programme à analyser, menaçant la capacité de l’analyse à passer à l’échelle.

4491 **Objectif et Défis.** *Notre objectif est de concevoir une technique automatique, efficace et*  
4492 *générique pour raisonner sur l’impact d’un attaquant avancé sur les propriétés de sécu-*  
4493 *rité d’un programme* alors que les analyses standards ne considèrent qu’un attaquant  
4494 capable de créer des entrées malicieuses.

4495 Notre premier défi consiste à modéliser un attaquant avancé et à imaginer une  
4496 formalisation de son impact sur l’exécution d’un programme. Notre second défi est  
4497 de concevoir un algorithme efficace et générique pour étudier la vulnérabilité d’un  
4498 programme à un modèle d’attaquant donné, tout en maîtrisant la complexité inhérente  
4499 à l’ajout d’actions possible pour l’attaquant, en particulier lorsque plusieurs actions  
4500 sont considérées durant une seule attaque.

4501 Les rares techniques existantes dans le domaine se concentrent sur les injections de  
4502 fautes matérielles pour des composants de sécurité. Elles se basent sur des techniques  
4503 comme la *génération de mutants* [CCG13, RG14, GWJLL17, CDFG18, GWJL20] ou  
4504 des *analyses branchantes* [PMPD14, BBC<sup>+</sup>14, BHE<sup>+</sup>19, LFBP21, Lan22]. Ces tech-  
4505 niques souffrent de difficultés de passage à l’échelle, en particulier pour considérer  
4506 plusieurs fautes lors d’une attaque. De plus, elles sont généralement limitées à quelques  
4507 modèles de fautes pré-définis.

4508 **Proposition.** Nous proposons un modèle caractérisant un attaquant avancé en fonction  
4509 de ses capacités à modifier le comportement du programme et de son objectif d’attaque.  
4510 Afin de pouvoir raisonner sur ce modèle d’attaquant, nous présentons une extension  
4511 du système de transitions représentant l’exécution d’un programme avec de nouvelles  
4512 transitions modélisant les capacités de l’attaquant. *L’atteignabilité adversariale* est un  
4513 formalisme exprimant l’atteignabilité de l’objectif de l’attaquant dans ce système de  
4514 transitions étendu.

4515 Nous construisons un nouvel algorithme basé sur l’exécution symbolique, appelé  
4516 *exécution symbolique adversariale*, pour répondre au problème de l’atteignabilité ad-  
4517 versariale selon l’angle de la détection de bogues (vérification bornée). Notre algo-  
4518 rithme est générique en termes de support de capacités attaquant et permet d’atténuer  
4519 l’explosion du nombre de chemins à analyser grâce à un encodage non-branchant des  
4520 capacités de l’attaquant, présentées comme des fautes injectées. Afin d’améliorer da-  
4521 vantage les performances, nous proposons deux optimisations permettant de réduire le  
4522 nombre de fautes injectées. Nous montrons enfin que notre algorithme est correct et  
4523 k-complet pour l’atteignabilité adversariale.

4524 Nous implémentons cet algorithme pour l’analyse de programme binaire dans la  
4525 plateforme BINSEC d’exécution symbolique, construisant BINSEC/ASE.

4526 **Contributions.** Dans cette thèse, nous revendiquons les contributions suivantes :

- 4527 – Nous proposons un modèle représentant un attaquant avancé et nous formalisons  
4528 son impact sur les propriétés d’un programme avec *l’atteignabilité adversariale*,

4529 une extension du concept d’atteignabilité en présence d’un attaquant avancé.  
 4530 Nous définissons les propriétés de correction et de complétude associées ;  
 4531 – Nous décrivons *l’exécution symbolique adversariale*, une nouvelle technique sym-  
 4532 bolique évaluant l’atteignabilité adversariale. L’objectif de cet algorithme est  
 4533 d’être automatique et efficace dans son raisonnement sur l’impact d’un attaquant  
 4534 avancé sur un programme, et d’être générique quant au support du modèle  
 4535 d’attaquant. Notre algorithme comprend un nouvel encodage non-branchant des  
 4536 capacités de l’attaquant pour atténuer l’explosion de chemins et deux optimi-  
 4537 sations pour limiter le nombre d’injections tout en préservant le même modèle  
 4538 d’attaquant. Nous montrons que notre algorithme est correct et k-complet pour  
 4539 l’atteignabilité adversariale ;  
 4540 – Nous proposons une implémentation de l’exécution symbolique adversariale pour  
 4541 l’analyse de programme binaire au sein de la plateforme BINSEC, construisant  
 4542 un outil appelé *BINSEC/ASE*. Nous procédons à une évaluation expérimentale  
 4543 de la performance de BINSEC/ASE comparé à l’état de l’art en utilisant des  
 4544 programmes de test issus du domaine de l’injection de fautes matérielles et des  
 4545 cartes à puce. Nous montrons un gain significatif en termes de temps d’analyse et  
 4546 de chemins explorés. Nous soulignons l’intérêt et la faisabilité de notre technique  
 4547 en explorant plusieurs scénarios de sécurité. Enfin, nous étudions le cas du pro-  
 4548 gramme de démarrage de WooKey, accomplissant une étude de vulnérabilité de  
 4549 deux implémentations et proposons un correctif pour une attaque non rapportée  
 4550 précédemment.

4551 Ces travaux sont un premier pas vers la conception d’analyse de programme efficace  
 4552 pour la sécurité, prenant en compte un attaquant avancé. Notre approche est suffisam-  
 4553 ment générique pour inclure de nombreuses capacités attaquant comme l’inversion de  
 4554 bit de Rowhammer, l’inversion de test, le saut d’instruction ou la corruption arbitraire  
 4555 de données. Cependant, la modification d’instructions reste hors de notre portée. Alors  
 4556 que nous avons investigué le côté détection de vulnérabilité (sous-approximation) de  
 4557 l’atteignabilité adversariale, le côté vérification (sur-approximation) serait intéressant  
 4558 à explorer également. Ce sont des directions passionnantes pour de futurs travaux de  
 4559 recherche.

4560 Une partie des travaux présentés dans cette thèse ont été publiés à ESOP 2023  
 4561 [DBP23].

## 4562 A.2 Résumés de chaque chapitre

### 4563 A.2.1 Chapitre 2 : Contexte et Motivation

4564 Ce chapitre décrit le contexte des systèmes d’information implémentant des fonctionnal-  
 4565 ités de sécurité, bien plus répandus que l’on peut le penser, ainsi que les vulnérabilités  
 4566 auxquelles ils font face et les processus de sécurisation mis en place. Nous motivons  
 4567 le besoin d’intégrer un modèle d’attaquant avancé dans les analyses de programme  
 4568 en listant quelques vecteurs d’attaque à la disposition d’un tel attaquant, notamment  
 4569 l’injection de faute, et les capacités obtenues.

### 4570 **A.2.2 Chapitre 3 : Préambule**

4571 Dans ce chapitre, nous discutons de quelques notions importantes pour la compréhens-  
4572 sion de la suite du manuscrit. L'analyse de programme, les différents type d'approximations  
4573 et quelques propriétés qui peuvent être vérifiées dans un programme sont présentés.  
4574 En particulier, la notion d'atteignabilité y est définie. Nous présentons une technique  
4575 d'analyse de programme, l'exécution symbolique, utilisée dans cette thèse. Nous abor-  
4576 dons ensuite les particularités de l'analyse de programme au niveau binaire, et de  
4577 l'analyse de programme pour l'injection de faute.

### 4578 **A.2.3 Chapitre 4 : Atteignabilité Adversariale**

4579 Ce chapitre débute par la formalisation du modèle d'attaquant utilisé, composé (1) des  
4580 capacités d'attaque, (2) d'un nombre maximal d'actions et (3) d'un objectif d'attaque.  
4581 Nous intégrons ce modèle d'attaquant dans la notion d'atteignabilité, définissant ainsi la  
4582 notion d'atteignabilité adversariale, en étendant la représentation d'un programme sous  
4583 forme d'un système de transition avec de nouvelles transitions adversariales. La correc-  
4584 tion et complétude d'une analyse pour l'atteignabilité adversariale sont définies. Un al-  
4585 gorithme de vérification de l'atteignabilité adversariale est ensuite proposée, l'exécution  
4586 symbolique adversariale, basée sur l'exécution symbolique. Les actions de l'attaquant  
4587 sont modélisées par des fautes, encodées de façon non branchante dans la construction  
4588 du prédicat de chemin. Deux optimisations permettant de réduire le nombre de fautes  
4589 dans les requêtes au solveur sont décrites.

### 4590 **A.2.4 Chapitre 5 : le Prototype BINSEC/ASE**

4591 Dans ce chapitre, nous décrivons l'implémentation de notre algorithme d'exécution  
4592 symbolique adversariale dans l'outil BINSEC, pour l'analyse de code au niveau bi-  
4593 naire. En particulier, nous changeons la manière dont sont traitées les instructions  
4594 pour y inclure la possibilité d'une faute. Nous détaillons également l'implémentation  
4595 de la technique branchant dans BINSEC, qui construit le compétiteur face auquel nous  
4596 comparer. Enfin, nous proposons un guide utilisateur pour utiliser BINSEC/ASE avec  
4597 un nouveau programme, et un guide développeur pour ajouter de nouveaux modèles  
4598 de fautes.

### 4599 **A.2.5 Chapitre 6 : Evaluation Expérimentale**

4600 Ce chapitre contient l'évaluation expérimentale de notre outil BINSEC/ASE. En par-  
4601 ticulier, nous montrons que notre technique d'injection améliore significativement les  
4602 performances de l'analyse, d'autant plus lorsque le nombre de fautes considéré est élevé.  
4603 Notre technique permet d'atténuer l'explosion du nombre de chemins à laquelle l'état  
4604 de l'art fait face.

4605 Nous montrons également l'intérêt et la faisabilité de notre technique en explo-  
4606 rant différents scénarios de sécurité : reproduire une attaque BellCoRe sur CRT-RSA,  
4607 étudier la différence de vulnérabilité de deux implémentations d'une boîte à secret,  
4608 chercher des attaques sur la contre-mesure SecSwift, étudier l'impact d'un attaquant  
4609 sur un réseau de neurones, et enfin, analyser le cas du programme de démarrage de  
4610 WooKey, un challenge de sécurité de l'ANSSI. Nous sommes capables de rejouer des  
4611 attaques et avons trouvé une attaque non rapportée dans un correctif récent.

## 4612 A.3 Chapitre 7 : Conclusion

4613 Les techniques d'analyse de programme se concentrent majoritairement sur l'aspect  
4614 sûreté et considèrent un modèle d'attaquant faible, seulement capable de créer des  
4615 entrées malicieuses pour les programmes. En pratique, un attaquant puissant peut  
4616 tirer parti de nombreux vecteurs d'attaques, menaçant la sécurité des programmes.  
4617 Dans cette thèse, nous visons à porter l'analyse de programme de la sûreté à la sécurité  
4618 par la modélisation et l'intégration d'un modèle d'attaquant avancé et de son impact  
4619 sur les propriétés de sécurité d'un programme.

4620 Nous avons d'abord défini un modèle d'attaquant représentant un attaquant avancé  
4621 et nous avons formalisé son impact sur l'exécution d'un programme au travers de  
4622 l'atteignabilité adversariale qui étend le système de transition représentant un pro-  
4623 gramme avec des transitions adversariales modélisant les actions de l'attaquant.

4624 Nous avons ensuite conçu un algorithme, l'exécution symbolique adversariale pour  
4625 répondre au problème de l'atteignabilité adversariale. Les techniques de l'état de  
4626 l'art souffrent d'une explosion du nombre de chemins à explorer. Nous atténuons  
4627 ce phénomène par un encodage non branchant des capacités de l'attaquant et deux  
4628 optimisations permettant de réduire la complexité des requêtes au solveur.

4629 Puis, nous avons implémenté cet algorithme dans le moteur d'exécution symbolique  
4630 BINSEC, créant une analyse de sécurité au niveau binaire. Notre évaluation expérimen-  
4631 tale montre une amélioration significative des performances de l'analyse par rapport  
4632 à la technique branchante de l'état de l'art, en parallèle d'une réduction manifeste du  
4633 nombre de chemins à explorer.

4634 Enfin, nous illustrons l'intérêt et la faisabilité de notre technique au travers de  
4635 différents scénarios de sécurité. Nous étudions également le cas du programme de  
4636 démarrage de WooKey, un challenge de sécurité de l'ANSSI. Nous sommes capables  
4637 de trouver des attaques connues, et avons même mis en évidence une attaque non  
4638 rapportée précédemment sur un correctif récent et proposé notre propre correctif.

4639 **Travaux Futurs.** Nous présentons maintenant quelques perspectives de travaux futurs  
4640 étendant ou améliorant les travaux présentés dans cette thèse :

- 4641 – Étendre le support de BINSEC/ASE à de nouveaux modèles de faute, avec un  
4642 besoin d'algorithmes efficaces. Par exemple, notre formalisation est générique et  
4643 permet de modéliser un attaquant avec plusieurs types d'actions différentes, ce  
4644 qui n'est, à l'heure actuelle, pas implémenté dans notre outil. Il serait égale-  
4645 ment intéressant de poursuivre la recherche d'optimisations limitant en amont le  
4646 nombre de points d'injection possibles ;
- 4647 – Certains programmes, modèles de fautes, ou peut-être même points d'injection  
4648 se révèlent plus long à traiter avec une analyse non branchante. À première vue,  
4649 une analyse non branchante semble être plus performante sur des fautes sur les  
4650 données, alors qu'une analyse branchante pourrait avoir de meilleur résultat pour  
4651 des fautes sur le flow de control. Il serait intéressant de chercher des heuristiques  
4652 pour déterminer quand utiliser l'une ou l'autre des techniques, créant une analyse  
4653 hybride ;
- 4654 – Dans cette thèse, nous nous intéressons à montrer la vulnérabilité ou la résistance  
4655 d'un programme à un modèle d'attaquant donné. À l'inverse, on pourrait chercher  
4656 quel est l'attaquant "minimal" pouvant attaquer un programme. Il serait alors  
4657 nécessaire de proposer une relation d'ordre entre les modèles de faute et leurs  
4658 nombres qui fasse sens.



## Additional Experimental Data

In this appendix chapter, we include additional experimental data that did not fit in Chapter 6. The interested reader will find:

- A deeper evaluation of the reset fault model in Section B.1 (RQ B1), assessing the impact of our optimizations (RQ B1.1) in Section B.1.1 and the impact of the forkless encoding (RQ B1.2) in Section B.1.2;
- An evaluation of the impact of our optimizations on bit-flip faults (RQ B2) in Section B.2;
- An evaluation of the impact of our optimizations on instruction skip faults (RQ B3) in Section B.3;
- A summary of our optimization evaluation across fault models in Section B.4;
- An evaluation of the affinity of different solvers to our various arbitrary data forkless encodings (RQ B4) in Section B.5.

### B.1 FASE Evaluation of Reset Faults (RQ B1)

We evaluate BINSEC/ASE for the reset fault model regarding the impact of our optimizations and the impact of the chosen forkless encoding. We assess them independently. Our aim is to verify the extensibility of arbitrary data fault results.

#### B.1.1 Impact of Optimizations (RQ B1.1)

We evaluate now the impact of the different optimizations, EDS, IOD and their combination, IOD+EDS, for reset faults.

**RQ B1.1.** What is the impact of our optimizations on reset faults?

**Goal.** The goal of this experiment is to evaluate our optimizations for the reset fault model and evaluate how they contribute to reducing query complexity. In particular, we want to check whether general trends for arbitrary data fault evaluation also hold for reset faults.

**Protocol.** We evaluate FASE on the performance benchmark, with an attacker model able to perform 1 to 10 reset faults.

**Results.** We start by checking which is the fastest optimization. Analysis time results are presented in Table B.1 (p.VIII). On average for all numbers of faults, FASE-EDS is x1.05 times faster than FASE, FASE-IOD x1.6 times faster and FASE-EDS+IOD x1.4 times faster.

Table B.1: Analysis time for reset faults (RQ B1.1)

		Analysis time (s)						
		1f	2f	3f	4f	6f	8f	10f
		average value						
FASE-Reset		0.815	2.32	4.53	7.48	12.4	17.5	19.6
FASE-Reset-EDS		0.561	1.85	4.42	7.56	13.9	20.0	24.0
FASE-Reset-IOD		0.363	1.15	2.81	5.04	9.78	13.3	15.2
FASE-Reset-EDS+IOD		0.414	1.3	3.15	5.65	10.8	14.9	16.9
Forking-Reset		0.922	18.8	255	2.82 <i>k</i>	3.69 <i>k</i>	27.8 <i>k</i>	11.3 <i>k</i>
		timeouts (24h) over 12 benchmarks in total						
FASE-Reset		0	0	0	0	0	0	0
FASE-Reset-EDS		0	0	0	0	0	0	0
FASE-Reset-IOD		0	0	0	0	0	0	0
FASE-Reset-EDS+IOD		0	0	0	0	0	0	0
Forking-Reset*		0	0	0	0	9	9	11

(Forking-Reset\*: values for 6 faults and above are computed for incomplete runs due to timeouts)

4691 To measure the impact of our optimizations on query complexity, we consider the  
 4692 ratio of queries created on the number of queries sent to the solver, presented in Table  
 4693 B.2 (p.IX). This provides information about the complexity of the created queries and  
 4694 how many can be solved to true or false with arithmetic rules inside BINSEC without  
 4695 a solver call. We can see a trend similar to arbitrary data faults, where optimizations  
 4696 increase the number of queries created by the analysis, but allow for simpler queries,  
 4697 as many more are simplified without the solver’s help.

4698 Finally, our second measure of query complexity is the average solving time per  
 4699 query, assessing the complexity of queries for those that need to be sent to the solver.  
 4700 Results for average solving time per query are presented in Table B.4 (p.X). The average  
 4701 solving time per query also decreases with optimizations. However, while FASE-IOD  
 4702 managed to get to the level of the forking technique for arbitrary data faults, here the  
 4703 forking technique remains x2.3 times faster than FASE-IOD on average for all numbers  
 4704 of faults.

**Conclusion RQ B1.1.** Reset faults experimentally show trends in performance similar to arbitrary data faults, with FASE-IOD still the best optimization.

Table B.2: Number of queries created and sent to the solver for reset faults (RQ B1.1)

	Number of queries created / Number of queries sent to the solver										
	1f	2f	3f	4f	6f	8f	10f				
	average value										
FASE-Reset	185 / 80.8	465 / 204	899 / 383	1.38k / 583	2.16k / 939	2.64k / 1.21k	2.9k / 1.36k				
FASE-Reset-EDS	198 / 65.6	518 / 211	1.04k / 441	1.65k / 702	2.66k / 1.19k	3.3k / 1.56k	3.66k / 1.8k				
FASE-Reset-IOD	198 / 48.2	527 / 176	1.08k / 410	1.76k / 712	2.9k / 1.29k	3.59k / 1.71k	3.96k / 1.95k				
FASE-Reset-EDS+IOD	205 / 53.0	555 / 191	1.15k / 444	1.87k / 765	3.11k / 1.38k	3.87k / 1.85k	4.32k / 2.15k				
Forking-Reset*	1.34k / 173	29.2k / 4.08k	476k / 69.0k	6.22M / 963k	8.27M / 2.04M	61.2M / 20.8M	13.1M / 6.0M				

(Forking-Reset\*: values for 6 faults and above are computed for incomplete runs due to timeouts)

Table B.3: Number of queries created and sent to the solver comparison for reset fault encodings (RQ B1.2)

	Number of queries created / Number of queries sent to the solver										
	1f	2f	3f	4f	6f	8f	10f				
	average value										
FASE-Reset-Ite	185 / 80.8	465 / 204	899 / 383	1.38k / 583	2.16k / 939	2.64k / 1.21k	2.9k / 1.36k				
FASE-Reset-Sub	185 / 103	465 / 271	899 / 521	1.38k / 801	2.16k / 1.26k	2.64k / 1.59k	2.9k / 1.79k				
FASE-Reset-And	185 / 103	465 / 271	899 / 521	1.38k / 801	2.16k / 1.26k	2.64k / 1.59k	2.9k / 1.79k				
FASE-Reset-Xor	185 / 103	465 / 271	899 / 521	1.38k / 801	2.16k / 1.26k	2.64k / 1.59k	2.9k / 1.79k				

Table B.4: Average solving time per query for reset faults (RQ B1.1)

Average solving time per query (s)							
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE-Reset	0.007	0.009	0.009	0.009	0.010	0.010	0.009
FASE-Reset-EDS	0.007	0.008	0.009	0.009	0.009	0.009	0.009
FASE-Reset-IOD	0.005	0.006	0.007	0.006	0.006	0.007	0.006
FASE-Reset-EDS+IOD	0.005	0.007	0.007	0.006	0.007	0.007	0.006
Forking-Reset*	0.003	0.003	0.003	0.003	0.001	0.001	0.002

(Forking-Reset\*: values for 6 faults and above are computed for incomplete runs due to timeouts)

## 4705 B.1.2 Comparison of the Different Forkless Encodings (RQ B1.2)

4706 Multiple forkless encodings have been implemented for reset faults (see Section 5.3.3.1).  
 4707 We compare them here.

4708 **RQ B1.2.** What is the impact of the different encodings on the solver’s performance?

4709 **Goal.** Our goal is to explore the impact of different operators in the forkless encoding  
 4710 and their affinity with the solver used by BINSEC/ASE. We check whether the ite  
 4711 operator also makes for the fastest forkless encoding with reset faults.

4712 **Protocol.** We evaluate FASE on the performance benchmark, with an attacker model  
 4713 able to perform 1 to 10 reset faults. We compare four encodings, FASE-Reset-Ite,  
 4714 FASE-Reset-Sub (with a munis operator), FASE-Reset-And (with a logical and opera-  
 4715 tor) and FASE-Reset-Xor (with a logical xor operator), on three metrics. No optimiza-  
 4716 tions were used.

4717 **Results.** First, analysis time results are presented in Table B.5 (p.X). In terms of  
 4718 overall analysis time, FASE-Reset-Ite is x2.1 times faster than FASE-Reset-Xor for 1  
 4719 fault on average, x3.5 times faster than FASE-Reset-And and x3.4 times faster than  
 FASE-Reset-Sub.

Table B.5: Analysis time for reset faults (RQ B1.2)

Analysis time (s)							
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE-Reset-Ite	0.815	2.32	4.53	7.48	12.4	17.5	19.6
FASE-Reset-Sub	2.79	7.03	14.0	23.5	39.4	52.0	55.2
FASE-Reset-And	2.86	7.78	15.8	25.9	45.7	59.7	71.6
FASE-Reset-Xor	1.72	4.87	10.0	16.2	26.9	34.2	39.3

4720

4721 Then, looking at query complexity, our first metric is the ratio of queries created on  
 4722 the number of queries sent, presented in Table B.3 (p.IX), reflecting query complexity  
 4723 at query creation inside BINSEC, before solver call. FASE-Reset-Ite enables more  
 4724 arithmetical simplifications preventing a solver call for all fault numbers. FASE-Reset-  
 4725 Ite simplifies 56% of queries, while the other 3 encodings only simplify 44% of queries.  
 4726 This difference is likely due to the various arithmetic simplification rules implemented

4727 in BINSEC, not treating concrete constant values (in reset faults) like more general  
4728 expressions (in arbitrary data faults).

4729 Finally, the average solving time per query results are presented in Table B.6 (p.XI),  
4730 to reflect the query complexity in solver calls. Due to variability in the measurements,  
4731 100 runs were averaged for each program of the benchmark. FASE-Reset-Ite is also  
4732 the encoding generating queries that take the least time to solve. Averaging over all  
4733 number of faults, FASE-Reset-Ite is x3.4 times faster than FASE-RESET-Sub, x3.7  
times faster than FASE-Reset-And and x1.8 times faster than FASE-Reset-Xor.

Table B.6: Average solving time per query for reset faults (RQ B1.2)

Average solving time per query (s)							
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE-Reset-Ite	0.007	0.009	0.009	0.009	0.010	0.010	0.009
FASE-Reset-Sub	0.031	0.030	0.029	0.035	0.028	0.028	0.029
FASE-Reset-And	0.032	0.036	0.035	0.030	0.031	0.033	0.035
FASE-Reset-Xor	0.014	0.016	0.017	0.017	0.017	0.017	0.017

4734

**Conclusion RQ B1.2.** The encoding using the ite operator is the best-performing encoding for reset faults and arbitrary data faults for our BINSEC/ASE implementation and the solver it uses.

## 4735 B.2 FASE Evaluation of Bit-Flip Faults (RQ B2)

4736 In this section, we assess the last data fault model implemented by evaluating our  
4737 different optimizations (FASE+EDS, FASE+IOD and FASE-EDS+IOD) and FASE.  
4738 We have only one forkless encoding implemented, using an ite and a shift operator, for  
4739 the bit-flip fault model.

4740 **RQ B2.** What is the impact of our optimizations on bit-flip faults?

4741 **Goal.** The goal of this experiment is to check that FASE-IOD also performs best for  
4742 bit-flips and evaluate how optimizations contribute to reducing query complexity.

4743 **Protocol.** We evaluate FASE on the performance benchmark, with an attacker model  
4744 able to perform 1 to 10 bit-flip faults.

4745 **Results.** We present the analysis time results in Table B.7 (p.XII). FASE IOD is x2.2  
4746 times faster than FASE, x1.4 times faster than FASE-EDS and x1.1 times faster than  
4747 FASE-EDS+IOD.

4748 To compare query complexity, we start with the ratio of queries created over queries  
4749 sent to the solver, presented in Table B.9 (p.XIII). As expected, optimizations generate  
4750 some more queries than FASE, x1.6 times for FASE-EDS and FASE-IOD for 1 fault,  
4751 and x1.1 times more for FASE-EDS+IOD. Despite this, they succeed in reducing query  
4752 complexity as more queries are arithmetically resolved to true or false by BINSEC  
4753 without the help of the solver. FASE-EDS sends x1.1 times fewer queries to the solver,  
4754 FASE-IOD x1.4 times fewer and FASE-EDS+IOD x1.2 times fewer. As the number of  
4755 faults increases, the ratio of queries simplified without a solver call increases also.

Table B.7: Analysis time for bit-flip faults (RQ3.2)

Analysis time (s)							
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE	2.71	5.42	7.97	9.84	13.0	13.8	14.2
FASE-EDS	1.71	4.58	7.87	10.3	15.1	15.2	16.0
FASE-IOD	1.25	2.89	4.69	6.35	8.76	9.27	9.15
FASE-EDS+IOD	1.4	3.82	5.25	6.9	9.51	10.4	10.5
Forking*	0.645	15.1	225	952	2.13k	13.4k	9.47k
timeouts (24h) over 12 benchmarks in total							
FASE	0	0	0	0	0	0	0
FASE-EDS	0	0	0	0	0	0	0
FASE-IOD	0	0	0	0	0	0	0
FASE-EDS+IOD	0	0	0	0	0	0	0
Forking	0	0	0	2	7	8	10

(Forking\*: values for 4 faults and above are computed for incomplete runs due to timeouts)

4756 Our second measure of query complexity is the average solving time per query,  
 4757 presented in Table B.8 (p.XII). On average over all faults, FASE-EDS queries are x1.04  
 faster to solve, FASE-IOD x1.9 times and FASE-EDS+IOD x1.8 times faster.

Table B.8: Average solving time per query for bit-flip faults (RQ B2)

Average solving time per query (s)							
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE	0.050	0.053	0.052	0.046	0.050	0.050	0.050
FASE-EDS	0.035	0.045	0.060	0.049	0.057	0.048	0.053
FASE-IOD	0.035	0.024	0.024	0.024	0.026	0.026	0.027
FASE-EDS+IOD	0.034	0.037	0.024	0.024	0.025	0.026	0.027
Forking*	0.003	0.004	0.003	0.002	0.003	0.004	0.004

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts)

4758

**Conclusion RQ B2.** FASE-IOD is still the fastest optimization for the bit-flip fault model.

Table B.9: Number of queries created and sent to the solver for bit-flip faults (RQ B2)

	Number of queries created / Number of queries sent to the solver						
	1f	2f	3f	4f	6f	8f	10f
	average value						
FASE	161 / 86.1	303 / 162	435 / 228	547 / 287	676 / 358	715 / 377	726 / 382
FASE-EDS	171 / 78.0	330 / 165	480 / 243	609 / 313	766 / 407	813 / 432	826 / 439
FASE-IOD	171 / 60.8	345 / 144	517 / 222	660 / 292	823 / 370	870 / 388	883 / 393
FASE-EDS+IOD	181 / 70.8	372 / 166	557 / 250	712 / 321	902 / 416	960 / 440	976 / 447
Forking*	1.06k / 237	17.8k / 5.24k	204k / 68.0k	1.05M / 349k	2.09M / 703k	8.18M / 3.07M	5.19M / 2.15M

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts)

Table B.10: Number of queries created and sent to the solver comparison for instruction skip faults (RQ B3)

	Number of queries created / Number of queries sent to the solver						
	1f	2f	3f	4f	6f	8f	10f
	average value						
FASE*	274 / 156	951 / 535	2.52k / 1.32k	5.41k / 2.76k	14.0k / 7.29k	20.8k / 10.5k	25.9k / 12.7k
FASE-IOD*	352 / 114	1.25k / 519	3.3k / 1.4k	6.99k / 2.94k	17.3k / 7.44k	24.5k / 10.1k	30.5k / 12.4k
Forking*	2.8k / 474	133k / 34.0k	2.76M / 496k	69.0M / 4.67M	46.7M / 13.6M	181M / 32.3M	133M / 33.8M

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts)

FASE\*: all value from incomplete runs)

### B.3 FASE Evaluation of Instruction Skip Faults (RQ B3)

We now evaluate BINSEC/ASE’s performance for instruction skips. Only one forkless encoding has been implemented, using an ite operator. FASE support instruction skips in its un-optimized version, and with IOD.

**RQ B3.** What is the impact of our optimizations on instruction skip faults?

**Goal.** The goal of this experiment is to check whether IOD improves FASE’s performance and evaluate how it contributes to reducing query complexity.

**Protocol.** We evaluate FASE on the performance benchmark, with an attacker model able to perform 1 to 10 instruction skip faults.

**Results.** We start with the analysis time results, presented in Table B.11 (p.XIV). First, we can see that there is 1 program for which the instruction skip forkless cannot finish in time for FASE and FASE-IOD similarly. It is the unrolled version of VerifyPIN with a PIN size of 16. This is a rather long sequence of assignments that seem to overload the analysis when faulted in a forkless manner. As FASE and FASE-IOD timeout for the same benchmark, we can still compare their results in analysis time. We see that FASE-IOD is x1.1 times faster than FASE on average for 1 fault, which is less than for data fault models. Injection On Demand was designed with data faults in mind, we leave to future work to imagine an optimization tailored for this particular fault model.

Table B.11: Analysis time for instruction skip faults (RQ B3)

		Analysis time (s)						
		1f	2f	3f	4f	6f	8f	10f
		average value						
FASE*		107	175	343	625	1.73k	2.94k	3.23k
FASE-IOD*		95.6	150	329	515	822	842	962
Forking*		9.6	857	3.52k	16.3k	32.1k	86.4k	86.4k
		timeouts (24h) over 12 benchmarks in total						
FASE		1	1	1	1	1	1	1
FASE-IOD		1	1	1	1	1	1	1
Forking		0	0	3	6	10	12	12

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts  
FASE\*: all value from incomplete runs)

Then, we evaluate FASE-IOD compared to FASE on query complexity metrics. The ratio of queries created over queries sent is presented in Table B.10 (p.XIII). As runs are incomplete for at least one benchmark, comparison of exact values does not make sense, but trends are preserved. Hence we can see that, as expected, FASE-IOD creates more queries but reduces query complexity enough so that it sends fewer to the solver than FASE.

Our second query complexity metric, the average solving time per query, is presented in Table B.12 (p.XV). As for other fault models, FASE-IOD allows to greatly reduce the average solving time per query. On average over all number of faults, FASE-IOD has queries x5.4 faster to solve than FASE.

Table B.12: Average solving time per query for instruction skip faults (RQ B3)

Average solving time per query (s)							
	1f	2f	3f	4f	6f	8f	10f
average value							
FASE*	0.051	0.085	0.138	0.143	0.135	0.132	0.140
FASE-IOD*	0.017	0.025	0.023	0.020	0.023	0.021	0.023
Forking*	0.010	0.012	0.006	0.005	0.002	0.003	0.003

(Forking\*: values for 3 faults and above are computed for incomplete runs due to timeouts  
 FASE\*: all value from incomplete runs)

**Conclusion RQ B3.** FASE-IOD also reduces query complexity for the instruction skip fault model, which allows for a faster analysis.

## 4788 B.4 FASE Optimizations Summary

4789 FASE-IOD is our best-performing optimization, and thus across all fault models for  
 4790 which it is implemented. This gives us hope that it will be efficient too for yet another  
 4791 fault model useful to a BINSEC/ASE user. It is interesting to see that it reduces  
 4792 query complexity in our two metrics, generating a high percentage of queries that  
 4793 can be resolved by BINSEC/ASE without a solver call, and the remainder is still less  
 4794 complex to solve.

## 4795 B.5 Influence of Solver on Encoding Operators (RQ B4)

4796 BINSEC/ASE inherits the native binding of BINSEC to the solver bitwuzla. Conse-  
 4797 quently, BINSEC/ASE results in a faster analysis with this native binding.

4798 In this section, we consider the influence of the operator used in the forkless encoding  
 4799 and the solver used, considering z3, boolector and bitwuzla without the native binding.

4800 **RQ B4.** What is the impact of the choice of the solver on the performance of the  
 4801 forkless encodings?

4802 **Goal.** The goal of this experiment is to check whether the forkless encoding using  
 4803 the ite operator results in the fastest analysis independently of the solver used, or if  
 4804 different solvers favor different operators.

4805 **Protocol.** We evaluate FASE without optimizations on the performance benchmark,  
 4806 with an attacker model able to perform 1 to 4 arbitrary data faults, with three different  
 4807 solvers (Bitwuzla without native binding, Boolector and z3), for four different encodings  
 4808 (Ite, Mul, And and Xor). We limit our experiment to 4 faults since some encodings  
 4809 start to experience timeouts after with the native binding that is faster than using  
 4810 external solvers.

4811 **Results.** Analysis time results and average solving time per query are presented in  
 4812 Table B.13. None of these analyses experimented a timeout.

4813 First, we can see that globally, Bitwuzla performs better than the other two solvers,  
 4814 even without the native binding, in analysis time and in average solving time per query.  
 4815 Bitwuzla is tuned for the specific SMT theory used by BINSEC (QF\_ABV), while

4816 Boolector and z3 are more generic in SMT theories, hence do not perform as well in  
 4817 this specific theory.

4818 Bitwuzla still favors the Ite encoding, so does z3, while Boolector favors the Xor  
 4819 encoding. This shows the encoding used and the chosen solver have a big influence  
 4820 on the analysis time. For instance, with z3, the best-performing encoding yields an  
 4821 analysis x3.4 times faster than the worst for 1 fault. The best encoding of Bitwuzla  
 results in an analysis x16.7 times faster than the slowest z3 encoding for 1 fault.

Table B.13: Analysis time and average solving time per query of arbitrary data faults with different solvers (RQ B4)

	1f	2f	3f	4f
Analysis time (s)				
Bitwuzla-Ite	10.4	38.7	130	358
Bitwuzla-Mul	13.8	56.7	252	979
Bitwuzla-And	14.1	59.6	311	1.26k
Bitwuzla-Xor	11.0	40.2	136	371
Boolector-Ite	25.0	125	423	1.22k
Boolector-Mul	38.8	198	717	2.21k
Boolector-And	40.4	219	836	2.4k
Boolector-Xor	23.5	113	377	1.05k
Z3-Ite	51.3	279	872	2.12k
Z3-Mul	147	761	1.97k	4.3k
Z3-And	154	790	2.03k	4.49k
Z3-Xor	174	877	2.12k	4.38k
Average solving time per query (s)				
Bitwuzla-Ite	0.065	0.052	0.055	0.054
Bitwuzla-Mul	0.089	0.074	0.088	0.087
Bitwuzla-And	0.09	0.087	0.087	0.107
Bitwuzla-Xor	0.071	0.056	0.066	0.06
Boolector-Ite	0.198	0.224	0.231	0.241
Boolector-Mul	0.28	0.294	0.301	0.32
Boolector-And	0.278	0.299	0.31	0.34
Boolector-Xor	0.183	0.202	0.207	0.214
Z3-Ite	0.273	0.305	0.324	0.341
Z3-Mul	0.663	0.791	0.751	0.74
Z3-And	0.688	0.801	0.773	0.75
Z3-Xor	0.707	0.834	0.772	0.747

4822

**Conclusion RQ B4.** As SMT queries are a bottleneck for symbolic execution, it is important to select a solver tuned for the SMT theory used and to select operators to yield good analysis performance.

## Instrumentation Details

This appendix chapter aims to provide more information about the way we instrumented programs of our benchmark for evaluation (Section 6.5).

**Attacker Model.** We preserve the attack goal written with an *assert* in each benchmark program and do not fault its computation. We consider here a user-defined number of arbitrary data faults.

**General C Instrumentation.** Programs are written in C. In Table C.1, we illustrate how each type of C statement is instrumented.

- The right-hand side of assignments are wrapped into a utility function *inject\_AD*, injecting the fault;

- Conditional statements are expended. First, a new variable is created and holds the faulted conditional expression. It is this new variable that is inserted into the conditional statement;

- Return statements are instrumented similarly to assignments.

Table C.1: Instrumentation examples

C code	Instrumented C code
<code>int x = expr;</code>	<code>int x = inject_AD(expr);</code>
<code>if (cdt_expr) {     \\... }</code>	<code>int cdt_1 = inject_AD(cdt_expr); if (cdt_1) {     \\... }</code>
<code>return expr;</code>	<code>return inject_AD(expr);</code>

**Instrumentation Settings.** We add instrumentation settings as calls to instrumentation library functions from the main of the instrumented program as illustrated in Figure C.1 for a *VerifyPIN* program.

- We initialize instrumentation with the fault budget and whether we consider forkless or forking faults, by calling the *init\_injection* function;

- The body of the *verifyPIN* function if instrumented as presented above;

- A check constraining the total number of faults is inserted before the *assert* for forkless faults.

**Instrumentation Library.** The functions described in Figure C.2 are written in a separate file, serving as a library.

---

```

1 #define MAX_FAULTS 0
2 #define IS_FORKLESS false
3
4 int main() {
5     init_injection(MAX_FAULTS, IS_FORKLESS);
6     initialize();
7     verifyPIN();
8     int o = oracle();
9     fault_condition();
10    assert(o);
11    return 0;
12 }

```

Figure C.1: Instrumented VerifyPIN main function

- 4848 – The function *init\_injection* sets global variables holding the analysis injection
- 4849 parameters;
- 4850 – The function *fault\_condition* adds the fault budget constraint for forkless faults;
- 4851 – The *inject\_AD* function first creates a new symbolic boolean variable for the
- 4852 fault activation and adds it to the *sum* variable. If we consider forkless faults,
- 4853 a new symbolic variable is created for the effect of the arbitrary data fault, and
- 4854 a forkless fault encoding is returned, taking the place of the original expression.
- 4855 The forkless encoding chosen uses a And operator. In the case of forking faults,
- 4856 the fault budget is checked first, then the control flow is split depending on the
- 4857 symbolic value of *b*. If the fault happens, a new symbolic value is returned,
- 4858 otherwise, the original expression is returned.

4859 **Instrumentation Extension.** This instrumentation framework has been written for the  
4860 evaluation of arbitrary data faults. However, it can be extended to support more:

- 4861 – It is possible to extend this instrumentation framework to other data fault mod-
- 4862 els or other forkless encodings simply by modifying the fault encoding in the
- 4863 *inject\_AD* function;
- 4864 – Fault models other than data faults can also be implemented based on this tem-
- 4865 plate, they may require slightly different C instrumentation than what is pre-
- 4866 sented in Table C.1 though;
- 4867 – We believe EDS optimization can be implemented as well in this framework.
- 4868 Some work or rethinking of optimizations might be needed before implementing
- 4869 other optimizations.

```
1 int sum, max_faults;
2 bool is_forkless;
3
4 void init_injection(int max, bool ib) {
5     sum = 0;
6     max_faults = max;
7     is_forkless = ib;
8 }
9
10 bool fault_condition() {
11     if (is_forkless) {
12         klee_assume(sum <= max_faults);
13     }
14     return true;
15 }
16
17 int inject_AD(int x) {
18     int b;
19     klee_make_symbolic(&b, sizeof(int), "b");
20     klee_assume(b <= 1);
21     klee_assume(b >= 0);
22     sum += b;
23
24     if (is_forkless) {
25         int non_det;
26         klee_make_symbolic(&non_det, sizeof(int),
27                             "non_det");
28         klee_assume(non_det != 0);
29         return x + ((-b) & non_det);
30     }
31     else {
32         klee_assume(sum <= max_faults);
33         if (b){
34             int non_det;
35             klee_make_symbolic(&non_det, sizeof(int),
36                                 "non_det");
37             klee_assume(x != non_det);
38             return non_det;
39         }
40         return x;
41     }
42 }
```

Figure C.2: Instrumentation utility functions



# Bibliography

4870

- 4871 [AAE<sup>+</sup>20] ANSSI, Amosys, EDSI, LETI, Lexfo, Oppida, Quarkslab, SERMA,  
4872 Synacktiv, Thales, and Trusted Labs. Inter-cesti: Methodological and  
4873 technical feedbacks on hardware devices evaluations. In *SSTIC 2020,*  
4874 *Symposium sur la sécurité des technologies de l'information et des com-*  
4875 *munications*, 2020. [133](#)
- 4876 [AARR03] Dakshi Agrawal, Bruce Archambeault, Josyula R Rao, and Pankaj Ro-  
4877 hatgi. The em side—channel (s). In *Cryptographic Hardware and Embed-*  
4878 *ded Systems-CHES 2002: 4th International Workshop Redwood Shores,*  
4879 *CA, USA, August 13–15, 2002 Revised Papers 4*, pages 29–45. Springer,  
4880 2003. [17](#)
- 4881 [ABEL09] Martín Abadi, Mihai Buiu, Ulfar Erlingsson, and Jay Ligatti. Control-  
4882 flow integrity principles, implementations, and applications. *ACM Trans-*  
4883 *actions on Information and System Security (TISSEC)*, 13(1):1–40, 2009.  
4884 [19](#)
- 4885 [AC04] Martín Abadi and Véronique Cortier. Deciding knowledge in security  
4886 protocols under equational theories. In *International Colloquium on Au-*  
4887 *tomata, Languages, and Programming*, pages 46–58. Springer, 2004. [64](#)
- 4888 [ACD<sup>+</sup>22] Ihab Alshaer, Brice Colombier, Christophe Deleuze, Vincent Beroulle,  
4889 and Paolo Maistri. Variable-length instruction set: Feature or bug? In  
4890 *2022 25th Euromicro Conference on Digital System Design (DSD)*, pages  
4891 464–471. IEEE, 2022. [60](#)
- 4892 [AKS06] Onur Aciçmez, Çetin Kaya Koç, and Jean-Pierre Seifert. Predicting secret  
4893 keys via branch prediction. In *Topics in Cryptology-CT-RSA 2007: The*  
4894 *Cryptographers' Track at the RSA Conference 2007, San Francisco, CA,*  
4895 *USA, February 5-9, 2007. Proceedings*, pages 225–242. Springer, 2006. [17](#)
- 4896 [ASA<sup>+</sup>15] Adnan Akhunzada, Mehdi Sookhak, Nor Badrul Anuar, Abdullah Gani,  
4897 Ejaz Ahmed, Muhammad Shiraz, Steven Furnell, Amir Hayat, and  
4898 Muhammad Khurram Khan. Man-at-the-end attacks: Analysis, taxon-  
4899 omy, human aspects, motivation and future directions. *Journal of Network*  
4900 *and Computer Applications*, 48:44–57, 2015. [23](#), [64](#)
- 4901 [AVFM07] Frederic Amiel, Karine Villegas, Benoit Feix, and Louis Marcel. Passive  
4902 and active combined attacks: Combining fault attacks and side channel

- 4903 analysis. In *Workshop on Fault Diagnosis and Tolerance in Cryptography*  
4904 (*FDTC 2007*), pages 92–102. IEEE, 2007. [20](#), [46](#)
- 4905 [BB09] Robert Brummayer and Armin Biere. Boolector: An efficient smt solver  
4906 for bit-vectors and arrays. In *Tools and Algorithms for the Construction*  
4907 *and Analysis of Systems: 15th International Conference, TACAS 2009,*  
4908 *Held as Part of the Joint European Conferences on Theory and Practice*  
4909 *of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings 15,*  
4910 pages 174–177. Springer, 2009. [71](#)
- 4911 [BBB<sup>+</sup>10] Alessandro Barenghi, Guido M Bertoni, Luca Breveglieri, Mauro Pellicioli,  
4912 and Gerardo Pelosi. Low voltage fault attacks to aes. In *2010 IEEE Inter-*  
4913 *national Symposium on Hardware-Oriented Security and Trust (HOST),*  
4914 pages 7–12. IEEE, 2010. [20](#)
- 4915 [BBC<sup>+</sup>14] Maël Berthier, Julien Bringer, Hervé Chabanne, Thanh-Ha Le, Lionel  
4916 Rivièrè, and Victor Servant. Idea: embedded fault injection simulator on  
4917 smartcard. In *International Symposium on Engineering Secure Software*  
4918 *and Systems*, pages 222–229. Springer, 2014. [3](#), [35](#), [62](#), [63](#), [II](#)
- 4919 [BBKN12] Alessandro Barenghi, Luca Breveglieri, Israel Koren, and David Naccache.  
4920 Fault injection attacks on cryptographic devices: Theory, practice, and  
4921 countermeasures. *Proceedings of the IEEE*, 100(11):3056–3076, 2012. [19](#),  
4922 [20](#)
- 4923 [BBY17] Sandrine Blazy, David Bühler, and Boris Jakobowski. Structuring ab-  
4924 stract interpreters through state and value abstractions. In *International*  
4925 *Conference on Verification, Model Checking, and Abstract Interpretation,*  
4926 pages 112–130. Springer, 2017. [61](#), [62](#)
- 4927 [BCD<sup>+</sup>18] Roberto Baldoni, Emilio Coppa, Daniele Cono D’elia, Camil Demetrescu,  
4928 and Irene Finocchi. A survey of symbolic execution techniques. *ACM*  
4929 *Computing Surveys (CSUR)*, 51(3):1–39, 2018. [30](#)
- 4930 [BCDK14] Sébastien Bardin, Omar Chebaro, Mickaël Delahaye, and Nikolai Kosma-  
4931 tov. An all-in-one toolkit for automated white-box testing. In *Internat-*  
4932 *ional Conference on Tests and Proofs*, pages 53–60. Springer, 2014. [36](#),  
4933 [61](#)
- 4934 [BCE08] Peter Boonstoppel, Cristian Cadar, and Dawson Engler. Rwsset: Attack-  
4935 ing path explosion in constraint-based test generation. In *Tools and Algo-*  
4936 *rithms for the Construction and Analysis of Systems: 14th International*  
4937 *Conference, TACAS 2008, Held as Part of the Joint European Confer-*  
4938 *ences on Theory and Practice of Software, ETAPS 2008, Budapest, Hun-*  
4939 *gary, March 29-April 6, 2008. Proceedings 14,* pages 351–366. Springer,  
4940 2008. [33](#)
- 4941 [BCL14] Gergei Bana and Hubert Comon-Lundh. A computationally complete  
4942 symbolic attacker for equivalence properties. In *Proceedings of the 2014*  
4943 *ACM SIGSAC Conference on Computer and Communications Security,*  
4944 pages 609–620, 2014. [3](#), [64](#)

- 4945 [BCLR04] Thomas Ball, Byron Cook, Vladimir Levin, and Sriram K Rajamani. Slam  
4946 and static driver verifier: Technology transfer of formal methods inside  
4947 microsoft. In *International Conference on Integrated Formal Methods*,  
4948 pages 1–20. Springer, 2004. [2](#), [1](#)
- 4949 [BCN<sup>+</sup>17] Nathan Burow, Scott A Carr, Joseph Nash, Per Larsen, Michael Franz,  
4950 Stefan Brunthaler, and Mathias Payer. Control-flow integrity: Precision,  
4951 security, and performance. *ACM Computing Surveys (CSUR)*, 50(1):1–33,  
4952 2017. [19](#), [64](#)
- 4953 [BCR<sup>+</sup>19] Cataldo Basile, Daniele Canavese, Leonardo Regano, Paolo Falcarin, and  
4954 Bjorn De Sutter. A meta-model for software protections and reverse en-  
4955 gineering attacks. *Journal of Systems and Software*, 150:3–21, 2019. [64](#)
- 4956 [BD12] Leyla Bilge and Tudor Dumitraş. Before we knew it: an empirical study  
4957 of zero-day attacks in the real world. In *Proceedings of the 2012 ACM  
4958 conference on Computer and communications security*, pages 833–844,  
4959 2012. [13](#)
- 4960 [BDL97] Dan Boneh, Richard A DeMillo, and Richard J Lipton. On the importance  
4961 of checking cryptographic protocols for faults. In *International conference  
4962 on the theory and applications of cryptographic techniques*, pages 37–51.  
4963 Springer, 1997. [128](#)
- 4964 [BDL01] Dan Boneh, Richard A DeMillo, and Richard J Lipton. On the impor-  
4965 tance of eliminating errors in cryptographic computations. *Journal of  
4966 cryptology*, 14:101–119, 2001. [20](#)
- 4967 [BDM17] Sébastien Bardin, Robin David, and Jean-Yves Marion. Backward-  
4968 bounded dse: targeting infeasibility questions on obfuscated codes. In  
4969 *2017 IEEE Symposium on Security and Privacy (SP)*, pages 633–651.  
4970 IEEE, 2017. [70](#)
- 4971 [BECN<sup>+</sup>06] Hagai Bar-El, Hamid Choukri, David Naccache, Michael Tunstall, and  
4972 Claire Whelan. The sorcerer’s apprentice guide to fault attacks. *Proceed-  
4973 ings of the IEEE*, 94(2):370–382, 2006. [19](#)
- 4974 [BEMP20] Etienne Boespflug, Cristian Ene, Laurent Mounier, and Marie-Laure  
4975 Potet. Countermeasures optimization in multiple fault-injection context.  
4976 In *2020 Workshop on Fault Detection and Tolerance in Cryptography  
4977 (FDTC)*, pages 26–34. IEEE, 2020. [130](#)
- 4978 [Ber05] Daniel J Bernstein. Cache-timing attacks on aes. 2005. [17](#)
- 4979 [BFP19] Claudio Bozzato, Riccardo Focardi, and Francesco Palmari. Shaping  
4980 the glitch: optimizing voltage fault injection attacks. *IACR transactions  
4981 on cryptographic hardware and embedded systems*, pages 199–224, 2019.  
4982 [20](#)
- 4983 [BG22] Sébastien Bardin and Guillaume Girol. A quantitative flavour of robust  
4984 reachability. *arXiv preprint arXiv:2212.05244*, 2022. [61](#)

- 4985 [BGM13] Ella Bounimova, Patrice Godefroid, and David Molnar. Billions and bil-  
4986 lions of constraints: Whitebox fuzz testing in production. In *2013 35th*  
4987 *International Conference on Software Engineering (ICSE)*, pages 122–131.  
4988 IEEE, 2013. [2](#), [1](#)
- 4989 [BGV11] Josep Balasch, Benedikt Gierlichs, and Ingrid Verbauwhede. An in-depth  
4990 and black-box characterization of the effects of clock glitches on 8-bit  
4991 mcus. In *2011 Workshop on Fault Diagnosis and Tolerance in Cryptog-*  
4992 *raphy*, pages 105–114. IEEE, 2011. [20](#)
- 4993 [BHE<sup>+</sup>19] Jean-Baptiste Bréjon, Karine Heydemann, Emmanuelle Encrenaz,  
4994 Quentin Meunier, and Son-Tuan Vu. Fault attack vulnerability assess-  
4995 ment of binary code. In *Proceedings of the Sixth Workshop on Cryptog-*  
4996 *raphy and Security in Computing Systems*, pages 13–18, 2019. [3](#), [35](#), [62](#),  
4997 [63](#), [II](#)
- 4998 [BHL<sup>+</sup>11] Sébastien Bardin, Philippe Herrmann, Jérôme Leroux, Olivier Ly, Renaud  
4999 Tabary, and Aymeric Vincent. The bincoa framework for binary code  
5000 analysis. In *International Conference on Computer Aided Verification*,  
5001 pages 165–170. Springer, 2011. [70](#)
- 5002 [BKC14] Sébastien Bardin, Nikolai Kosmatov, and François Cheynier. Efficient  
5003 leveraging of symbolic execution to advanced coverage criteria. In *2014*  
5004 *IEEE Seventh International Conference on Software Testing, Verification*  
5005 *and Validation*, pages 173–182. IEEE, 2014. [36](#)
- 5006 [BLLL18] Sebanjila K Bukasa, Ronan Lashermes, Jean-Louis Lanet, and Axel Leqay.  
5007 Let’s shock our iot’s heart: Armv7-m under (fault) attacks. In *Proceed-*  
5008 *ings of the 13th International Conference on Availability, Reliability and*  
5009 *Security*, pages 1–6, 2018. [20](#)
- 5010 [BMPM13] Stephen P Buchner, Florent Miller, Vincent Pouget, and Dale P McMor-  
5011 row. Pulsed-laser testing for single-event effects investigations. *IEEE*  
5012 *Transactions on Nuclear Science*, 60(3):1852–1875, 2013. [19](#)
- 5013 [BR10] Gogul Balakrishnan and Thomas Reps. Wysinwyx: What you see is not  
5014 what you execute. *ACM Transactions on Programming Languages and*  
5015 *Systems (TOPLAS)*, 32(6):1–84, 2010. [34](#)
- 5016 [BRT<sup>+</sup>18] Ryad Benadjila, Mathieu Renard, Philippe Trebuchet, Philippe Thierry,  
5017 Arnauld Michelizza, and Jérémy Lefaire. Wookey: Usb devices strike  
5018 back. *Proceedings of SSTIC*, 2018. [132](#)
- 5019 [BS97] Eli Biham and Adi Shamir. Differential fault analysis of secret key cryp-  
5020 tosystems. In *Advances in Cryptology—CRYPTO’97: 17th Annual Inter-*  
5021 *national Cryptology Conference Santa Barbara, California, USA August*  
5022 *17–21, 1997 Proceedings 17*, pages 513–525. Springer, 1997. [20](#)
- 5023 [BSGD09] Shivam Bhasin, Nidhal Selmane, Sylvain Guilley, and Jean-Luc Dan-  
5024 ger. Security evaluation of different aes implementations against practical  
5025 setup time violation attacks in fpgas. In *2009 IEEE International Work-*  
5026 *shop on Hardware-Oriented Security and Trust*, pages 15–21. IEEE, 2009.  
5027 [20](#)

- 5028 [BT18] Clark Barrett and Cesare Tinelli. Satisfiability modulo theories. In *Hand-*  
5029 *book of model checking*, pages 305–343. Springer, 2018. 30
- 5030 [CBRZ01] Edmund Clarke, Armin Biere, Richard Raimi, and Yunshan Zhu. Bounded  
5031 model checking using satisfiability solving. *Formal methods in system*  
5032 *design*, 19:7–34, 2001. 2, 29, I
- 5033 [CC2] Common criteria for information technology security evaluation. part  
5034 3: Security assurance components, cc:2022, revision 1. [https://www.](https://www.commoncriteriaportal.org/files/ccfiles/CC2022PART5R1.pdf)  
5035 [commoncriteriaportal.org/files/ccfiles/CC2022PART5R1.pdf](https://www.commoncriteriaportal.org/files/ccfiles/CC2022PART5R1.pdf). 12,  
5036 15
- 5037 [CCF<sup>+</sup>05] Patrick Cousot, Radhia Cousot, Jérôme Feret, Laurent Mauborgne, An-  
5038 toine Miné, David Monniaux, and Xavier Rival. The astrée analyzer. In  
5039 *Programming Languages and Systems: 14th European Symposium on Pro-*  
5040 *gramming, ESOP 2005, Held as Part of the Joint European Conferences*  
5041 *on Theory and Practice of Software, ETAPS 2005, Edinburgh, UK, April*  
5042 *4-8, 2005. Proceedings 14*, pages 21–30. Springer, 2005. 2, I
- 5043 [CCG13] Maria Christofi, Boutheina Chetali, and Louis Goubin. Formal verifica-  
5044 tion of an implementation of crt-rsa vigilant’s algorithm. In *PROOFS*  
5045 *workshop: pre-proceedings*, volume 28, 2013. 3, 35, 62, 63, II
- 5046 [CDE<sup>+</sup>08] Cristian Cadar, Daniel Dunbar, Dawson R Engler, et al. Klee: unassisted  
5047 and automatic generation of high-coverage tests for complex systems pro-  
5048 grams. In *OSDI*, volume 8, pages 209–224, 2008. 62
- 5049 [CdFB21] Hervé Chauvet, François de Ferrière, and Thomas Bizet. Software fault  
5050 injection for secswift qualification, 2021. 130
- 5051 [CDFG18] Sébastien Carré, Matthieu Desjardins, Adrien Facon, and Sylvain Guilley.  
5052 Openssl bellcore’s protection helps fault attack. In *2018 21st Euromicro*  
5053 *Conference on Digital System Design (DSD)*, pages 500–507. IEEE, 2018.  
5054 3, 35, 62, 63, II
- 5055 [CDSL20] Domenico Cotroneo, Luigi De Simone, Pietro Liguori, and Roberto  
5056 Natella. Profipy: Programmable software fault injection as-a-service. In  
5057 *2020 50th annual IEEE/IFIP international conference on dependable sys-*  
5058 *tems and networks (DSN)*, pages 364–372. IEEE, 2020. 36
- 5059 [Cer01] Iliano Cervesato. The dolev-yao intruder is the most powerful attacker. In  
5060 *16th Annual Symposium on Logic in Computer Science—LICS*, volume 1,  
5061 pages 1–2. Citeseer, 2001. 3
- 5062 [CGJ<sup>+</sup>03] Edmund Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut  
5063 Veith. Counterexample-guided abstraction refinement for symbolic model  
5064 checking. *Journal of the ACM (JACM)*, 50(5):752–794, 2003. 2, I
- 5065 [CGMN12] Juan Caballero, Gustavo Grieco, Mark Marron, and Antonio Nappa. Un-  
5066 dangle: early detection of dangling pointers in use-after-free and double-  
5067 free vulnerabilities. In *Proceedings of the 2012 International Symposium*  
5068 *on Software Testing and Analysis*, pages 133–143, 2012. 18

- 5069 [CGP<sup>+</sup>08] Cristian Cadar, Vijay Ganesh, Peter M Pawlowski, David L Dill, and  
5070 Dawson R Engler. Exe: Automatically generating inputs of death. *ACM*  
5071 *Transactions on Information and System Security (TISSEC)*, 12(2):1–38,  
5072 2008. [30](#)
- 5073 [CGV<sup>+</sup>22] Brice Colombier, Paul Grandamme, Julien Vernay, Émilie Chanavat, Lil-  
5074 ian Bossuet, Lucie de Laulanié, and Bruno Chassagne. Multi-spot laser  
5075 fault injection setup: New possibilities for fault injection attacks. In *Smart*  
5076 *Card Research and Advanced Applications: 20th International Confer-*  
5077 *ence, CARDIS 2021, Lübeck, Germany, November 11–12, 2021, Revised*  
5078 *Selected Papers*, pages 151–166. Springer, 2022. [19](#)
- 5079 [Cla97] Edmund M Clarke. Model checking. In *Foundations of Software Tech-*  
5080 *nology and Theoretical Computer Science: 17th Conference Kharagpur,*  
5081 *India, December 18–20, 1997 Proceedings 17*, pages 54–56. Springer, 1997.  
5082 [29](#)
- 5083 [CMD<sup>+</sup>19] Brice Colombier, Alexandre Menu, Jean-Max Dutertre, Pierre-Alain  
5084 Moëllic, Jean-Baptiste Rigaud, and Jean-Luc Danger. Laser-induced  
5085 single-bit faults in flash memory: Instructions corruption on a 32-bit  
5086 microcontroller. In *2019 IEEE International Symposium on Hardware*  
5087 *Oriented Security and Trust (HOST)*, pages 1–10. IEEE, 2019. [19](#)
- 5088 [CN13] Domenico Cotroneo and Roberto Natella. Fault injection for software  
5089 certification. *IEEE Security & Privacy*, 11(4):38–45, 2013. [36](#)
- 5090 [Cou21] Patrick Cousot. *Principles of Abstract Interpretation*. MIT Press, 2021.  
5091 [28](#)
- 5092 [CPHR21] Ludovic Claudepierre, Pierre-Yves Péneau, Damien Hardy, and Erven  
5093 Rohou. Traitor: a low-cost evaluation platform for multifault injection.  
5094 In *Proceedings of the 2021 International Symposium on Advanced Security*  
5095 *on Software and Systems*, pages 51–56, 2021. [20](#)
- 5096 [CS13] Cristian Cadar and Koushik Sen. Symbolic execution for software testing:  
5097 three decades later. *Communications of the ACM*, 56(2):82–90, 2013. [2](#),  
5098 [29](#), [30](#), [I](#)
- 5099 [CTB<sup>+</sup>17] Mariano Ceccato, Paolo Tonella, Cataldo Basile, Bart Coppens, Bjorn  
5100 De Sutter, Paolo Falcarin, and Marco Torchiano. How professional hack-  
5101 ers understand protected code while performing attack tasks. In *2017*  
5102 *IEEE/ACM 25th International Conference on Program Comprehension*  
5103 *(ICPC)*, pages 154–164. IEEE, 2017. [64](#)
- 5104 [CTB<sup>+</sup>19] Mariano Ceccato, Paolo Tonella, Cataldo Basile, Paolo Falcarin, Marco  
5105 Torchiano, Bart Coppens, and Bjorn De Sutter. Understanding the be-  
5106 haviour of hackers while performing attack tasks in a professional setting  
5107 and in a public challenge. *Empirical Software Engineering*, 24(1):240–286,  
5108 2019. [64](#)
- 5109 [CW96] Edmund M Clarke and Jeannette M Wing. Formal methods: State of the  
5110 art and future directions. *ACM Computing Surveys (CSUR)*, 28(4):626–  
5111 643, 1996. [27](#)

- 5112 [CWP<sup>+</sup>00] Crispin Cowan, F Wagle, Calton Pu, Steve Beattie, and Jonathan  
5113 Walpole. Buffer overflows: Attacks and defenses for the vulnerability  
5114 of the decade. In *Proceedings DARPA Information Survivability Confer-*  
5115 *ence and Exposition. DISCEX'00*, volume 2, pages 119–129. IEEE, 2000.  
5116 [18](#)
- 5117 [Dan21] Lesly-Ann Daniel. *Symbolic binary-level code analysis for security. Ap-*  
5118 *plication to the detection of microarchitectural timing attacks in crypto-*  
5119 *graphic code*. PhD thesis, Université Côte d’Azur, 2021. [ix](#), [27](#), [28](#)
- 5120 [DB15] Adel Djoudi and Sébastien Bardin. Binsec: Binary code analysis with  
5121 low-level regions. In *International Conference on Tools and Algorithms*  
5122 *for the Construction and Analysis of Systems*, pages 212–217. Springer,  
5123 2015. [70](#)
- 5124 [DBP23] Soline Ducouso, Sébastien Bardin, and Marie-Laure Potet. Adversarial  
5125 reachability for program-level security analysis. *Programming Languages*  
5126 *and Systems LNCS 13990*, page 59, 2023. [4](#), [38](#), [65](#), [III](#)
- 5127 [DBR20] Lesly-Ann Daniel, Sébastien Bardin, and Tamara Rezk. Binsec/rel: Ef-  
5128 ficient relational symbolic execution for constant-time at binary-level. In  
5129 *2020 IEEE Symposium on Security and Privacy (SP)*, pages 1021–1038.  
5130 IEEE, 2020. [31](#), [61](#), [70](#)
- 5131 [DBR21] Lesly-Ann Daniel, Sébastien Bardin, and Tamara Rezk. Hunting the  
5132 haunter-efficient relational symbolic execution for spectre with haunted  
5133 relse. In *NDSS*, 2021. [70](#)
- 5134 [DBT<sup>+</sup>16] Robin David, Sébastien Bardin, Thanh Dinh Ta, Laurent Mounier, Jos-  
5135 selin Feist, Marie-Laure Potet, and Jean-Yves Marion. Binsec/se: A dy-  
5136 namic symbolic execution toolkit for binary-level analysis. In *SANER*,  
5137 2016. [4](#), [70](#)
- 5138 [DDCS<sup>+</sup>14] Jean-Max Dutertre, Stephan De Castro, Alexandre Sarafianos, Noémie  
5139 Boher, Bruno Rouzeyre, Mathieu Lisart, Joel Damiens, Philippe Cande-  
5140 lier, Marie-Lise Flottes, and Giorgio Di Natale. Laser attacks on integrated  
5141 circuits: from cmos to fd-soi. In *2014 9th IEEE International Conference*  
5142 *on Design & Technology of Integrated Systems in Nanoscale Era (DTIS)*,  
5143 pages 1–6. IEEE, 2014. [19](#)
- 5144 [DDRT12] Amine Dehbaoui, Jean-Max Dutertre, Bruno Robisson, and Assia Tria.  
5145 Electromagnetic transient faults injection on a hardware and a software  
5146 implementations of aes. In *2012 Workshop on Fault Diagnosis and Tol-*  
5147 *erance in Cryptography*, pages 7–15. IEEE, 2012. [20](#)
- 5148 [dF21] François de Ferrière. Software countermeasures in the llvm risc-v compiler,  
5149 2021. [130](#)
- 5150 [DHM<sup>+</sup>23] Mathieu Dumont, Kevin Hector, Pierre-Alain Moellic, Jean-Max  
5151 Dutertre, and Simon Pontié. Evaluation of parameter-based attacks  
5152 against embedded neural networks with laser injection. *arXiv preprint*  
5153 *arXiv:2304.12876*, 2023. [131](#)

- 5154 [DLG21] Zhen Yu Ding and Claire Le Goues. An empirical study of oss-fuzz bugs.  
5155 In *2021 IEEE/ACM 18th International Conference on Mining Software*  
5156 *Repositories (MSR)*, pages 131–142. IEEE, 2021. [26](#)
- 5157 [DLRA15] Will Dietz, Peng Li, John Regehr, and Vikram Adve. Understanding  
5158 integer overflow in c/c++. *ACM Transactions on Software Engineering*  
5159 *and Methodology (TOSEM)*, 25(1):1–29, 2015. [18](#)
- 5160 [DMB08] Leonardo De Moura and Nikolaj Bjørner. Z3: An efficient smt solver.  
5161 In *Tools and Algorithms for the Construction and Analysis of Systems:*  
5162 *14th International Conference, TACAS 2008, Held as Part of the Joint*  
5163 *European Conferences on Theory and Practice of Software, ETAPS 2008,*  
5164 *Budapest, Hungary, March 29–April 6, 2008. Proceedings 14*, pages 337–  
5165 340. Springer, 2008. [71](#)
- 5166 [DMM<sup>+</sup>13] Amine Dehbaoui, Amir-Pasha Mirbaha, Nicolas Moro, Jean-Max  
5167 Dutertre, and Assia Tria. Electromagnetic glitch on the aes round counter.  
5168 In *Constructive Side-Channel Analysis and Secure Design: 4th Interna-*  
5169 *tional Workshop, COSADE 2013, Paris, France, March 6–8, 2013, Re-*  
5170 *vised Selected Papers 4*, pages 17–31. Springer, 2013. [20](#)
- 5171 [DOL<sup>+</sup>10] Amine Dehbaoui, Thomas Ordas, Victor Lomné, Philippe Maurine, Li-  
5172 onel Torres, and Michel Robert. Incoherence analysis and its application  
5173 to time domain em analysis of secure circuits. In *2010 Asia-Pacific Inter-*  
5174 *national Symposium on Electromagnetic Compatibility*, pages 1039–1042.  
5175 IEEE, 2010. [17](#)
- 5176 [DPdC<sup>+</sup>15] Louis Dureuil, Marie-Laure Potet, Philippe de Choudens, Cécile Dumas,  
5177 and Jessy Clédière. From code review to fault injection attacks: Filling  
5178 the gap using fault model inference. In *International conference on smart*  
5179 *card research and advanced applications*, pages 107–124. Springer, 2015.  
5180 [34](#), [60](#), [64](#)
- 5181 [DPP<sup>+</sup>16] Louis Dureuil, Guillaume Petiot, Marie-Laure Potet, Thanh-Ha Le, Aude  
5182 Crohen, and Philippe de Choudens. Fissc: A fault injection and simula-  
5183 tion secure collection. In *International Conference on Computer Safety,*  
5184 *Reliability, and Security*, pages 3–11. Springer, 2016. [ix](#), [15](#), [16](#), [20](#), [86](#),  
5185 [87](#), [105](#)
- 5186 [DRPR19] Jean-Max Dutertre, Timothé Riom, Olivier Potin, and Jean-Baptiste  
5187 Rigaud. Experimental analysis of the laser-induced instruction skip fault  
5188 model. In *Secure IT Systems: 24th Nordic Conference, NordSec 2019,*  
5189 *Aalborg, Denmark, November 18–20, 2019, Proceedings 24*, pages 221–  
5190 237. Springer, 2019. [19](#), [20](#)
- 5191 [Dul17] Thomas Dullien. Weird machines, exploitability, and provable unex-  
5192 ploitable. *IEEE Transactions on Emerging Topics in Computing*,  
5193 8(2):391–403, 2017. [129](#), [130](#)
- 5194 [DY83] Danny Dolev and Andrew Yao. On the security of public key protocols.  
5195 *IEEE Transactions on information theory*, 29(2):198–208, 1983. [64](#)
- 5196 [Fac] Facebook. Infer static analyzer. <https://fbinfer.com/>. [2](#), [I](#)

- 
- 5197 [FDBL18] Benjamin Farinier, Robin David, Sébastien Bardin, and Matthieu  
5198 Lemerre. Arrays made simpler: An efficient, scalable and thorough pre-  
5199 processing. In *LPAR*, pages 363–380, 2018. [70](#)
- 5200 [FMB<sup>+</sup>16] Josselin Feist, Laurent Mounier, Sébastien Bardin, Robin David, and  
5201 Marie-Laure Potet. Finding the needle in the heap: combining static  
5202 analysis and dynamic symbolic execution to trigger use-after-free. In  
5203 *Proceedings of the 6th Workshop on Software Security, Protection, and*  
5204 *Reverse Engineering*, pages 1–12, 2016. [31](#)
- 5205 [FR08] Cédric Fournet and Tamara Rezk. Cryptographically sound implemen-  
5206 tations for typed information-flow security. In George C. Necula and  
5207 Philip Wadler, editors, *Proceedings of the 35th ACM SIGPLAN-SIGACT*  
5208 *Symposium on Principles of Programming Languages, POPL 2008, San*  
5209 *Francisco, California, USA, January 7-12, 2008*. ACM, 2008. [65](#)
- 5210 [GDTM21a] Joseph Gravellier, Jean-Max Dutertre, Yannick Teglia, and Philippe Lou-  
5211 bet Moundi. Faultline: Software-based fault injection on memory trans-  
5212 fers. In *2021 IEEE International Symposium on Hardware Oriented Se-*  
5213 *curity and Trust (HOST)*, pages 46–55. IEEE, 2021. [21](#)
- 5214 [GDTM21b] Joseph Gravellier, Jean-Max Dutertre, Yannick Teglia, and Philippe Lou-  
5215 bet Moundi. Sideline: How delay-lines (may) leak secrets from your soc.  
5216 In *Constructive Side-Channel Analysis and Secure Design: 12th Interna-*  
5217 *tional Workshop, COSADE 2021, Lugano, Switzerland, October 25–27,*  
5218 *2021, Proceedings 12*, pages 3–30. Springer, 2021. [21](#)
- 5219 [GFB21] Guillaume Girol, Benjamin Farinier, and Sébastien Bardin. Not all bugs  
5220 are created equal, but robust reachability can tell the difference. In *In-*  
5221 *ternational Conference on Computer Aided Verification*, pages 669–693.  
5222 Springer, 2021. [31](#), [61](#)
- 5223 [GHHR23] Antoine Gicquel, Damien Hardy, Karine Heydemann, and Erven Rohou.  
5224 Samva: Static analysis for multi-fault attack paths determination. In  
5225 *Constructive Side-Channel Analysis and Secure Design: 14th Interna-*  
5226 *tional Workshop, COSADE 2023, Munich, Germany, April 3–4, 2023,*  
5227 *Proceedings*, pages 3–22. Springer, 2023. [62](#), [63](#)
- 5228 [Gir05] Christophe Giraud. Dfa on aes. In *Advanced Encryption Standard–AES:*  
5229 *4th International Conference, AES 2004, Bonn, Germany, May 10-12,*  
5230 *2004, Revised Selected and Invited Papers 4*, pages 27–41. Springer, 2005.  
5231 [20](#)
- 5232 [GKS05] Patrice Godefroid, Nils Klarlund, and Koushik Sen. Dart: Directed au-  
5233 tomated random testing. In *Proceedings of the 2005 ACM SIGPLAN*  
5234 *conference on Programming language design and implementation*, pages  
5235 213–223, 2005. [30](#)
- 5236 [GLM12] Patrice Godefroid, Michael Y Levin, and David Molnar. Sage: whitebox  
5237 fuzzing for security testing. *Communications of the ACM*, 55(3):40–44,  
5238 2012. [2](#), [I](#)

- 5239 [GMA22] Aakash Gangolli, Qusay H Mahmoud, and Akramul Azim. A systematic  
5240 review of fault injection attacks on iot systems. *Electronics*, 11(13):2023,  
5241 2022. [64](#)
- 5242 [God11] Patrice Godefroid. Higher-order test generation. In *Proceedings of the*  
5243 *32nd ACM SIGPLAN conference on Programming language design and*  
5244 *implementation*, pages 258–269, 2011. [31](#)
- 5245 [God20] Patrice Godefroid. Fuzzing: Hack, art, and science. *Communications of*  
5246 *the ACM*, 63(2):70–76, 2020. [28](#)
- 5247 [GST17] Daniel Genkin, Adi Shamir, and Eran Tromer. Acoustic cryptanalysis.  
5248 *Journal of Cryptology*, 30:392–443, 2017. [17](#)
- 5249 [GSV03] Bharat Goyal, Sriranjani Sitaraman, and S Venkatesan. A unified ap-  
5250 proach to detect binding based race condition attacks. In *Int'l Workshop*  
5251 *on Cryptology & Network Security (CANS)*, page 16, 2003. [23](#)
- 5252 [GWJL20] Thomas Given-Wilson, Nisrine Jafri, and Axel Legay. Combined soft-  
5253 ware and hardware fault injection vulnerability detection. *Innovations in*  
5254 *Systems and Software Engineering*, 16(2):101–120, 2020. [3](#), [35](#), [62](#), [63](#), [II](#)
- 5255 [GWJLL17] Thomas Given-Wilson, Nisrine Jafri, Jean-Louis Lanet, and Axel Legay.  
5256 An automated formal process for detecting fault injection vulnerabilities  
5257 in binaries and case study on present. In *2017 IEEE Trustcom/Big-*  
5258 *DataSE/ICISS*, pages 293–300. IEEE, 2017. [3](#), [35](#), [62](#), [63](#), [II](#)
- 5259 [GWL20] Thomas Given-Wilson and Axel Legay. Formalising fault injection and  
5260 countermeasures. In *Proceedings of the 15th International Conference on*  
5261 *Availability, Reliability and Security*, pages 1–11, 2020. [65](#)
- 5262 [HH19] Reiner Hähnle and Marieke Huisman. Deductive software verification:  
5263 from pen-and-paper proofs to industrial tools. *Computing and Software*  
5264 *Science: State of the Art and Perspectives*, pages 345–373, 2019. [27](#)
- 5265 [HS14] Michael Hutter and Jörn-Marc Schmidt. The temperature side channel  
5266 and heating fault attacks. In *Smart Card Research and Advanced Ap-*  
5267 *plications: 12th International Conference, CARDIS 2013, Berlin, Ger-*  
5268 *many, November 27-29, 2013. Revised Selected Papers 12*, pages 219–235.  
5269 Springer, 2014. [20](#)
- 5270 [HSP20] Max Hoffmann, Falk Schellenberg, and Christof Paar. Armory: Fully  
5271 automated and exhaustive fault simulation on arm-m binaries. *IEEE*  
5272 *Transactions on Information Forensics and Security*, 16:1058–1073, 2020.  
5273 [34](#)
- 5274 [HTS<sup>+</sup>17] Siva Kumar Sastry Hari, Timothy Tsai, Mark Stephenson, Stephen W  
5275 Keckler, and Joel Emer. Sassifi: An architecture-level fault injection tool  
5276 for gpu application resilience evaluation. In *2017 IEEE International*  
5277 *Symposium on Performance Analysis of Systems and Software (ISPASS)*,  
5278 pages 249–258. IEEE, 2017. [36](#)

- 5279 [JGH<sup>+</sup>22] Zhiyuan Jiang, Shuitao Gan, Adrian Herrera, Flavio Toffalini, Lucio  
5280 Romerio, Chaojing Tang, Manuel Egele, Chao Zhang, and Mathias Payer.  
5281 Evocatio: Conjuring bug capabilities from a single poc. In *Proceedings of*  
5282 *the 2022 ACM SIGSAC Conference on Computer and Communications*  
5283 *Security*, pages 1599–1613, 2022. [64](#)
- 5284 [KDK<sup>+</sup>14] Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk  
5285 Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu. Flipping bits in  
5286 memory without accessing them: An experimental study of dram distur-  
5287 bance errors. *ACM SIGARCH Computer Architecture News*, 42(3):361–  
5288 372, 2014. [21](#)
- 5289 [KHF<sup>+</sup>20] Paul Kocher, Jann Horn, Anders Fogh, Daniel Genkin, Daniel Gruss,  
5290 Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas  
5291 Prescher, et al. Spectre attacks: Exploiting speculative execution. *Com-*  
5292 *munications of the ACM*, 63(7):93–101, 2020. [22](#)
- 5293 [Kin76] James C King. Symbolic execution and program testing. *Communications*  
5294 *of the ACM*, 19(7):385–394, 1976. [30](#)
- 5295 [KJJ99] Paul Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analy-  
5296 sis. In *Advances in Cryptology—CRYPTO’99: 19th Annual International*  
5297 *Cryptology Conference Santa Barbara, California, USA, August 15–19,*  
5298 *1999 Proceedings 19*, pages 388–397. Springer, 1999. [17](#)
- 5299 [KKBC12] Volodymyr Kuznetsov, Johannes Kinder, Stefan Bucur, and George Can-  
5300 dea. Efficient state merging in symbolic execution. *Acm Sigplan Notices*,  
5301 47(6):193–204, 2012. [33](#)
- 5302 [KKJ<sup>+</sup>21] Kyounggon Kim, Jun Seok Kim, Seonghoon Jeong, Jo-Hee Park, and  
5303 Huy Kang Kim. Cybersecurity for autonomous vehicles: Review of attacks  
5304 and defense. *Computers & Security*, 103:102150, 2021. [11](#)
- 5305 [KKP<sup>+</sup>15] Florent Kirchner, Nikolai Kosmatov, Virgile Prevosto, Julien Signoles,  
5306 and Boris Yakobowski. Frama-c: A software analysis perspective. *Formal*  
5307 *aspects of computing*, 27(3):573–609, 2015. [2](#), [61](#), [62](#), [I](#)
- 5308 [Koc96] Paul C Kocher. Timing attacks on implementations of diffie-hellman, rsa,  
5309 dss, and other systems. In *Advances in Cryptology—CRYPTO’96: 16th*  
5310 *Annual International Cryptology Conference Santa Barbara, California,*  
5311 *USA August 18–22, 1996 Proceedings 16*, pages 104–113. Springer, 1996.  
5312 [17](#)
- 5313 [KP05] Alfred Koelbl and Carl Pixley. Constructing efficient formal models from  
5314 high-level descriptions using symbolic simulation. *International Journal*  
5315 *of Parallel Programming*, 33:645–666, 2005. [33](#)
- 5316 [KSV13] Duško Karaklajić, Jörn-Marc Schmidt, and Ingrid Verbauwhede. Hard-  
5317 ware designer’s guide to fault attacks. *IEEE Transactions on Very Large*  
5318 *Scale Integration (VLSI) Systems*, 21(12):2295–2306, 2013. [19](#)

- 5319 [Lan22] Julien Lancia. Detecting fault injection vulnerabilities in binaries with  
5320 symbolic execution. In *2022 14th International Conference on Electronics,  
5321 Computers and Artificial Intelligence (ECAI)*, pages 1–8. IEEE, 2022. [3](#),  
5322 [35](#), [62](#), [63](#), [II](#)
- 5323 [LBC<sup>+</sup>15] Marc Lacruche, Nicolas Borrel, Clement Champeix, Cyril Roscian,  
5324 Alexandre Sarafianos, Jean-Baptiste Rigaud, Jean-Max Dutertre, and  
5325 Edith Kussener. Laser fault injection into sram cells: Picosecond ver-  
5326 sus nanosecond pulses. In *2015 IEEE 21st International On-Line Testing  
5327 Symposium (IOLTS)*, pages 13–18. IEEE, 2015. [19](#)
- 5328 [LBD<sup>+</sup>18] Johan Laurent, Vincent Berouille, Christophe Deleuze, Florian Pebay-  
5329 Peyroula, and Athanasios Papadimitriou. On the importance of analysing  
5330 microarchitecture for accurate software fault models. In *2018 21st Euromi-  
5331 cro Conference on Digital System Design (DSD)*, pages 561–564. IEEE,  
5332 2018. [60](#), [64](#)
- 5333 [LBDPP19] Johan Laurent, Vincent Berouille, Christophe Deleuze, and Florian Pebay-  
5334 Peyroula. Fault injection on hidden registers in a risc-v rocket processor  
5335 and software countermeasures. In *2019 Design, Automation & Test in  
5336 Europe Conference & Exhibition (DATE)*, pages 252–255. IEEE, 2019. [60](#)
- 5337 [LFBP21] Guilhem Lacombe, David Feliot, Etienne Boespflug, and Marie-Laure  
5338 Potet. Combining static analysis and dynamic symbolic execution in a  
5339 toolchain to detect fault injection vulnerabilities. In *PROOFS WORK-  
5340 SHOP (SECURITY PROOFS FOR EMBEDDED SYSTEMS)*, 2021. [ix](#),  
5341 [3](#), [35](#), [61](#), [62](#), [63](#), [133](#), [135](#), [136](#), [137](#), [II](#)
- 5342 [LH07] Daniel Larsson and Reiner Hähnle. Symbolic fault injection. In *Inter-  
5343 national Verification Workshop (VERIFY)*, volume 259, pages 85–103.  
5344 Citeseer, 2007. [36](#)
- 5345 [LHGD18] Hoang M Le, Vladimir Herdt, Daniel Große, and Rolf Drechsler. Resilience  
5346 evaluation via symbolic fault injection on intermediate code. In *2018  
5347 Design, Automation & Test in Europe Conference & Exhibition (DATE)*,  
5348 pages 845–850. IEEE, 2018. [36](#), [62](#), [105](#)
- 5349 [LM18] Yangdi Lyu and Prabhat Mishra. A survey of side-channel attacks on  
5350 caches and countermeasures. *Journal of Hardware and Systems Security*,  
5351 2:33–50, 2018. [17](#)
- 5352 [LRV<sup>+</sup>22] Quang Loc Le, Azalea Raad, Jules Villard, Josh Berdine, Derek Dreyer,  
5353 and Peter W O’Hearn. Finding real bugs in big programs with incor-  
5354 rectness logic. *Proceedings of the ACM on Programming Languages*,  
5355 6(OOPSLA1):1–27, 2022. [2](#), [I](#)
- 5356 [LWH<sup>+</sup>20] Guanjun Lin, Sheng Wen, Qing-Long Han, Jun Zhang, and Yang Xiang.  
5357 Software vulnerability detection using deep neural networks: a survey.  
5358 *Proceedings of the IEEE*, 108(10):1825–1848, 2020. [29](#)
- 5359 [MBK<sup>+</sup>18] Michaël Marcozzi, Sébastien Bardin, Nikolai Kosmatov, Mike Papadakis,  
5360 Virgile Prevosto, and Loïc Correnson. Time to clean your test objec-  
5361 tives. In *Proceedings of the 40th International Conference on Software  
5362 Engineering*, pages 456–467, 2018. [36](#)

- 5363 [MDH<sup>+</sup>13] Nicolas Moro, Amine Dehbaoui, Karine Heydemann, Bruno Robisson, and  
5364 Emmanuelle Encrenaz. Electromagnetic fault injection: towards a fault  
5365 model on a 32-bit microcontroller. In *2013 Workshop on Fault Diagnosis  
5366 and Tolerance in Cryptography*, pages 77–88. Ieee, 2013. [20](#)
- 5367 [MDP<sup>+</sup>20] Alexandre Menu, Jean-Max Dutertre, Olivier Potin, Jean-Baptiste  
5368 Rigaud, and Jean-Luc Danger. Experimental analysis of the electromag-  
5369 netic instruction skip fault model. In *2020 15th Design & Technology of  
5370 Integrated Systems in Nanoscale Era (DTIS)*, pages 1–7. IEEE, 2020. [20](#)
- 5371 [MK19] Onur Mutlu and Jeremie S Kim. Rowhammer: A retrospective. *IEEE  
5372 Transactions on Computer-Aided Design of Integrated Circuits and Sys-  
5373 tems*, 39(8):1555–1571, 2019. [21](#)
- 5374 [MKP22] Thibault Martin, Nikolai Kosmatov, and Virgile Prevosto. Verifying  
5375 redundant-check based countermeasures: a case study. In *Proceedings of  
5376 the 37th ACM/SIGAPP Symposium on Applied Computing*, pages 1849–  
5377 1852, 2022. [61](#), [62](#), [63](#), [137](#)
- 5378 [MOG<sup>+</sup>20] Kit Murdock, David Oswald, Flavio D Garcia, Jo Van Bulck, Daniel  
5379 Gruss, and Frank Piessens. Plundervolt: Software-based fault injection  
5380 attacks against intel sgx. In *2020 IEEE Symposium on Security and Pri-  
5381 vacy (SP)*, pages 1466–1482. IEEE, 2020. [22](#)
- 5382 [NHH<sup>+</sup>17] Shoei Nashimoto, Naofumi Homma, Yu-ichi Hayashi, Junko Takahashi,  
5383 Hitoshi Fuji, and Takafumi Aoki. Buffer overflow attack with multiple  
5384 fault injection and a proven countermeasure. *Journal of Cryptographic  
5385 Engineering*, 7:35–46, 2017. [20](#)
- 5386 [NP20] Aina Niemetz and Mathias Preiner. Bitwuzla at the SMT-COMP 2020.  
5387 *CoRR*, abs/2006.01621, 2020. [71](#)
- 5388 [OGSM15] Sébastien Ordas, Ludovic Guillaume-Sage, and Philippe Maurine. Em  
5389 injection: Fault model and locality. In *2015 Workshop on Fault Diagnosis  
5390 and Tolerance in Cryptography (FDTC)*, pages 3–13. IEEE, 2015. [20](#)
- 5391 [O’H19] Peter W O’Hearn. Incorrectness logic. *Proceedings of the ACM on Pro-  
5392 gramming Languages*, 4(POPL):1–32, 2019. [64](#)
- 5393 [PCHR20] Pierre-Yves Péneau, Ludovic Claudepierre, Damien Hardy, and Erven  
5394 Rohou. Nop-oriented programming: Should we care? In *2020 IEEE  
5395 European Symposium on Security and Privacy Workshops (EuroS&PW)*,  
5396 pages 694–703. IEEE, 2020. [19](#)
- 5397 [PHB<sup>+</sup>19] Julien Proy, Karine Heydemann, Alexandre Berzati, Fabien Majéric, and  
5398 Albert Cohen. A first isa-level characterization of em pulse effects on  
5399 superscalar microarchitectures: a secure software perspective. In *Proce-  
5400 edings of the 14th International Conference on Availability, Reliability and  
5401 Security*, pages 1–10, 2019. [20](#)
- 5402 [PIK<sup>+</sup>18] Goran Petrovic, Marko Ivankovic, Bob Kurtz, Paul Ammann, and René  
5403 Just. An industrial application of mutation testing: Lessons, challenges,

- 5404 and research directions. In *2018 IEEE International Conference on Soft-*  
5405 *ware Testing, Verification and Validation Workshops (ICSTW)*, pages 47–  
5406 53. IEEE, 2018. [36](#)
- 5407 [PLFP19] Lucas Palazzi, Guanpeng Li, Bo Fang, and Karthik Pattabiraman. A tale  
5408 of two injectors: End-to-end comparison of ir-level and assembly-level  
5409 fault injection. In *2019 IEEE 30th International Symposium on Software*  
5410 *Reliability Engineering (ISSRE)*, pages 151–162. IEEE, 2019. [36](#)
- 5411 [PM10] Mike Papadakis and Nicos Malevris. Automatic mutation test case gen-  
5412 eration via dynamic symbolic execution. In *2010 IEEE 21st International*  
5413 *Symposium on Software Reliability Engineering*, pages 121–130. IEEE,  
5414 2010. [36](#)
- 5415 [PMPD14] Marie-Laure Potet, Laurent Mounier, Maxime Puys, and Louis Dureuil.  
5416 Lazart: A symbolic approach for evaluation the robustness of secured  
5417 codes against control flow injections. In *2014 IEEE Seventh International*  
5418 *Conference on Software Testing, Verification and Validation*, pages 213–  
5419 222. IEEE, 2014. [3](#), [35](#), [62](#), [63](#), [124](#), [II](#)
- 5420 [PNKI08] Karthik Pattabiraman, Nithin Nakka, Zbigniew Kalbarczyk, and Ravis-  
5421 hankar Iyer. Symplified: Symbolic program-level fault injection and error  
5422 detection framework. In *2008 IEEE International Conference on Depend-*  
5423 *able Systems and Networks With FTCS and DCC (DSN)*, pages 472–481.  
5424 IEEE, 2008. [36](#)
- 5425 [PRBL14] Maxime Puys, Lionel Riviere, Julien Bringer, and Thanh-ha Le. High-  
5426 level simulation for multiple fault injection evaluation. In *Data Privacy*  
5427 *Management, Autonomous Spontaneous Security, and Security Assur-*  
5428 *ance*, pages 293–308. Springer, 2014. [128](#)
- 5429 [QS02] Jean-Jacques Quisquater and David Samyde. Eddy current for magnetic  
5430 analysis with active sensor. In *Proceedings of eSMART*, volume 2002,  
5431 2002. [20](#)
- 5432 [RBB<sup>+</sup>19] Frédéric Recoules, Sébastien Bardin, Richard Bonichon, Laurent Mounier,  
5433 and Marie-Laure Potet. Get rid of inline assembly through verification-  
5434 oriented lifting. In *2019 34th IEEE/ACM International Conference on*  
5435 *Automated Software Engineering (ASE)*, pages 577–589. IEEE, 2019. [70](#)
- 5436 [RBB<sup>+</sup>21] Frédéric Recoules, Sébastien Bardin, Richard Bonichon, Matthieu  
5437 Lemerre, Laurent Mounier, and Marie-Laure Potet. Interface compli-  
5438 ance of inline assembly: Automatically check, patch and refine. In  
5439 *2021 IEEE/ACM 43rd International Conference on Software Engineer-*  
5440 *ing (ICSE)*, pages 1236–1247. IEEE, 2021. [70](#)
- 5441 [RBSG22] Jan Richter-Brockmann, Pascal Sasdrich, and Tim Guneyusu. Revisiting  
5442 fault adversary models—hardware faults in theory and practice. *IEEE*  
5443 *Transactions on Computers*, 2022. [64](#)
- 5444 [RG14] Pablo Rauzy and Sylvain Guilley. A formal proof of countermeasures  
5445 against fault injection attacks on crt-rsa. *Journal of Cryptographic Engi-*  
5446 *neering*, 4(3):173–185, 2014. [3](#), [35](#), [62](#), [63](#), [II](#)

- 
- 5447 [RGB<sup>+</sup>16] Kaveh Razavi, Ben Gras, Erik Bosman, Bart Preneel, Cristiano Giuffrida,  
5448 and Herbert Bos. Flip feng shui: Hammering a needle in the software  
5449 stack. In *USENIX Security symposium*, volume 25, pages 1–18, 2016. [21](#)
- 5450 [Ric53] Henry Gordon Rice. Classes of recursively enumerable sets and their  
5451 decision problems. *Transactions of the American Mathematical society*,  
5452 74(2):358–366, 1953. [27](#)
- 5453 [RNR<sup>+</sup>15] Lionel Riviere, Zakaria Najm, Pablo Rauzy, Jean-Luc Danger, Julien  
5454 Bringer, and Laurent Sauvage. High precision fault injections on the  
5455 instruction cache of armv7-m architectures. In *2015 IEEE International  
5456 Symposium on Hardware Oriented Security and Trust (HOST)*, pages 62–  
5457 67. IEEE, 2015. [19](#), [20](#)
- 5458 [RY20] Xavier Rival and Kwangkeun Yi. *Introduction to static analysis: an ab-  
5459 stract interpretation perspective*. Mit Press, 2020. [28](#)
- 5460 [SD15] Mark Seaborn and Thomas Dullien. Exploiting the dram rowhammer bug  
5461 to gain kernel privileges. *Black Hat*, 15:71, 2015. [21](#)
- 5462 [Sen07] Koushik Sen. Concolic testing. In *Proceedings of the twenty-second  
5463 IEEE/ACM international conference on Automated software engineering*,  
5464 pages 571–572, 2007. [29](#)
- 5465 [SMA05] Koushik Sen, Darko Marinov, and Gul Agha. Cute: A concolic unit testing  
5466 engine for c. *ACM SIGSOFT Software Engineering Notes*, 30(5):263–272,  
5467 2005. [30](#)
- 5468 [SNK<sup>+</sup>12] Alexander Schlösser, Dmitry Nedospasov, Juliane Krämer, Susanna Or-  
5469 lic, and Jean-Pierre Seifert. Simple photonic emission analysis of aes:  
5470 photonic side channel analysis for the rest of us. In *Cryptographic Hard-  
5471 ware and Embedded Systems—CHES 2012: 14th International Workshop,  
5472 Leuven, Belgium, September 9-12, 2012. Proceedings 14*, pages 41–57.  
5473 Springer, 2012. [17](#)
- 5474 [TAC<sup>+</sup>22] Simon Tollec, Mihail Asavoae, Damien Couroussé, Karine Heydemann,  
5475 and Mathieu Jan. Exploration of fault effects on formal risc-v microar-  
5476 chitecture models. In *2022 Workshop on Fault Detection and Tolerance  
5477 in Cryptography (FDTC)*, pages 73–83. IEEE, 2022. [60](#), [61](#), [62](#), [63](#), [98](#)
- 5478 [TM17] Niek Timmers and Cristofaro Mune. Escalating privileges in linux us-  
5479 ing voltage fault injection. In *2017 Workshop on Fault Diagnosis and  
5480 Tolerance in Cryptography (FDTC)*, pages 1–8. IEEE, 2017. [20](#)
- 5481 [TS16] Niek Timmers and Albert Spruyt. Bypassing secure boot using fault  
5482 injection. *Black Hat Europe*, 2016, 2016. [20](#)
- 5483 [TSS17] Adrian Tang, Simha Sethumadhavan, and Salvatore J Stolfo. Clkscrew:  
5484 Exposing the perils of security-oblivious energy management. In *USENIX  
5485 Security Symposium*, volume 2, pages 1057–1074, 2017. [22](#)
- 5486 [Van22] Julien Vanegue. Adversarial logic. In *Static Analysis: 29th International  
5487 Symposium, SAS 2022, Auckland, New Zealand, December 5–7, 2022,  
5488 Proceedings*, pages 422–448. Springer, 2022. [64](#)

- 5489 [VBMS<sup>+</sup>20] Jo Van Bulck, Daniel Moghimi, Michael Schwarz, Moritz Lippi, Marina  
5490 Minkin, Daniel Genkin, Yuval Yarom, Berk Sunar, Daniel Gruss, and  
5491 Frank Piessens. Lvi: Hijacking transient execution through microarchi-  
5492 tectural load value injection. In *2020 IEEE Symposium on Security and*  
5493 *Privacy (SP)*, pages 54–72. IEEE, 2020. 23
- 5494 [VKS11] Ingrid Verbauwhede, Dusko Karaklajic, and Jorn-Marc Schmidt. The  
5495 fault attack jungle—a classification model to guide you. In *2011 Workshop*  
5496 *on Fault Diagnosis and Tolerance in Cryptography*, pages 3–8. IEEE, 2011.  
5497 20
- 5498 [VTM<sup>+</sup>17] Aurélien Vasselle, Hugues Thiebauld, Quentin Maouhoub, Adele Moris-  
5499 set, and Sébastien Ermeneux. Laser-induced fault injection on smartphone  
5500 bypassing the secure boot. In *2017 Workshop on Fault Diagnosis and Tol-*  
5501 *erance in Cryptography (FDTC)*, pages 41–48. IEEE, 2017. 19, 20
- 5502 [Woo] <https://github.com/wookey-project>. accessed july 2021. 132
- 5503 [WP05] Jinpeng Wei and Calton Pu. Toctou vulnerabilities in unix-style file  
5504 systems: An anatomical study. In *FAST*, volume 5, pages 12–12, 2005.  
5505 23
- 5506 [WTSS13] Stefan Winter, Michael Tretter, Benjamin Sattler, and Neeraj Suri. simfi:  
5507 From single to simultaneous software fault injections. In *2013 43rd An-*  
5508 *ual IEEE/IFIP International Conference on Dependable Systems and*  
5509 *Networks (DSN)*, pages 1–12. IEEE, 2013. 36
- 5510 [WZ18] Maximilian Wohrer and Uwe Zdun. Smart contracts: security patterns in  
5511 the ethereum ecosystem and solidity. In *2018 International Workshop on*  
5512 *Blockchain Oriented Software Engineering (IWBOSE)*, pages 2–8. IEEE,  
5513 2018. 10
- 5514 [XTDHS09] Tao Xie, Nikolai Tillmann, Jonathan De Halleux, and Wolfram Schulte.  
5515 Fitness-guided path exploration in dynamic symbolic execution. In *2009*  
5516 *IEEE/IFIP International Conference on Dependable Systems & Networks*,  
5517 pages 359–368. IEEE, 2009. 33
- 5518 [YGS<sup>+</sup>16] Bilgiday Yuce, Nahid Farhady Ghalaty, Harika Santapuri, Chinmay Desh-  
5519 pande, Conor Patrick, and Patrick Schaumont. Software fault resistance  
5520 is futile: Effective single-glitch attacks. In *2016 Workshop on Fault Diag-*  
5521 *nosis and Tolerance in Cryptography (FDTC)*, pages 47–58. IEEE, 2016.  
5522 20
- 5523 [ZDT<sup>+</sup>14] Loic Zussa, Amine Dehbaoui, Karim Tobich, Jean-Max Dutertre, Philippe  
5524 Maurine, Ludovic Guillaume-Sage, Jessy Clediere, and Assia Tria. Effi-  
5525 ciency of a glitch detector against electromagnetic fault injection. In *2014*  
5526 *Design, Automation & Test in Europe Conference & Exhibition (DATE)*,  
5527 pages 1–6. IEEE, 2014. 20
- 5528 [ZGWL<sup>+</sup>21] Igor Zavalysyn, Thomas Given-Wilson, Axel Legay, Ramin Sadre, and  
5529 Etienne Riviere. Chaos duck: A tool for automatic iot software fault-  
5530 tolerance analysis. In *2021 40th International Symposium on Reliable*  
5531 *Distributed Systems (SRDS)*, pages 46–55. IEEE, 2021. 36