

# Contents

<b>ACKNOWLEDGEMENTS</b>	<b>xiii</b>
<b>SAMENVATTING</b>	<b>xv</b>
<b>1 GENERAL INTRODUCTION</b>	<b>1</b>
1.1 The Need for Dependability . . . . .	2
1.2 The Problem with Device Drivers . . . . .	4
1.3 Why do Systems Crash? . . . . .	6
1.3.1 Software Complexity . . . . .	6
1.3.2 Design Flaws . . . . .	8
1.4 Improving OS Dependability . . . . .	9
1.4.1 A Modular OS Design . . . . .	10
1.4.2 Fault-tolerance Strategies . . . . .	12
1.4.3 Other Benefits of Modularity . . . . .	13
1.5 Preview of Related Work . . . . .	14
1.6 Focus of this Thesis . . . . .	16
1.7 Outline of this Thesis . . . . .	18
<b>2 ARCHITECTURAL OVERVIEW</b>	<b>19</b>
2.1 The MINIX Operating System . . . . .	19
2.1.1 Historical Perspective . . . . .	19
2.1.2 Multiserver OS Structure . . . . .	21
2.1.3 Interprocess Communication . . . . .	22
2.2 Driver Management . . . . .	24
2.3 Isolating Faulty Drivers . . . . .	26
2.3.1 Isolation Architecture . . . . .	26
2.3.2 Hardware Considerations . . . . .	30
2.4 Recovering Failed Drivers . . . . .	32
2.4.1 Defect Detection and Repair . . . . .	32
2.4.2 Assumptions and Limitations . . . . .	34
2.5 Fault and Failure Model . . . . .	35

<b>3</b>	<b>FAULT ISOLATION</b>	<b>37</b>
3.1	Isolation Principles . . . . .	37
3.1.1	The Principle of Least Authority . . . . .	37
3.1.2	Classification of Privileged Operations . . . . .	38
3.1.3	General Rules for Isolation . . . . .	41
3.2	User-level Driver Framework . . . . .	42
3.2.1	Moving Drivers to User Level . . . . .	42
3.2.2	Supporting User-level Drivers . . . . .	43
3.3	Isolation Techniques . . . . .	44
3.3.1	Restricting CPU Usage . . . . .	44
3.3.2	Restricting Memory Access . . . . .	45
3.3.3	Restricting Device I/O . . . . .	50
3.3.4	Restricting IPC . . . . .	51
3.4	Case Study: Living in Isolation . . . . .	54
<b>4</b>	<b>FAILURE RESILIENCE</b>	<b>57</b>
4.1	Defect Detection Techniques . . . . .	57
4.1.1	Unexpected Process Exits . . . . .	58
4.1.2	Periodic Status Monitoring . . . . .	58
4.1.3	Explicit Update Requests . . . . .	59
4.2	On-the-fly Repair . . . . .	60
4.2.1	Recovery Scripts . . . . .	60
4.2.2	Restarting Failed Components . . . . .	61
4.2.3	State Management . . . . .	63
4.3	Effectiveness of Recovery . . . . .	65
4.3.1	Recovering Device Drivers . . . . .	66
4.3.2	Recovering System Servers . . . . .	70
4.4	Case Study: Monitoring Driver Correctness . . . . .	70
4.5	Case Study: Automating Server Recovery . . . . .	73
<b>5</b>	<b>EXPERIMENTAL EVALUATION</b>	<b>75</b>
5.1	Software-implemented Fault Injection . . . . .	75
5.1.1	SWIFI Test Methodology . . . . .	75
5.1.2	Network-device Driver Results . . . . .	79
5.1.3	Block-device Driver Results . . . . .	85
5.1.4	Character-device Driver Results . . . . .	87
5.2	Performance Measurements . . . . .	89
5.2.1	Costs of Fault Isolation . . . . .	89
5.2.2	Costs of Failure Resilience . . . . .	94
5.3	Source-code Analysis . . . . .	96
5.3.1	Evolution of MINIX 3 . . . . .	96
5.3.2	Evolution of Linux 2.6 . . . . .	99

<b>6</b>	<b>RELATED WORK</b>	<b>101</b>
6.1	In-kernel Sandboxing . . . . .	101
6.1.1	Hardware-enforced Protection . . . . .	102
6.1.2	Software-based Isolation . . . . .	104
6.2	Virtualization Techniques . . . . .	107
6.2.1	Full Virtualization . . . . .	107
6.2.2	Paravirtualization . . . . .	109
6.3	Formal Methods . . . . .	111
6.3.1	Language-based Protection . . . . .	112
6.3.2	Driver Synthesis . . . . .	115
6.4	User-level Frameworks . . . . .	117
6.4.1	Process Encapsulation . . . . .	117
6.4.2	Split-driver Architectures . . . . .	120
6.5	Comparison . . . . .	123
<b>7</b>	<b>SUMMARY AND CONCLUSION</b>	<b>125</b>
7.1	Summary of this Thesis . . . . .	125
7.1.1	Problem Statement and Approach . . . . .	125
7.1.2	Fault-tolerance Techniques . . . . .	128
7.2	Lessons Learned . . . . .	130
7.2.1	Dependability Challenges . . . . .	131
7.2.2	Performance Perspective . . . . .	132
7.2.3	Engineering Effort . . . . .	133
7.3	Epilogue . . . . .	135
7.3.1	Contribution of this Thesis . . . . .	135
7.3.2	Application of this Research . . . . .	136
7.3.3	Directions for Future Research . . . . .	137
7.4	Availability of MINIX 3 . . . . .	139
	<b>REFERENCES</b>	<b>141</b>
	<b>ABBREVIATIONS</b>	<b>161</b>
	<b>PUBLICATIONS</b>	<b>163</b>
	<b>BIOGRAPHY</b>	<b>165</b>