Simula SpringerBriefs on Computing

Volume 1

More information about this series at http://www.springer.com/series/13548

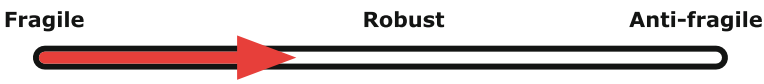
Fragile
Robust
Anti-fragile

Kjell Jørgen Hole

Anti-fragile ICT Systems

Springer Open

Kjell Jørgen Hole
Department of Informatics
University of Bergen
Bergen
Norway

Simula SpringerBriefs on Computing
ISBN 978-3-319-30068-9 ISBN 978-3-319-30070-2 (eBook)
DOI 10.1007/978-3-319-30070-2

Library of Congress Control Number: 2016931422

Mathematics Subject Classification (2010): 68-02

Printed on acid-free paper

Springer International Publishing AG Switzerland is part of Springer Science+Business Media (www.springer.com)

To my wife, Winnie Busiku

Foreword

Dear reader,

Our aim with the series *Simula SpringerBriefs on Computing* is to provide compact introductions to selected fields of computing. Entering a new field of research can be quite demanding for graduate students, postdocs, and experienced researchers alike: the process often involves reading hundreds of papers, and the methods, results and notation styles used often vary considerably, which makes for a time-consuming and potentially frustrating experience. The briefs in this series are meant to ease the process by introducing and explaining important concepts and theories in a relatively narrow field, and by posing critical questions on the fundamentals of that field. A typical brief in this series should be around 100 pages and should be well suited as material for a research seminar in a well-defined and limited area of computing.

We have decided to publish all items in this series under the SpringerOpen framework, as this will allow authors to use the series to publish an initial version of their manuscript that could subsequently evolve into a full-scale book on a broader theme. Since the briefs are freely available online, the authors will not receive any direct income from the sales; however, remuneration is provided for every completed manuscript. Briefs are written on the basis of an invitation from a member of the editorial board. Suggestions for possible topics are most welcome and can be sent to aslak@simula.no.

January 2016

Prof. Aslak Tveito
CEO

Dr. Martin Peters
Executive Editor Mathematics
Springer Heidelberg, Germany

Preface

As information and communications technology (ICT) becomes increasingly important to modern societies, there is a growing need to understand how to design and operate very large ICT systems. How should a huge system be designed and operated to support both high availability and rapid change? Will some of the system's stakeholders be exposed to events with intolerable impact? Is the system fragilizing a service of importance to millions of users? These are questions that need answers.

According to conventional wisdom, the opposite of a fragile system is a robust system. While stressors or perturbations can easily damage fragile systems, robust systems can withstand a great deal of pressure. This is why we write *handle with care* on a box with fragile contents and nothing on a box with robust contents. In 2012, essayist and scholar Nassim N. Taleb published his landmark book *Antifragility: Things That Gain from Disorder*, pointing out that the opposite of a fragile system is really a system that needs stressors to thrive. We would write *please mishandle* on a box with anti-fragile contents. Unlike robust systems, anti-fragile systems learn from events with negative impact how to adjust themselves and become stronger in a changing world. An example of an anti-fragile system is the human immune system, with its ability to adapt and self-repair. While Taleb's book discusses many natural and man-made systems that are anti-fragile, it says nothing about how to design and operate anti-fragile ICT systems.

Anti-fragile ICT Systems

This book you hold in your hands or are reading on a computing device models large distributed ICT systems as complex adaptive systems to determine fundamental properties that make systems anti-fragile to different classes of events with a

negative impact.[1] For example, a system can be anti-fragile to downtime or the spreading of malicious software or malware. Because there are many types of ICT systems and because each type can be anti-fragile to many classes of events, we cannot study all possible anti-fragile ICT systems. Instead, this book examines different aspects of anti-fragile systems carefully selected to show that the concept of anti-fragility offers a novel and useful approach to the design and operation of complex adaptive ICT systems.

The book first discusses rare events with a large negative impact and argues that it is, at best, very hard to predict all such events in complex adaptive ICT systems. It explains why it is necessary to limit the impact of these events to gain robustness and why learning from the remaining events with a small impact is necessary to achieve anti-fragility. Since loss of trust is an inherent and general threat to any ICT system, the book also models why it is vital for an organization operating an anti-fragile ICT system to build and maintain a strong trust relationship with its customer base. Next, the book discusses four design principles, namely, *modularity*, *weak links*, *redundancy*, and *diversity*, and one operational principle, the *fail fast* principle. While each principle by itself is well known and does not provide any new fundamental insight, collectively the five principles outline a novel way to design and operate anti-fragile ICT systems.

We apply the five principles in studies of how anti-fragile systems can (i) achieve high availability, (ii) prevent malware epidemics, and (iii) detect anomalies. Analyses of real ICT systems such as Netflix's media streaming solution, Norway's telecommunication (telecom) infrastructure, electronic government platforms, banking systems, and Numenta's anomaly detection software show that cloud computing is central to achieving all three goals. The book therefore concentrates on the design and operation of anti-fragile systems running on cloud computing platforms.

There are good reasons why the goals (i)–(iii) were selected. We study systems that are anti-fragile to downtime because prolonged outages constitute a serious problem in a world where users are increasingly dependent on ICT systems. Malware of many different types represents another serious problem affecting the security and well-being of all Internet users. Since "classical" signature-based malware detection techniques are inadequate, we study novel solutions to cope with the large negative impact of malware. Finally, to react quickly to local failures before they have time to spread, it is necessary to detect system anomalies early. This is a difficult challenge, since complex ICT systems have many interconnected entities. Consequently, we study a powerful and general learning algorithm to detect anomalies.

At the time of this writing, there are no general methods or theories on how to develop or operate anti-fragile ICT systems. The book studies select philosophical and practical aspects of anti-fragile ICT systems to gain an initial understanding

[1]The book should be printed in color or read on a device with a color screen because some of the figures are hard to understand when reproduced in black and white.

of them. The main message is that we should stop building fragile ICT systems of national or international importance and start building anti-fragile ICT systems. The book's contents are deeply influenced by Taleb's work on anti-fragile systems that thrive in a world dominated by large-impact, hard-to-predict, and rare events, Daniel E. Geer Jr.'s keynote speech at the Source 2008 Conference,[2] and Jeff Hawkins' still evolving theory on how the brain learns. The individual chapters are based on my own published research, basic results in complexity and network science, presentations by Neil Hunt[3] and Adrian Cockcroft[4] on Netflix's web-scale solution, and talks by Subutai Ahmad[5] and Scott Purdy[6] on Numenta's technology for anomaly detection.

Who Should Read This Book

While this introductory book is, first and foremost, written for undergraduate students in computer science, the first half should be understandable to any technically educated individual interested in the design, development, and operation of large ICT systems. The first half introduces the concept of anti-fragility, describes the design and operational principles, and outlines how the principles can be applied to achieve anti-fragility to downtime. The book's second half is more technical and assumes that the reader has an elementary understanding of graphs. It describes how to achieve anti-fragility against malware spreading and how to detect anomalies. The whole book should be of interest to new graduate students looking for a research topic.

The book contains few abbreviations and formal definitions, background knowledge is introduced as needed, and studies of real systems help clarify concepts and insights. Each chapter is short and to the point, enabling reading in one or two sittings. Key information is repeated to make chapters easier to understand and the definitions of central abbreviations are repeated in each chapter they are used. An effort was made to reference easy-to-understand books, papers, reports, and webpages for readers wanting more background information. While the book argues that anti-fragile ICT solutions in the cloud should have a microservice architecture, it is not a textbook on cloud computing and microservices. More information on cloud computing platforms and how to implement microservices can be found in the References and on the Web.

[2]See geer.tinho.net/geer.sourceboston.txt.

[3]See youtube.com/watch?v=jCanhyFDopQ.

[4]See youtube.com/watch?v=dekV3Oq7pH8.

[5]See youtube.com/watch?v=nVCKjZWYavM.

[6]See youtube.com/watch?v=I5lSEHvngaI.

Table 1 This book is partly based on articles published by the Institute of Electrical and Electronics Engineers (IEEE)

Title	Authors	IEEE citation
Toward Risk Assessment of Large-Impact and Rare Events	K.J. Hole and L.-H. Netland	*Security & Privacy*, vol. 8 no. 3, 2010, pp. 21–27
Building and Maintaining Trust in Internet Voting	L.H. Nestås and K.J. Hole	*Computer*, vol. 45, no. 5, 2012, pp. 74–80
Management of Hidden Risks	K.J. Hole	*Computer*, vol. 46, no. 1, 2013, pp. 65–70
Diversity Reduces the Impact of Malware	K.J. Hole	*Security & Privacy*, vol. 13, no. 3, 2015, pp. 48–54
Towards Anti-fragility: A Malware-Halting Technique	K.J. Hole	*Security & Privacy*, vol. 13, no. 4, 2015, pp. 40–46
Building Trust in E-Government Services	K.J. Hole	*Computer*, vol. 49, no. 1, 2016, pp. 66–74

Acknowledgments

I am grateful to my colleagues Olav Lysne, Øyvind Ytrehus, and Håvard Raddum, as well as my students Tetiana Yarygina, Christian W. Otterstad, and Alexandre Vivmond for illuminating discussions and comments on early versions of the manuscript. A special thanks to the external expert reviewers Chief Information Security Officer Daniel E. Geer Jr. at In-Q-Tel and Vice President of Research Subutai Ahmad at Numenta. Thanks are also due to the internal expert reviewers at Simula Research Laboratory, Head of Department Ernst Gunnar Gran, Research Scientist Ahmed Elmokashi, and Senior Research Scientist Leon Moonen. The expert reviewers pointed out embarrassing mistakes, suggested much needed changes, and asked important questions leading to significant improvements of the text. Of course, I take full responsibility for all remaining mistakes and ambiguities in the book.

Some chapters are based on my own work published in the IEEE magazines *Security & Privacy* and *Computer*. I am grateful to the IEEE for allowing me to reuse material from the articles listed in Table 1. Thanks also to the articles' anonymous reviewers for the many useful comments and good suggestions that improved the presentation of the material. Finally, thanks to Lars-Helge Netland and Lars Hopland Nestås, my coauthors of the two first articles in Table 1.

Bergen, Norway
December 2015

Kjell Jørgen Hole

Contents

About the Author

Kjell Jørgen Hole lives in Norway. He holds a full-time position as a Professor in the Department of Informatics, University of Bergen (UiB) and a part-time position as the Head of the Security Department at the Simula Research Laboratory in Oslo. At the time of this writing, he is part of a joint effort between UiB and Simula to build a new cybersecurity research group.

While Kjell completed his Ph.D. in Coding Theory at UiB, he did most of his thesis work at the University of California, San Diego, where he worked at the Center for Magnetic Recording Research (CMRR). At CMRR, he was fortunate enough to join the "Wolf pack," led by the late Prof. Jack K. Wolf. Professor Wolf was an outstanding teacher, a dedicated thesis advisor, and, above all, a great human being.

Kjell was a postdoctoral researcher at IBM Almaden Research Center in Silicon Valley, where he conducted research on convolutional codes and artificial neural networks. Later, he worked on trellis coded modulation at the Norwegian University of Science and Technology. During this period, Kjell mainly published his research in *IEEE Transactions on Information Theory* and *IEEE Transactions on Communications*.

Kjell eventually switched fields from coding theory to cybersecurity because he wanted to do more applied research. His first research group in security became "infamous" in Norway for a few years because it shattered local myths about information technology systems' high degree of security and privacy. The group published articles in the IEEE magazines *Security & Privacy* and *Computer*. Kjell has continued to publish in these magazines.

Kjell enjoys teaching and working with students. Through the years he has supervised many master's and Ph.D. students. At UiB, he developed and taught courses and seminars in introductory programming, coding theory, information security, anti-fragile systems, and communication standards such as Wi-Fi and Bluetooth. He has also given many talks to Norwegian industry and written feature articles about security in Norwegian newspapers.

Together with two of his Ph.D. students, Kjell founded a security consultancy. While the consultancy was never a big financial success, the effort increased his appreciation of how difficult and time-consuming it is to build a successful company. Today, he is a board member of mCASH, a company operating a mobile payment solution in Norway.