Intrusion Tolerance for Networked Systems through Two-Level Feedback Control

IEEE DSN 2024, Brisbane, Australia
International Conference on Dependable Systems and Networks

Kim Hammar and Rolf Stadler

kimham@kth.se and *stadler@kth.se*
KTH Royal Institute of Technology

June 27, 2024

Use Case: Intrusion Tolerance

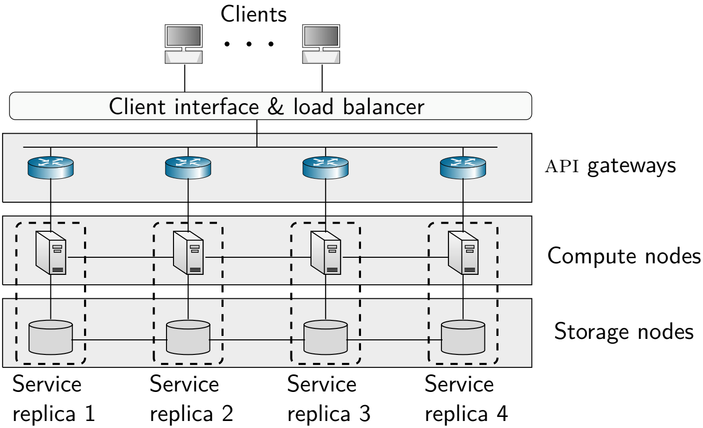

- A **replicated system** offers a service to a client population.
- The system should provide service without disruption.

Use Case: Intrusion Tolerance

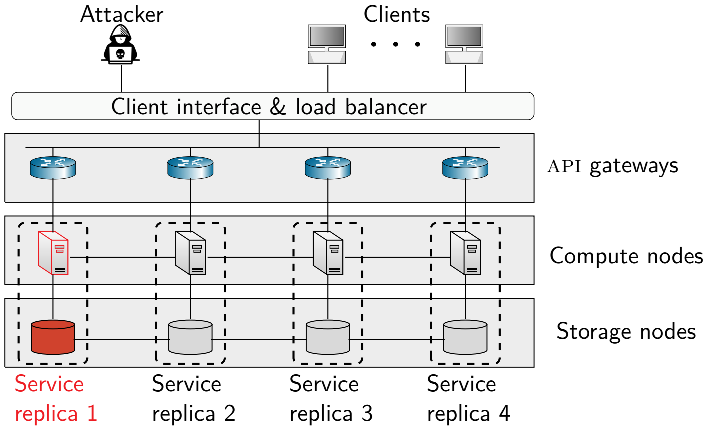

- An attacker seeks to intrude on the system and disrupt service.
- The system should **tolerate intrusions.**

Intrusion Tolerance (Simplified)

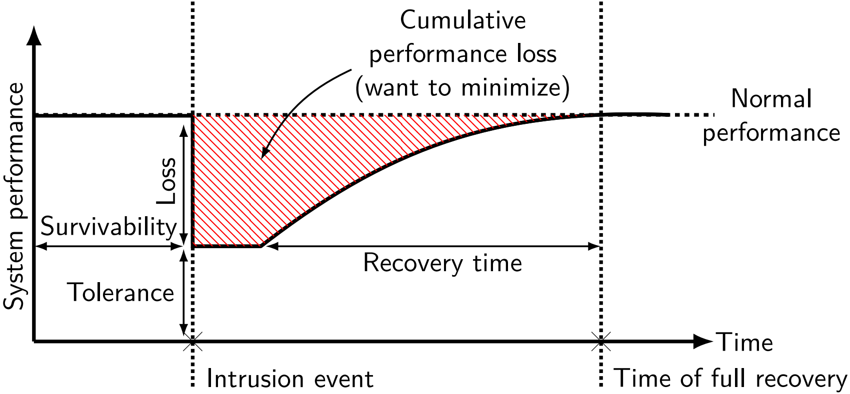

Increasing Demand for Intrusion-Tolerant Systems

- As our **reliance on online services grows**, there is an increasing demand for intrusion-tolerant systems.
- Example applications:

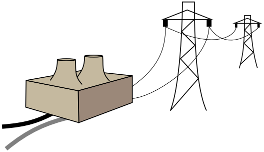

Power grids
e.g., SCADA systems[1].

Safety-critical IT systems
e.g., banking systems, e-commerce applications[2], healthcare systems, etc.

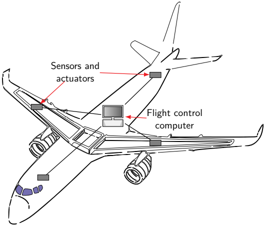

Real-time control systems
e.g., flight control computer[3].

[1] Amy Babay et al. "Network-Attack-Resilient Intrusion-Tolerant SCADA for the Power Grid". In: *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 2018, pp. 255–266. DOI: 10.1109/DSN.2018.00036.

[2] Jukka Soikkeli et al. "Redundancy Planning for Cost Efficient Resilience to Cyber Attacks". In: *IEEE Transactions on Dependable and Secure Computing* 20.2 (2023), pp. 1154–1168. DOI: 10.1109/TDSC.2022.3151462.

[3] J.H. Wensley et al. "SIFT: Design and analysis of a fault-tolerant computer for aircraft control". In: *Proceedings of the IEEE* 66.10 (1978), pp. 1240–1255. DOI: 10.1109/PROC.1978.11114.

Theoretical Foundations of Intrusion Tolerance

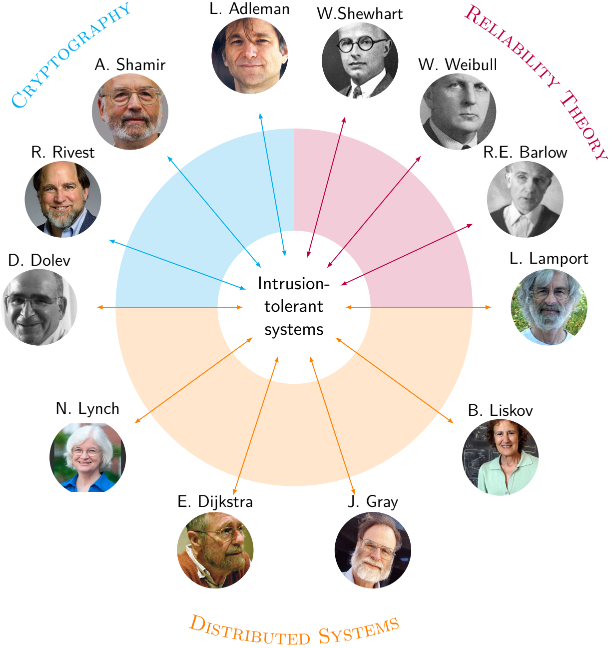

Our Contribution

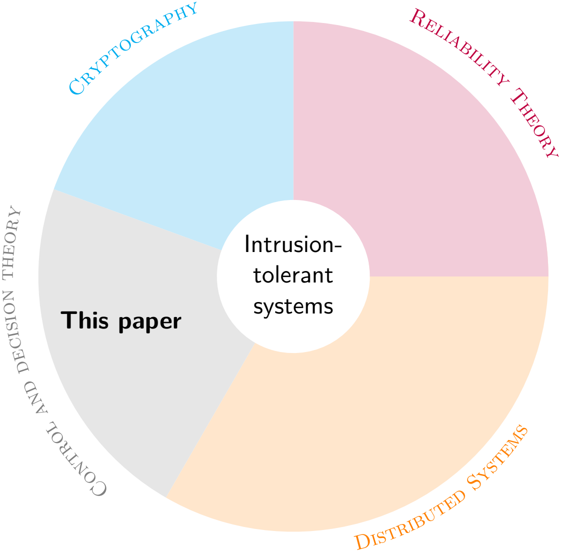

Building Blocks of An Intrusion-Tolerant System

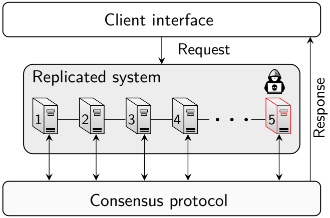

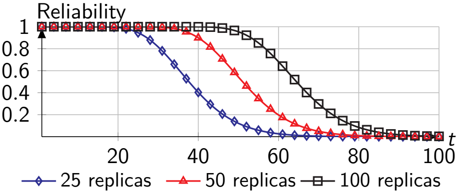

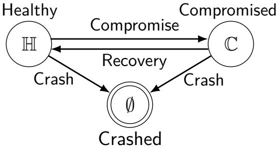

1. Intrusion-tolerant consensus protocol

A quorum needs to reach agreement to tolerate f compromised replicas.

2. Replication strategy

Cost-reliability trade-off.

3. Recovery strategy

Compromises will occur as $t \to \infty$.

Prior Work on Intrusion-Tolerant Systems

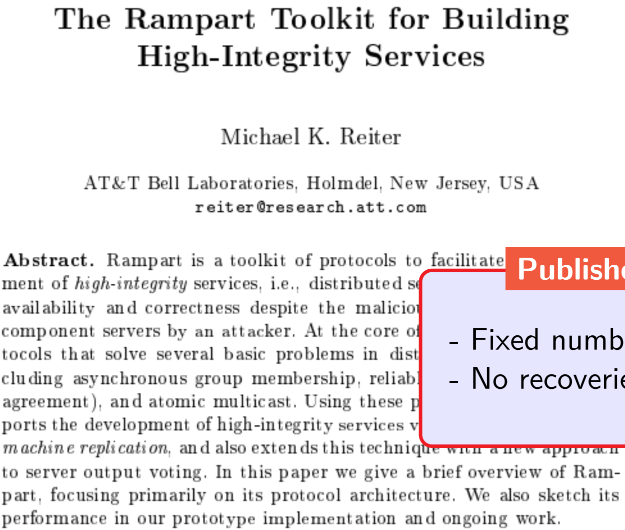

The Rampart Toolkit for Building High-Integrity Services

Michael K. Reiter

AT&T Bell Laboratories, Holmdel, New Jersey, USA
reiter@research.att.com

Abstract. Rampart is a toolkit of protocols to facilitate
ment of *high-integrity* services, i.e., distributed se
availability and correctness despite the maliciou
component servers by an attacker. At the core of
tocols that solve several basic problems in dist
cluding asynchronous group membership, reliab
agreement), and atomic multicast. Using these p
ports the development of high-integrity services v
machine replication, and also extends this technique
to server output voting. In this paper we give a brief overview of Rampart, focusing primarily on its protocol architecture. We also sketch its performance in our prototype implementation and ongoing work.

Published 1995

- Fixed number of replicas
- No recoveries

Prior Work on Intrusion-Tolerant Systems

The Rampart Toolkit for Building High-Integrity Services

Michael K. Reiter

AT&T Bell Laboratories, Holmdel, New Jersey, USA
reiter@research.att.com

Abstract. Rampart is a toolkit of protocols to facilitate the development of *high-integrity* services, i.e., distributed services that retain their availability and correctness despite the malicious penetration of some component servers by an attacker. At the core of Rampart are new protocols that solve several basic problems in distributed computing, including asynchronous group membership, reliable multicast (Byzantine agreement), and atomic multicast. Using these protocols, Rampart supports the development of high-integrity services via the technique of *state machine replication*, and also extends this technique with a new approach to server output voting. In this paper we give a brief overview of Rampart, focusing primarily on its protocol architecture. We also sketch its performance in our prototype implementation and ongoing work.

The SecureRing Protocols for Securing Group Communication*

Kim Potter Kihlstrom, L. E. Moser, P. M. Melliar-Smith
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
kimk@alpha.ece.ucsb.edu, moser@ece.ucsb.edu, pmms@ece.ucsb.edu

Abstract

The SecureRing group communication protocols prov... reliable ordered message delivery and group members... services despite Byzantine faults such as might be caused... modifications to the programs of a group member follow... illicit access to, or capture of, a group member.

Published 1998

- Fixed number of replicas
- No recoveries

Prior Work on Intrusion-Tolerant Systems

The Rampart Toolkit for Building High-Integrity Services

Michael K. Reiter

AT&T Bell Laboratories, Holmdel, New Jersey, USA
reiter@research.att.com

The SecureRing Protocols for Securing Group Communication*

Kim Potter Kihlstrom, L. E. Moser, P. M. Melliar-Smith
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
kimk@alpha.ece.ucsb.edu, moser@ece.ucsb.edu, pmms@ece.ucsb.edu

Abstract

The SecureRing group communication protocols provide reliable ordered message delivery and group membership services despite Byzantine faults such as might be caused by modifications to the programs of a group member following illicit access to, or capture of, a group member. The

processors within an asynchronous distributed syst pose a consistent total order on messages, and r consistent group memberships.

The approach adopted by SecureRing to protect Byzantine faults is to optimize the performance mal (fault-free) operation and to pay a performance

Practical Byzantine Fault Tolerance and Proactive Recovery

MIGUEL CASTRO
Microsoft Research
and
BARBARA LISKOV
MIT Laboratory for Computer Science

Our growing reliance on online service
tems that provide correct service with
malicious attacks are a major cause of
ior, that is, Byzantine faults. This artic
used to build highly available systems
to implement real services: it performs
Internet, it incorporates mechanisms
replicas proactively. The recovery mech

Published 2002

- Fixed number of replicas
- **Periodic** recoveries

Prior Work on Intrusion-Tolerant Systems

The Rampart Toolkit for Building High-Integrity Services

Michael K. Reiter
AT&T Bell Laboratories, Holmdel, New Jersey, USA
reiter@research.att.com

The SecureRing Protocols for Securing Group Communication*

Kim Potter Kihlstrom, L. E. Moser, P. M. Melliar-Smith
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
kimk@alpha.ece.ucsb.edu, moser@ece.ucsb.edu, pmms@ece.ucsb.edu

Abstract
The SecureRing group communication protocols provide reliable ordered message delivery and group membership services despite Byzantine faults such as might be caused by modifications to the programs of a group member following illicit access to, or capture of, a group member. The

processors within an asynchronous distributed syst pose a consistent total order on messages, and r consistent group memberships.
The approach adopted by SecureRing to protect Byzantine faults is to optimize the performance mal (fault-free) operation and to pay a performance

Practical Byzantine Fault Tolerance and Proactive Recovery

MIGUEL CASTRO
Microsoft Research
and
BARBARA LISKOV
MIT Laboratory for Computer Science

Our growing reliance on online services accessible on the Internet demands highly available systems that provide correct service without interruptions. Software bugs, operator mistakes, and malicious attacks are a major cause of service interruptions and they can cause arbitrary behavior, that is, Byzantine faults. This article describes a new replication algorithm, BFT, that can be used to build highly available systems that tolerate Byzantine faults. BFT can be used in practice to implement real services: it performs well, it is safe in asynchronous environments such as the Internet, it incorporates mechanisms to defend against Byzantine-faulty clients, and it recovers replicas proactively. The recovery mechanism allows the algorithm to tolerate any number of faults

A Qualitative Analysis of the Intrusion-Tolerance Capabilities of the MAFTIA Architecture

Robert Stroud, Ian Welch[1], John Warne, Peter Ryan,
School of Computing Science, University of Newcastle upon Tyne, UK
{R.J.Stroud, J.P.Warne, Peter.Ryan}@ncl.ac.uk
Ian.Welch@mcs.vu

Abstract
MAFTIA was a three-year European resear project that explored the use of fault-toleran techniques to build intrusion-tolerant systems. T MAFTIA architecture embodies a number of key desi

Published 2004

- Fixed number of replicas
- Periodic recoveries

Prior Work on Intrusion-Tolerant Systems

The Rampart Toolkit for Building High-Integrity Services

Michael K. Reiter

AT&T Bell Laboratories, Holmdel, New Jersey, USA
reiter@research.att.com

The SecureRing Protocols for Securing Group Communication*

Kim Potter Kihlstrom, L. E. Moser, P. M. Melliar-Smith
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
kimk@alpha.ece.ucsb.edu, moser@ece.ucsb.edu, pmms@ece.ucsb.edu

Abstract

The SecureRing group communication protocols provide reliable ordered message delivery and group membership services despite Byzantine faults such as might be caused by modifications to the programs of a group member following illicit access to, or capture of, a group member. The

processors within an asynchronous distributed syst pose a consistent total order on messages, and consistent group memberships.

The approach adopted by SecureRing to protect Byzantine faults is to optimize the performance mal (fault-free) operation and to pay a performance

Practical Byzantine Fault Tolerance and Proactive Recovery

MIGUEL CASTRO
Microsoft Research
and
BARBARA LISKOV
MIT Laboratory for Computer Science

A Qualitative Analysis of the Intrusion-Tolerance Capabilities of the MAFTIA Architecture

Robert Stroud, Ian Welch[1], John Warne, Peter Ryan,
School of Computing Science, University of Newcastle upon Tyne, UK
{R.J.Stroud, J.P.Warne, Peter.Ryan}@ncl.ac.uk
Ian.Welch@mcs.vuw.ac.nz

Abstract

MAFTIA was a three-year European research project that explored the use of fault-tolerance techniques to build intrusion-tolerant systems. The MAFTIA architecture embodies a number of key design

presence of malicious faults, i.e., deliberate attack the security of the system by both insiders outsiders. Such faults are perpetrated by attackers make unauthorised attempts to access, modify destroy information in a system, and/or to rende system unreliable or unusable. Attacks are facili by vulnerabilities and a successful attack results i

An architecture for adaptive intrusion-tolerant applications

Partha Pal[1,*] and Paul Rubel[1], Michael Atighetchi[1], Franklin Webber[1], William H. Sanders[2], Mouna Seri[2], HariGovind Ramasamy[3], James Lyons[2], Tod Courtney[3], Adnan Agbaria[2], Michel Cukier[3], Jeanna Gossett[4], Idit Keidar[5]

[1] *BBN Technologies, Cambridge, Massachusetts. {ppal, prubel, matighet, fwebber}@bbn.com*
[2] *University of Illinois at Urbana-Champaign. {whs, seri, ramasamy, jlyons, tod, adnan}@crhc.uiuc.edu*
[3] *University of Maryland at College Park, Maryland. mcukier@eng.umd.edu* [4] *The Boeing Company. jeanna.m.gossett@* ... *ectrical Engineering,*

Published 2006

- **Adaptive replication** based on heuristics
- Periodic recoveries

Prior Work on Intrusion-Tolerant Systems

An architecture for adaptive intrusion-tolerant applications

Partha Pal[1,*] and Paul Rubel[1], Michael Atighetchi[1], Franklin Webber[1], William H. Sanders[2], Mouna Seri[2], HariGovind Ramasamy[3], James Lyons[2], Tod Courtney[3], Adnan Agbaria[2], Michel Cukier[3], Jeanna Gossett[4], Idit Keidar[5]

[1] *BBN Technologies, Cambridge, Massachusetts. {ppal, prubel, matighet, fwebber}@bbn.com*
[2] *University of Illinois at Urbana-Champaign. {whs, seri, ramasamy, jlyons, tod, adnan}@crhc.uiuc.edu*
[3] *University of Maryland at College Park, Maryland. mcukier@eng.umd.edu* [4] *The Boeing Company. jeanna.m.gossett@MW.Boeing.com* [5] *Department of Electrical Engineering, Technion – Israel Institute of Technology. idish@technion.ac.il*

The Rampart Toolkit for Building High-Integrity Services

Michael K. Reiter

AT&T Bell Laboratories, Holmdel, New Jersey, USA
reiter@research.att.com

The SecureRing Protocols for Securing Group Communica

Kim Potter Kihlstrom, L. E. Moser, P. M. Melliar-Smith
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
kimk@alpha.ece.ucsb.edu, moser@ece.ucsb.edu, pmms@ece.ucsb.edu

Abstract

The SecureRing group communication protocols provide reliable ordered message delivery and group membership services despite Byzantine faults such as might be caused by modifications to the programs of a group member following illicit access to, or capture of, a group member. The

processors within an asynchronous distributed syst pose a consistent total order on messages, and consistent group memberships.

The approach adopted by SecureRing to protec Byzantine faults is to optimize the performance mal (fault-free) operation and to pay a performance

Practical Byzantine Fault Tolerance and Proactive Recovery

MIGUEL CASTRO
Microsoft Research
and
BARBARA LISKOV
MIT Laboratory for Computer Science

A Qualitative Analysis of the Intrusion-Tolerance Capabilities of the MAFTIA Architecture

Robert Stroud, Ian Welch[1], John Warne, Peter Ryan,
School of Computing Science, University of Newcastle upon Tyne, UK
{R.J.Stroud, J.P.Warne, Peter.Ryan}@ncl.ac.uk
Ian.Welch@mcs.vuw.ac.nz

Abstract

MAFTIA was a three-year European research project that explored the use of fault-tolerance techniques to build intrusion-tolerant systems. The MAFTIA architecture embodies a number of key design

presence of malicious faults, i.e., deliberate attack the security of the system by both insiders outsiders. Such faults are perpetrated by attackers make unauthorised attempts to access, modify destroy information in a system, and/or to rende system unreliable or unusable. Attacks are facili by vulnerabilities and a successful attack results i

Worm-IT – A wormhole-based intrusion-tolerant group communication system

Miguel Correia [a,*], Nuno Ferreira Neves [a], Lau Cheuk Lung [b], Paulo Veríssimo [a]

[a] *Faculdade de Ciências da Universidade de Lisboa, Departamento de Informática, Campo Grande, Bloco C6, Piso 3, 1749-016 Lisboa, Portugal*
[b] *Programa de Pós-Graduação em Informática Aplicada, Pontifícia Univ… Conceição, 1155, 80.215-901, Brazil*

Received 26 October 2005; r…

Published 2006

- Fixed number of replicas
- Periodic recoveries

Prior Work on Intrusion-Tolerant Systems

The Rampart Toolkit for Building High-Integrity Services

Michael K. Reiter

AT&T Bell Laboratories, Holmdel, New Jersey, USA
reiter@research.att.com

An architecture for adaptive intrusion-tolerant applications

Partha Pal[1,*] and Paul Rubel[1], Michael Atighetchi[1], Franklin Webber[1], William H. Sanders[2], Mouna Seri[2], HariGovind Ramasamy[3], James Lyons[2], Tod Courtney[3], Adnan Agbaria[2], Michel Cukier[3], Jeanna Gossett[4], Idit Keidar[5]

[1] BBN Technologies, Cambridge, Massachusetts. {ppal, prubel, matighet, fwebber}@bbn.com [2] University of Illinois at Urbana-Champaign. {whs, seri, ramasamy, jlyons, tod, adnan}@crhc.uiuc.edu [3] University of Maryland at College Park, Maryland. mcukier@eng.umd.edu [4] The Boeing Company. jeanna.m.gossett@MW.Boeing.com [5] Department of Electrical Engineering, Technion – Israel Institute of Technology. idish@technion.ac.il

The SecureRing Protocols for Securing Group Communica

Kim Potter Kihlstrom, L. E. Moser, P. M. Melliar-Smith
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
kimk@alpha.ece.ucsb.edu, moser@ece.ucsb.edu, pmms@ece.ucsb.edu

Abstract

The SecureRing group communication protocols provide reliable ordered message delivery and group membership services despite Byzantine faults such as might be caused by modifications to the programs of a group member following illicit access to, or capture of, a group member. The

processors within an asynchronous distributed syst pose a consistent total order on messages, and consistent group memberships.

The approach adopted by SecureRing to protec Byzantine faults is to optimize the performance mal (fault-free) operation and to pay a performance

Practical Byzantine Fault Tolerance and Proactive Recovery

MIGUEL CASTRO
Microsoft Research
and
BARBARA LISKOV
MIT Laboratory for Computer Science

A Qualitative Analysis of the Intrusion-Tolerance Capabilities of the MAFTIA Architecture

Robert Stroud, Ian Welch[1], John Warne, Peter Ryan,
School of Computing Science, University of Newcastle upon Tyne, UK
{R.J.Stroud, J.P.Warne, Peter.Ryan}@ncl.ac.uk
Ian.Welch@mcs.vuw.ac.nz

Worm-IT – A wormhole-based intrusion-tolerant group communication system

Miguel Correia [a,*], Nuno Ferreira Neves [a], Lau Cheuk Lung [b], Paulo Veríssimo [a]

[a] Faculdade de Ciências da Universidade de Lisboa, Departamento de Informática, Campo Grande, Bloco C6, Piso 3, 1749-016 Lisboa, Portugal
[b] Programa de Pós-Graduação em Informática Aplicada, Pontifícia Universidade Católica do Paraná, Rua Imaculada Conceição, 1155, 80.215-901, Brazil

Received 26 October 2005; received in revised form 28 March 2006; accepted 30 March 2006

Resilient Intrusion Tolerance through Proactive and Reactive Recovery*

Paulo Sousa Alysson Neves Bessani Miguel Correia
Nuno Ferreira Neves Paulo Verissimo
LASIGE, Faculdade de Ciências da Universidade de Lisboa – Portugal
{pjsousa, bessani, mpc, nuno, pjv}@di.fc.ul.pt

Published 2007

- Fixed number of replicas
- **Supports both periodic and reactive recoveries**
- Does not provide reactive recovery strategies

Prior Work on Intrusion-Tolerant Systems

The Rampart Toolkit for Building High-Integrity Services

Michael K. Reiter
AT&T Bell Laboratories, Holmdel, New Jersey, USA
reiter@research.att.com

An architecture for adaptive intrusion-tolerant applications

Partha Pal[1,*] and Paul Rubel[1], Michael Atighetchi[1], Franklin Webber[1], William H. Sanders[2], Mouna Seri[2], HariGovind Ramasamy[3], James Lyons[2], Tod Courtney[3], Adnan Agbaria[2], Michel Cukier[3], Jeanna Gossett[4], Idit Keidar[5]

[1] BBN Technologies, Cambridge, Massachusetts. {ppal, prubel, matighet, fwebber}@bbn.com
[2] University of Illinois at Urbana-Champaign. {whs, seri, ramasamy, jlyons, tod, adnan}@crhc.uiuc.edu
[3] University of Maryland at College Park, Maryland. mcukier@eng.umd.edu [4] The Boeing Company. jeanna.m.gossett@MW.Boeing.com [5] Department of Electrical Engineering, Technion – Israel Institute of Technology. idish@technion.ac.il

The SecureRing Protocols for Securing Group Communica

Kim Potter Kihlstrom, L. E. Moser, P. M. Melliar-Smi
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
kimk@alpha.ece.ucsb.edu, moser@ece.ucsb.edu, pmms@ece.ucsb.

Abstract
The SecureRing group communication protocols provide reliable ordered message delivery and group membership services despite Byzantine faults such as might be caused by modifications to the programs of a group member following illicit access to, or capture of, a group member. The

processors within an asynchronous distributed syst pose a consistent total order on messages, and consistent group memberships.
The approach adopted by SecureRing to prote Byzantine faults is to optimize the performance mal (fault-free) operation and to pay a performanc

Resilient Intrusion Tolerance through Proactive and Reactive Recovery*

Paulo Sousa Alysson Neves Bessani Miguel Correia
Nuno Ferreira Neves Paulo Verissimo
LASIGE, Faculdade de Ciências da Universidade de Lisboa – Portugal
{pjsousa, bessani, mpc, nuno, pjv}@di.fc.ul.pt

Practical Byzantine Fault Tolerance and Proactive Recovery

MIGUEL CASTRO
Microsoft Research
and
BARBARA LISKOV
MIT Laboratory for Computer Science

A Qualitative Analysis of the Intrusion-Tolerance Capabilities of the MAFTIA Architecture

Robert Stroud, Ian Welch[1], John Warne, Peter Ryan,
School of Computing Science, University of Newcastle upon Tyne, UK
{R.J.Stroud, J.P.Warne, Peter.Ryan}@ncl.ac.uk
Ian.Welch@mcs.vuw.ac.nz

Worm-IT – A wormhole-based intrusion-tolerant group communication system

Miguel Correia [a,*], Nuno Ferreira Neves [a], Lau Cheuk Lung [b], Paulo Veríssimo [a]

[a] Faculdade de Ciências da Universidade de Lisboa, Departamento de Informática, Campo Grande, Bloco C6, Piso 3, 1749-016 Lisboa, Portugal
[b] Programa de Pós-Graduação em Informática Aplicada, Pontifícia Universidade Católica do Paraná, Rua Imaculada Conceição, 1155, 80.215-901, Brazil

Received 26 October 2005; received in revised form 28 March 2006; accepted 30 March 2006

State Transfer for Hypervisor-Based Proactive Recovery of Heterogeneous Replicated Services

Tobias Distler Rüdiger Kapitza
Friedrich-Alexander University
Erlangen-Nuremberg, Germany
{distler,rrkapitza}@cs.fau.de

Hans P. Reiser
LASIGE
Universidade de Lisboa, Portu
hans@di.fc.ul.pt

Published 2011

- Fixed number of replicas
- Periodic recoveries

Prior Work on Intrusion-Tolerant Systems

The Rampart Toolkit for Building High-Integrity Services

Michael K. Reiter
AT&T Bell Laboratories, Holmdel, New Jersey, USA
reiter@research.att.com

An architecture for adaptive intrusion-tolerant applications

Partha Pal[1,*] and Paul Rubel[1], Michael Atighetchi[1],
William H. Sanders[2], Mouna Seri[2], HariGovind Rama
Tod Courtney[3], Adnan Agbaria[2], Michel Cukier[3], Jea

[1] BBN Technologies, Cambridge, Massachusetts. {ppal, prubel,
[2] University of Illinois at Urbana-Champaign. {whs, seri, rama
adnan}@crhc.uiuc.edu
[3] University of Maryland at College Park, Maryland. mcukier@
Company. jeanna.m.gossett@MW.Boeing.com [5] Department of Electrical Engineering,
Technion – Israel Institute of Technology. idish@technion.ac.il

State Transfer for Hypervisor-Based Proactive Recovery of Heterogeneous Replicated Services

Tobias Distler Rüdiger Kapitza
Friedrich-Alexander University
Erlangen-Nuremberg, Germany
{distler,rrkapitz}@cs.fau.de

Hans P. Reiser
LASIGE
Universidade de Lisboa, Portugal
hans@di.fc.ul.pt

The SecureRing Protocols for Securing Group Communica

Kim Potter Kihlstrom, L. E. Moser, P. M. Melliar-Smi
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
kimk@alpha.ece.ucsb.edu, moser@ece.ucsb.edu, pmms@ece.ucsb.

Abstract

The SecureRing group communication protocols provide reliable ordered message delivery and group membership services despite Byzantine faults such as might be caused by modifications to the programs of a group member following illicit access to, or capture of, a group member. The

processors within an asynchronous distributed syst
pose a consistent total order on messages, and
consistent group memberships.
The approach adopted by SecureRing to protec
Byzantine faults is to optimize the performance
mal (fault-free) operation and to pay a performance

Resilient Intrusion Tolerance through Proactive and Reactive Recovery*

Paulo Sousa Alysson Neves Bessani Miguel Correia
Nuno Ferreira Neves Paulo Verissimo
LASIGE, Faculdade de Ciências da Universidade de Lisboa – Portugal
{pjsousa, bessani, mpc, nuno, pjv}@di.fc.ul.pt

Practical Byzantine Fault Tolerance and Proactive Recovery

MIGUEL CASTRO
Microsoft Research
and
BARBARA LISKOV
MIT Laboratory for Computer Science

A Qualitative Analysis of the Intrusion-Tolerance Capabilities of the MAFTIA Architecture

Robert Stroud, Ian Welch[1], John Warne, Peter Ryan,
School of Computing Science, University of Newcastle upon Tyne, UK
{R.J.Stroud, J.P.Warne, Peter.Ryan}@ncl.ac.uk
Ian.Welch@mcs.vuw.ac.nz

Worm-IT – A wormhole-based intrusion-tolerant group communication system

Miguel Correia [a,*], Nuno Ferreira Neves [a], Lau Cheuk Lung [b], Paulo Veríssimo [a]

[a] *Faculdade de Ciências da Universidade de Lisboa, Departamento de Informática, Campo Grande, Bloco C6, Piso 3, 1749-016 Lisboa, Portugal*
[b] *Programa de Pós-Graduação em Informática Aplicada, Pontifícia Universidade Católica do Paraná, Rua Imaculada Conceição, 1155, 80.215-901, Brazil*

Received 26 October 2005; received in revised form 28 March 2006; accepted 30 March 2006

Network-Attack-Resilient Intrusion-Tolerant SCADA for the Power Grid

Amy Babay*, Thomas Tantillo*, Trevor Aron, Marco Platania, and Yair Amir
Johns Hopkins University — {babay, tantillo, taron1, yairamir}@cs.jhu.edu
AT&T Labs — {platania}@research.att.com
Spread Concepts LLC — {yairamir}@spreadconcepts.com

Published 2018

- Fixed number of replicas
- Periodic recoveries

Prior Work on Intrusion-Tolerant Systems

The Rampart Toolkit for Building High-Integrity Services

Michael K. Reiter
AT&T Bell Laboratories, Holmdel, New Jersey, USA
reiter@research.att.com

An architecture for adaptive intrusion-tolerant applications

Partha Pal[1,*] and Paul Rubel[1], Michael Atighetchi[1], William H. Sanders[2], Mouna Seri[2], HariGovind Rama… Tod Courtney[3], Adnan Agbaria[2], Michel Cukier[3], Jea…

[1] BBN Technologies, Cambridge, Massachusetts. {ppal, prubel, … [2] University of Illinois at Urbana-Champaign. {whs, seri, rama… adnan}@crhc.uiuc.edu [3] University of Maryland at College Park, Maryland. mcukier@… Company. jeanna.m.gossett@MW.Boeing.com [5] Department of Electr… Technion – Israel Institute of Technology. idish@technion.ac.il

State Transfer for Hypervisor-Based Proactive Recovery of Heterogeneous Replicated Services

Tobias Distler Rüdiger Kapitza
Friedrich-Alexander University
Erlangen-Nuremberg, Germany
{distler,rrkapitz}@cs.fau.de

Hans P. Reiser
LASIGE
Universidade de Lisboa, Portugal
hans@di.fc.ul.pt

The SecureRing Protocols for Securing Group Communica…

Kim Potter Kihlstrom, L. E. Moser, P. M. Melliar-Smi…
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
kimk@alpha.ece.ucsb.edu, moser@ece.ucsb.edu, pmms@ece.ucsb…

Resilient Intrusion Tolerance through Proactive a…

Paulo Sousa Alysson Neves Bessani Mig…
Nuno Ferreira Neves Paulo Veríssim…
LASIGE, Faculdade de Ciências da Universidade de Lisboa – Portugal
{pjsousa, bessani, mpc, nuno, pjv}@di.fc.ul.pt

Network-Attack-Resilient Intrusion-Tolerant SCADA for the Power Grid

Amy Babay*, Thomas Tantillo*, Trevor Aron, Marco Platania, and Yair Amir
Johns Hopkins University — {babay, tantillo, taron1, yairamir}@cs.jhu.edu
AT&T Labs — {platania}@research.att.com
Spread Concepts LLC — {yairamir}@spreadconcepts.com

Abstract

The SecureRing group communication protocols provide reliable ordered message delivery and group membership services despite Byzantine faults such as might be caused by modifications to the programs of a group member following illicit access to, or capture of, a group member. The

processors within an asynchronous distributed syst… pose a consistent total order on messages, and … consistent group memberships.

The approach adopted by SecureRing to protect… Byzantine faults is to optimize the performance … mal (fault-free) operation a…

Practical Byzantine Fault Tolerance and Proactive Recovery

MIGUEL CASTRO
Microsoft Research
and
BARBARA LISKOV
MIT Laboratory for Computer Science

Skynet: a Cyber-Aware Intrusion Tolerant Overseer

Tadeu Freitas, João Soares, Manuel E. Correia, Rolando Martins
Department of Computer Science, Faculty of Science, University of Porto
Email:{tadeufreitas, joao.soares, mdcorrei, rmartins}@fc.up.pt

A Qualitative Analysis of the Intrusion-Tolerance Capabilities of the MAFTIA Architecture

Robert Stroud, Ian Welch[1], John Warne, Peter Ryan,
School of Computing Science, University of Newcastle upon Tyne, UK
{R.J.Stroud, J.P.Warne, Peter.Ryan}@ncl.ac.uk
Ian.Welch@mcs.vuw.ac.nz

Published 2023

- Fixed number of replicas
- Periodic recoveries

Worm-IT – A wormhole-based intrusion-tolerant group communication system

Miguel Correia [a,*], Nuno Ferreira Neves [a], Lau Cheuk Lung [b], Paulo Veríssimo [a]

[a] Faculdade de Ciências da Universidade de Lisboa, Departamento de Informática, Campo Grande, Bloco C6, Piso 3, 1749-016 Lisboa, Portugal
[b] Programa de Pós-Graduação em Informática Aplicada, Pontifícia Universidade Católica do Paraná, Rua Imaculada Conceição, 1155, 80.215-901, Brazil

Received 26 October 2005; received in revised form 28 March 2006; accepted 30 March 2006

Prior Work on Intrusion-Tolerant Systems

The Rampart Toolkit for Building High-Integrity Services

Michael K. Reiter

AT&T Bell Laboratories, Holmdel, New Jersey, USA
reiter@research.att.com

An architecture for adaptive intrusion-tolerant applications

Partha Pal and Paul Rubel, Michael Atighetchi, William H. Sanders, Mouna Seri, HariGovind Rama, Tod Courtney, Adnan Agbaria, Michel Cukier, Jea

State Transfer for Hypervisor-Based Proactive Recovery of Heterogeneous Replicated Services

Tobias Distler, Rüdiger Kapitza, Hans P. Reiser

Friedrich-Alexander University Erlangen-Nuremberg, Germany {distler,rrkapitz}@cs.fau.de
LASIGE Universidade de Lisboa, Portugal hans@di.fc.ul.pt

The SecureRing Protocols for Securing Group Communica

Kim Potter Kihlstrom, L. E. Moser, P. M. Melliar-Smi
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
kimk@alpha.ece.ucsb.edu, moser@ece.ucsb.edu, pmms@ece.ucsb.

Resilient Intrusion Tolerance through Proactive a

Paulo Sousa, Alysson Neves Bessani, Mig, Nuno Ferreira Neves, Paulo Verissim
LASIGE, Faculdade de Ciências da Universidade de Lisboa – Portugal
{pjsousa, bessani, mpc, nuno, pjv}@di.fc.ul.pt

Network-Attack-Resilient Intrusion-Tolerant SCADA for the Power Grid

Amy Babay, Thomas Tantillo, Trevor Aron, Marco Platania, and Yair Amir

Can we do better by leveraging decision theory and optimal control?

A Qualitative Analysis of the Intrusion-Tolerance Capabilities of the MAFTIA Architecture

Robert Stroud, Ian Welch[1], John Warne, Peter Ryan,
School of Computing Science, University of Newcastle upon Tyne, UK
{R.J.Stroud, J.P.Warne, Peter.Ryan}@ncl.ac.uk
Ian.Welch@mcs.vuw.ac.nz

Worm-IT – A wormhole-based intrusion-tolerant group communication system

Miguel Correia[a,*], Nuno Ferreira Neves[a], Lau Cheuk Lung[b], Paulo Verissimo[a]

Received 26 October 2005; received in revised form 28 March 2006; accepted 30 March 2006

Published 2023

- Fixed number of replicas
- Periodic recoveries

The TOLERANCE Architecture

Tw**o**-**le**vel **r**ecovery **an**d replication **c**ontrol with f**e**edback.

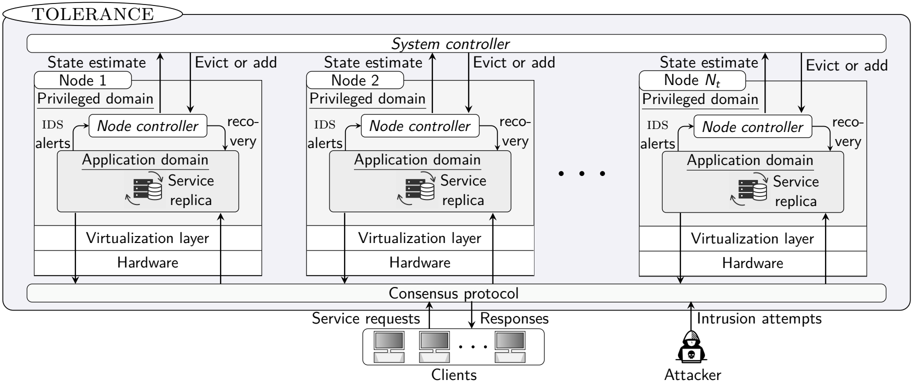

Definition 1 (Correct service)

The system provides **correct service** if the healthy replicas satisfy the following properties:

Each request is eventually executed. (Liveness)

Each executed request was sent by a client. (Validity)

Each replica executes the same request sequence. (Safety)

Proposition 1 (Correctness of TOLERANCE)

A system that implements the TOLERANCE architecture ***provides correct service*** *if*

- *Network links are authenticated.*
- *At most f nodes are compromised or crashed simultaneously.*
- $N_t \geq 2f + 1$.
- *The system is partially synchronous.*

Intrusion Tolerance as a Two-Level Control Problem

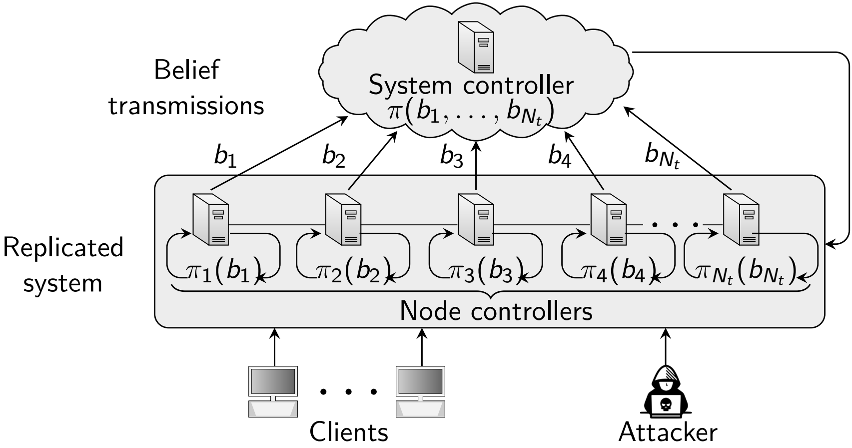

- The local level models intrusion recovery.
- The global level models replication control.

Assumption 1

The probability that the system controller fails is negligible.

Assumption 2

Compromise and crash events are statistically independent across nodes.

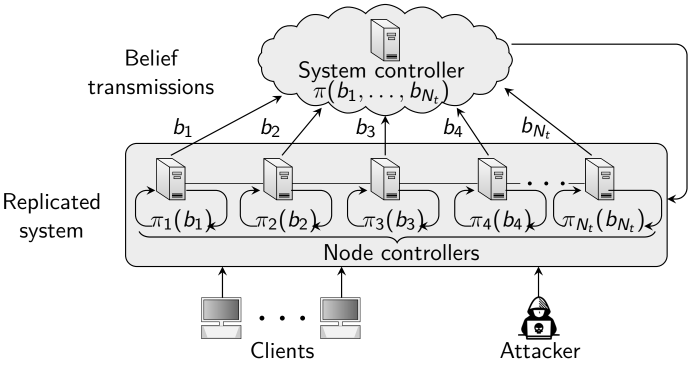

The Local Control Problem

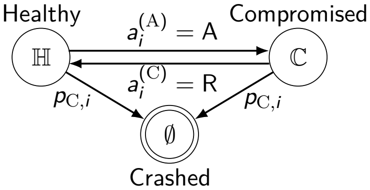

- **Partially observed Markov decision process Γ_i.**
- Controller actions: (R)ecover and (W)ait. $a_{i,t} \in \{\mathsf{R}, \mathsf{W}\}$.

- Node states: $\mathcal{S}_{\mathrm{N}} = \{(\mathbb{H})\text{ealthy}, (\mathbb{C})\text{ompromised}, \emptyset\}$. $s_{i,t} \in \mathcal{S}_{\mathrm{N}}$.
- State transition function: $f(s_{i,t} \mid s_{i,t}, a_{i,t})$.
- $p_{\mathrm{C},i}$: crash probability, $p_{\mathrm{A},i}$: intrusion probability.

- **Observation** $o_{i,t} \sim z_i(\cdot|s_{i,t})$: e.g., IDS alerts at time t.

Node Controller Strategy

- The controller computes the **belief**

$$b_{i,t}(s) \triangleq \mathbb{P}[S_{i,t} = \mathbb{C} | \mathbf{h}_t].$$

$$\mathbf{h}_t \triangleq (b_{i,1}, a_{i,1}, o_{i,2}, a_{i,2}, o_{i,3}, \ldots, a_{i,t-1}, o_{i,t}).$$

- Controller strategy:

$$\pi : [0, 1] \rightarrow \{\mathrm{W}, \mathrm{R}\}.$$

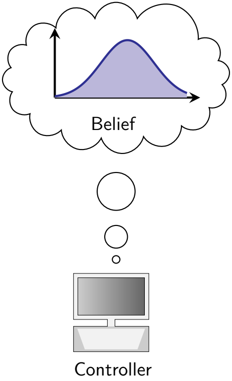

Node Controller Objective

- **Cost:** $J_i \triangleq \eta T_i^{(\mathrm{R})} + F_i^{(\mathrm{R})}$.
  - $T_i^{(\mathrm{R})}$ is the average *time-to-recovery*.
  - $F_i^{(\mathrm{R})}$ is the *recovery frequency*.
  - $\eta > 1$ is a scaling factor.

- Bounded-time-to-recovery constraint: The time between two recoveries can be at most Δ_{R}.

Failure (crash or compromise) probability.

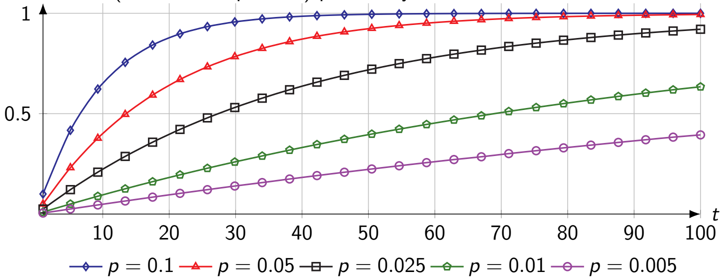

p is the failure probability per time-step.

Threshold Structure of the Optimal Control Strategy

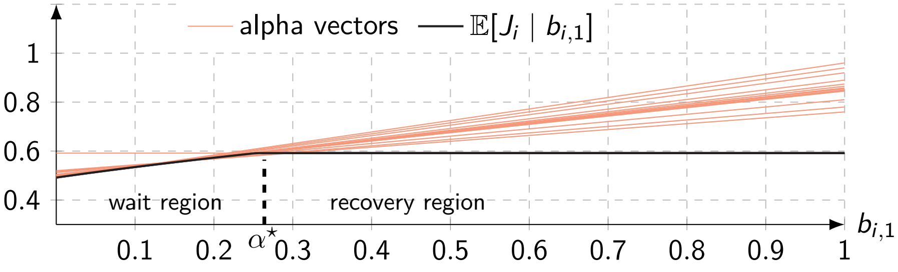

The controller's optimal cost function.

Theorem 2

There exists an optimal control strategy that satisfies

$$\pi^\star_{i,t}(b_{i,t}) = \mathsf{R} \iff b_{i,t} \geq \alpha^\star_{i,t} \qquad \forall t,$$

where $\alpha^\star_{i,t} \in [0,1]$ *is a threshold.*

Efficient Computation of Optimal Recovery Strategies

Algorithm 1: Threshold Optimization

Input: Objective function J_i, parametric optimizer po.

Output: An approximate optimal control strategy $\hat{\pi}_{i,\theta}$.

Algorithm

$\Theta \leftarrow [0, 1]$.

For each $\theta \in \Theta$, define $\pi_{i,\theta}(b_{i,t})$ as

$$\pi_{i,\theta}(b_{i,t}) \triangleq \begin{cases} \mathsf{R} & \text{if } b_{i,t} \geq \theta \\ \mathsf{W} & \text{otherwise.} \end{cases}$$

$J_\theta \leftarrow \mathbb{E}_{\pi_{i,\theta}}[J_i]$.

$\hat{\pi}_{i,\theta} \leftarrow \mathrm{po}(\Theta, J_\theta)$.

return $\hat{\pi}_{i,\theta}$.

- Examples of parameteric optimization algorithmns: CEM, BO, CMA-ES, DE, SPSA, etc.

Efficient Computation of Optimal Recovery Strategies

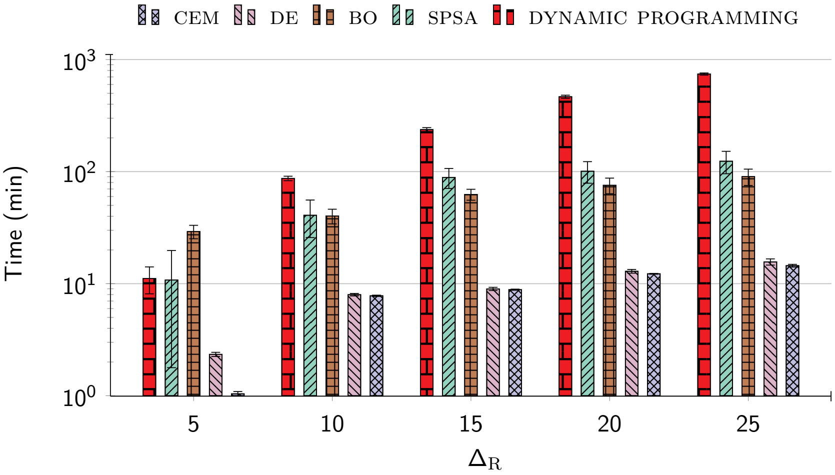

Mean compute time to obtain an optimal recovery strategy for different values of the bounded-time-to-recovery constraint Δ_{R}.

The Benefit of Optimal Recovery Control

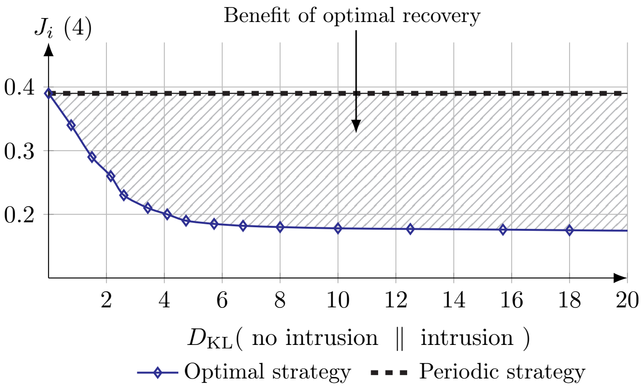

Key insight

Optimal recovery control can significantly reduce operational cost given that an **intrusion detection model is available**.

Intrusion Tolerance as a Two-Level Control Problem

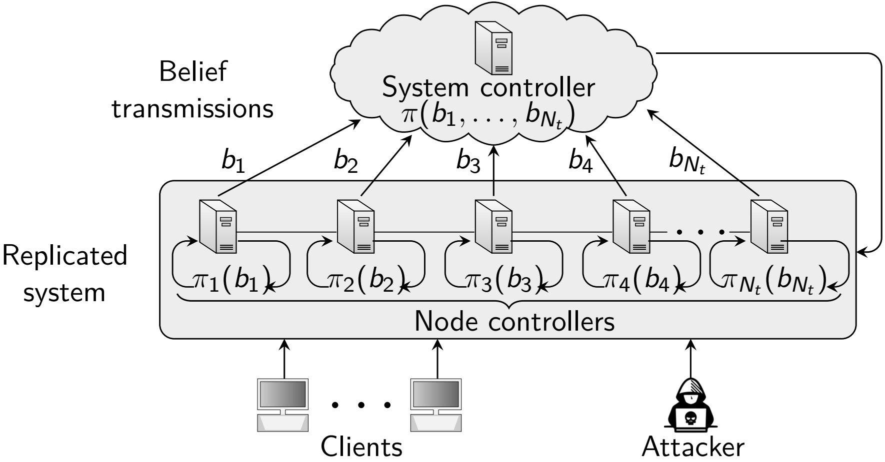

The Global Control Problem

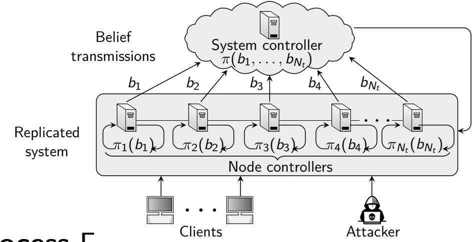

- **Constrained Markov decision process** Γ.
- States: $\mathcal{S}_{\mathrm{S}} = \{0, 1, \ldots, s_{\max}\}$, the number of healthy nodes.
- Controller actions: Add $a_t^{(\mathrm{C})} \in \{0, 1\}$ nodes.
- Dynamics f: depend on the local nodes.
- **Markov** strategy:

$$\pi : \mathcal{S}_{\mathrm{S}} \to \{0, 1\}.$$

System Controller Objective

- Cost: $J \triangleq \lim_{T\to\infty} \sum_{t=1}^{T} \frac{a_t}{T}$.
- Constraint: $T^{(\mathrm{A})} \geq \epsilon_{\mathrm{A}}$, where $T^{(\mathrm{A})}$ is the availability.

ϵ_{A}	*Allowed service downtime per year*
0.9	36 days
0.95	18 days
0.99	3 days
0.999	8 hours
0.9999	52 minutes
0.99999	5 minutes
1	0 minutes

System Reliability Analysis

- The **Mean-time-to-failure** (MTTF) is the mean hitting time of a state where $s_t \leq f$:

$$\mathbb{E}[T^{(\mathrm{F})} \mid S_1 = s_1] = \mathbb{E}_{(S_t)_{t\geq 1}}\Big[\inf\{t \geq 1 \mid S_t \leq f\} \mid S_1 = s_1\Big].$$

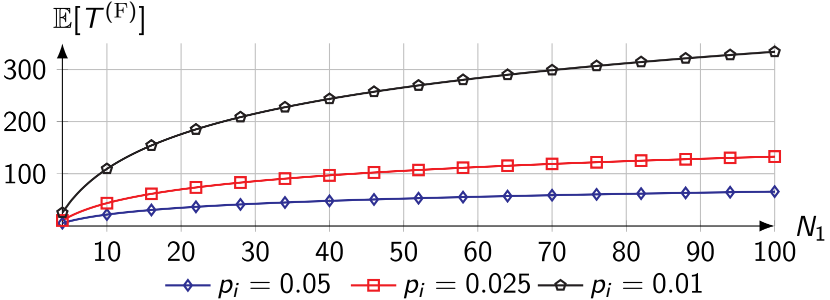

The MTTF in function of the number of initial nodes N_1 and failure probability per node p_i.

Theorem 3 (Optimal Control Strategy Existence)

Assuming

(A) *The Markov chain induced by any control strategy is **unichain**.*

(B) *The availability constraint is feasible.*

Then the following holds.

1. *There exists an optimal stationary replication control strategy.*
2. *The optimal strategy has a threshold structure.*
3. *An optimal replication control strategy can be computed by using **linear programming**.*

Efficient Computation of Optimal Replication Control Strategies

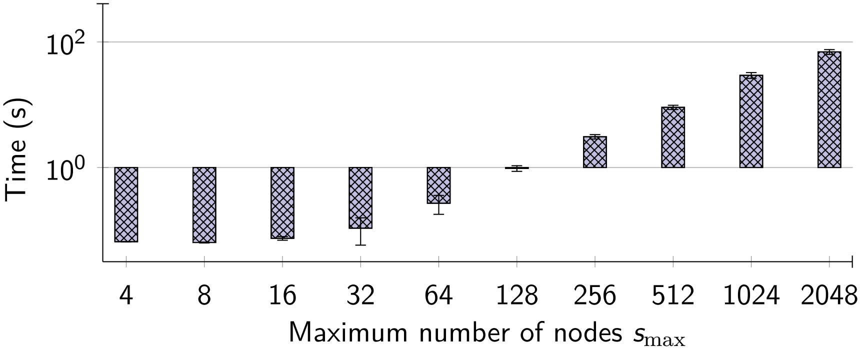

Mean compute time to obtain an optimal replication control strategy.

The Benefit of Optimal Replication Control

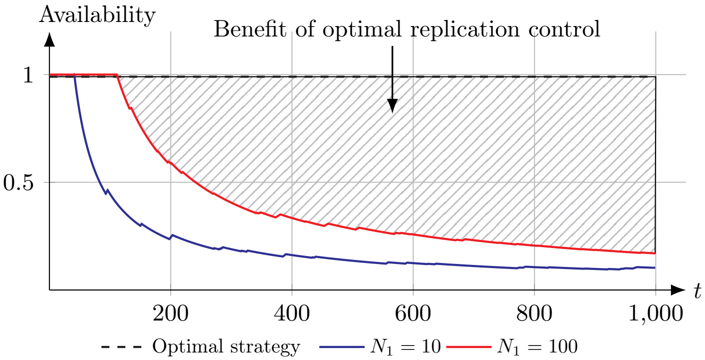

Key insight

Optimal replication control can **guarantee a high service availability in expectation**. The benefit of optimal replication is mainly prominent for long-running systems.

Summary of the Control-Theoretic Model

- **Intrusion recovery control**.
  - Partially observed Markov decision process.
  - Threshold structure of optimal control strategies.
  - Efficient computation through *stochastic approximation*.
- **Replication control**.
  - Constrained Markov decision process.
  - Threshold structure of optimal control strategies.
  - Efficient computation through *linear programming*.

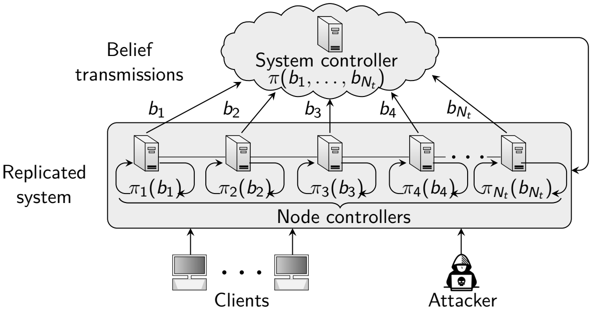

Experiment Setup - Testbed

The TOLERANCE Architecture

Tw**o**-**le**vel **r**ecovery **an**d replication **c**ontrol with f**e**edback.

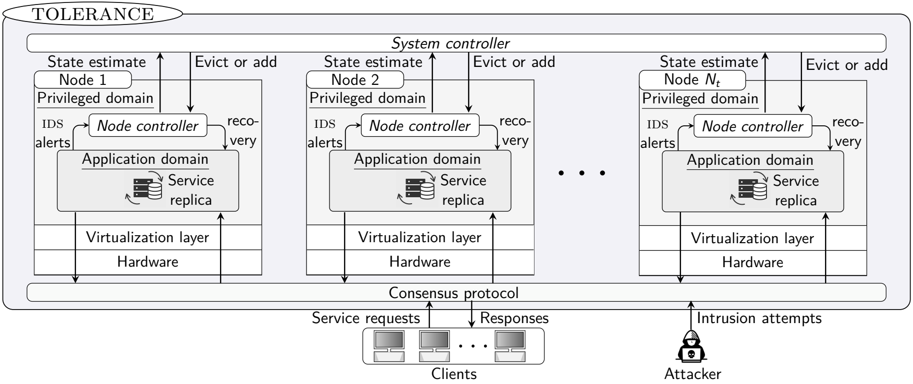

- A replicated web service which offers two operations:
  - A **read** operation that returns the service state.
  - A **write** operation that updates the state.

Intrusion-Tolerant Consensus Protocol (MINBFT)

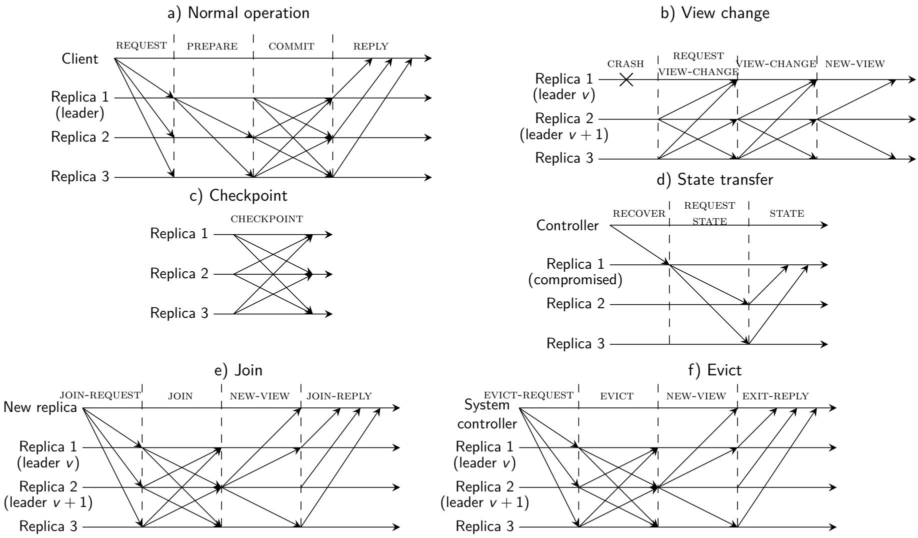

Experiment Setup - Emulated Intrusions

Replica ID	*Intrusion steps*
1	TCP SYN scan, FTP brute force
2	TCP SYN scan, SSH brute force
3	TCP SYN scan, TELNET brute force
4	ICMP scan, exploit of CVE-2017-7494
5	ICMP scan, exploit of CVE-2014-6271
6	ICMP scan, exploit of CWE-89 on DVWA
7	ICMP scan, exploit of CVE-2015-3306
8	ICMP scan, exploit of CVE-2016-10033
9	ICMP scan, SSH brute force, exploit of CVE-2010-0426
10	ICMP scan, SSH brute force, exploit of CVE-2015-5602

Table 1: Intrusion steps.

Experiment Setup - Background Traffic

Background services	*Replica ID(s)*
FTP, SSH, MONGODB, HTTP, TEAMSPEAK	1
SSH, DNS, HTTP	2
SSH, TELNET, HTTP	3
SSH, SAMBA, NTP	4
SSH	5, 7, 8, 10
DVWA, IRC, SSH	6
TEAMSPEAK, HTTP, SSH	9

Table 2: Background services.

Estimated Distributions of Intrusion Alerts

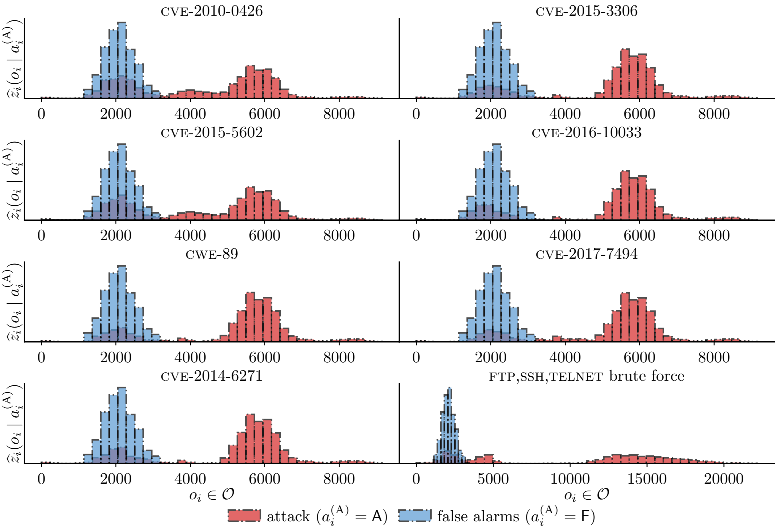

- **We estimate the observation distribution** z with the empirical distribution $\widehat{Z}$ based on M samples.
- $\widehat{z} \rightarrow^{\text{a.s}} z$ as $M \rightarrow \infty$ (Glivenko-Cantelli theorem).

Comparison with State-of-the-art Intrusion-Tolerant Systems

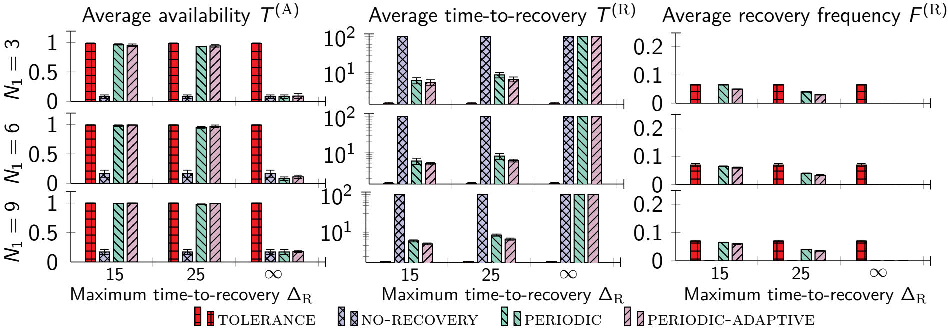

Comparison between our optimal control strategies and the baselines; x-axes indicate values of Δ_R; rows relate to the number of initial nodes N_1.

Conclusion

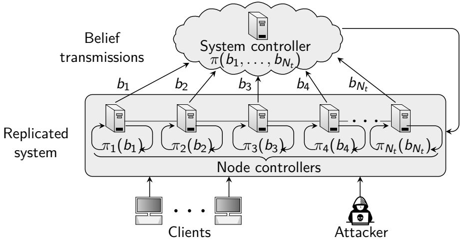

- We present a **control-theoretic model of intrusion tolerance**.
- We establish structural results.
- We evaluate the optimal control strategies on a testbed.
- Our control-theoretic strategies have stronger theoretical guarantees and significantly better practical performance than the heuristic control strategies used in state-of-the-art intrusion-tolerant systems.