

A Clean Slate Approach to Secure Wireless Networking

Jonathan Ponniah

Texas A&M University
jonathan.ponniah@gmail.com

Yih-Chun Hu

University of Illinois at Urbana-Champaign
yihchun@illinois.edu

P. R. Kumar

Texas A&M University
prk@tamu.edu

now

the essence of knowledge

Boston — Delft

Foundations and Trends[®] in Networking

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
United States
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is

J. Ponniah, Y.-C. Hu, and P. R. Kumar. *A Clean Slate Approach to Secure Wireless Networking*. Foundations and Trends[®] in Networking, vol. 9, no. 1, pp. 1–105, 2014.

This Foundations and Trends[®] issue was typeset in L^AT_EX using a class file designed by Neal Parikh. Printed on acid-free paper.

ISBN: 9781680830491

© 2015 J. Ponniah, Y.-C. Hu, and P. R. Kumar

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

Foundations and Trends® in Networking
Volume 9, Issue 1, 2014
Editorial Board

Editor-in-Chief

Anthony Ephremides
University of Maryland
United States

Editors

François Baccelli
University of Texas, Austin

Victor Bahl
Microsoft Research

Helmut Bölcskei
ETH Zurich

J.J. Garcia-Luna Aceves
UC Santa Cruz

Andrea Goldsmith
Stanford University

Roch Guerin
Washington University in Saint Louis

Bruce Hajek
UIUC

Jean-Pierre Hubaux
EPFL

Frank Kelly
University of Cambridge

P.R. Kumar
Texas A&M University

Steven Low
Caltech

Eytan Modiano
MIT

Keith Ross
Polytechnic Institute of NYU

Henning Schulzrinne
Columbia University

Mani Srivastava
UCLA

Leandros Tassioulas
Yale University

Lang Tong
Cornell University

Ozan Tonguz
Carnegie Mellon University

Don Towsley
University of Massachusetts, Amherst

Nitin Vaidya
UIUC

Pravin Varaiya
UC Berkeley

Roy Yates
Rutgers University

Raymond Yeung
Chinese University of Hong Kong

Editorial Scope

Topics

Foundations and Trends[®] in Networking publishes survey and tutorial articles in the following topics:

- Modeling and analysis of:
 - Ad hoc wireless networks
 - Sensor networks
 - Optical networks
 - Local area networks
 - Satellite and hybrid networks
 - Cellular networks
 - Internet and web services
- Protocols and cross-layer design
- Network coding
- Energy-efficiency incentives/pricing/utility-based
- Games (co-operative or not)
- Security
- Scalability
- Topology
- Control/Graph-theoretic models
- Dynamics and asymptotic behavior of networks

Information for Librarians

Foundations and Trends[®] in Networking, 2014, Volume 9, 4 issues. ISSN paper version 1554-057X. ISSN online version 1554-0588. Also available as a combined paper and online subscription.

Foundations and Trends® in Networking
Vol. 9, No. 1 (2014) 1–105
© 2015 J. Ponniah, Y.-C. Hu, and P. R. Kumar
DOI: 10.1561/1300000037



A Clean Slate Approach to Secure Wireless Networking

Jonathan Ponniah
Texas A&M University
jonathan.ponniah@gmail.com

Yih-Chun Hu
University of Illinois at Urbana-Champaign
yihchun@illinois.edu

P. R. Kumar
Texas A&M University
prk@tamu.edu

Contents

1	Introduction	2
2	The Model	13
2.1	The Model Assumptions	18
3	A Game Between Protocols and Adversarial Strategies	21
3.1	A Game Between Protocols and Adversarial Strategies . . .	24
4	Closed Synchronized Networks	28
4.1	Overview	29
4.2	The Neighbor Discovery Phase	32
4.3	The Network Discovery Phase	33
4.4	The Scheduling Phase	34
4.5	The Data Transfer Phase	34
4.6	The Verification Phase	35
4.7	Steady State Behavior	35
4.8	Feasibility of Protocol and Optimality Proof	36
4.9	Concluding Remarks	48
5	Secure Clock Synchronization	50
5.1	Previous Work	52
5.2	A Motivating Example: The Man-In-The-Middle Attack . . .	53

5.3	The Clock Model	54
5.4	Impossibility Results in Clock Synchronization	55
5.5	Clock Synchronization in a Cycle	58
5.6	The Consistency Check	59
5.7	Main Result	63
5.8	Application to General Networks	75
5.9	Concluding Remarks	76
6	Closed Unsynchronized Networks	78
6.1	Model Assumptions	81
6.2	An Orthogonal MAC code	82
6.3	Forming a Rudimentary Synchronized Network	88
6.4	Controlling the Overhead	94
6.5	Conclusions	97
7	Conclusion	98
	Acknowledgements	100
	References	101

Abstract

The design of secure protocols for wireless ad-hoc networks is an important problem in communication systems research. A seemingly fundamental limitation of the design process is that any new protocol only addresses the vulnerabilities detected in its predecessors, leaving the remaining vulnerabilities unaffected. Hence, the design process amounts to an arms race between more sophisticated attacks and protocol fixes. To change this situation, a framework is needed for secure protocol design that offers provable performance and security guarantees against all possible attacks on the network.

This monograph proposes such a framework, contingent on some underlying model assumptions. The framework consists of a game defined between protocols and adversarial strategies in which the adversarial strategy is selected after the protocol has been revealed to all of the nodes. Each choice of protocol and adversarial strategy results in a payoff that corresponds to the functionality retained by the network, despite the adversarial activity. The design imperative is to choose the protocol that maximizes this payoff.

Two scenarios are considered: networks in which the nodes are initially synchronized and unsynchronized respectively. In each scenario, a protocol is described and three results are proved. First, the protocol is max-min optimal with respect to the payoff. The max-min payoff is the best that can be achieved because the protocol is always known to the adversarial nodes before the adversarial strategy is chosen. Second, the protocol is min-max optimal; there is a Nash equilibrium in the space of protocols and adversarial strategies. By implication, the adversarial nodes gain no strategic advantage from knowing protocol a priori. Finally, the adversarial nodes are effectively confined to one of two behavior modes: either jam or conform to the protocol, neither of which can be prevented by any protocol.

1

Introduction

A communication network is a set of terminal nodes (sources and destinations), intermediate nodes, and telecommunication links through which certain pairs of nodes are directly connected. The purpose of such a network is to route information generated by the source terminal nodes in the form of bits, via a connecting path of links, to the corresponding destination terminal nodes. These bits are transmitted electromagnetically across each link in the connecting path.

In wireless networks, as opposed to wired networks, the communication medium is simultaneously shared by all terminal nodes; the signals transmitted at one terminal can potentially reach all other terminals in the network. As a result, the signals of two terminals that transmit simultaneously may interfere with each other at the corresponding destination nodes, preventing either node from decoding the intended message. An *independent set* is a set of simultaneously activated links in which such mutual interference does not occur, and is determined by the physical properties of the wireless channel.

In order for messages to be reliably transmitted between any source-destination pair, either a centralized controller or the nodes themselves must know or discover the set of links in the network as well as the

independent sets of links, and by extension, the routes connecting each source-destination pair, and the joint selection of transmit power levels that correspond to an independent set. Furthermore, the centralized controller or the nodes themselves must decide on a desired throughput vector between all source-destination pairs and allocate the network resources to support this desired throughput.

Wireless *ad-hoc* networks, the focus of this monograph, are wireless networks that carry out these tasks without the aid of a centralized controller. Instead the nodes themselves must jointly determine a schedule of transmissions and receptions, consistent with the independent sets in the network, without a priori knowledge of the network topology. In addition, the nodes may be half-duplex in the sense that they cannot simultaneously transmit and receive, thus complicating the process by which the nodes can even acquire the minimal knowledge to form a rudimentary network.

A set of instructions given to each node, known as a *protocol*, enables the nodes as a collective to form a reliable network in which data can be transported. A *legitimate node*, by definition, follows the protocol exactly. An *adversarial node*, on the other hand, can behave arbitrarily or with malicious intent to prevent the legitimate nodes from forming a network. In this monograph, we confront the challenge of designing protocols for wireless ad-hoc networks infiltrated with adversarial nodes. Part of the challenge in designing such protocols is that many of the protocols already in widespread use predate contemporary demands for security, and are instead tailored for performance and low complexity.

We will provide a general description of the kinds of protocols that have been developed over the past few decades before focusing on protocols that specifically address issues of security.

Typically, protocols are either categorized as *reactive* or *pro-active*. Reactive (or “on-demand”) protocols determine a route from a source to a destination, if it is not already known, only after a source requests it. On the other hand, table-driven or proactive protocols maintain tables of all the known routes in the network, and ensure at considerable overhead, that these tables are periodically updated. Consequently, data packets can be routed more quickly since the routes are known a pri-

ori. The choice of a pro-active or re-active protocol reflects a trade-off between faster routing times and larger overhead, and is based on the computational power available in the network and the quality of service demanded by the users.

Another way of categorizing protocols is by “link-state” or “distance-vector routing”. A link-state protocol enables each node to retain a view of the entire network topology, whereas a distance-vector routing protocol allows each node to monitor its distance (in hops) from every destination node. The latter requires less overhead but is more sensitive to topological changes that alter the routes.

In general, a protocol caters to the specific computational limitations and environmental constraints in which the network operates. As a result, some protocols best fall into hybrid categories such as pro-active link state or on-demand distance vector routing. The diverse range of protocols available today is the outcome of extensive research and development to tailor protocols for every environmental contingency.

Within this larger set of protocols is a smaller subset developed explicitly with security and the possibility in mind that some of the nodes, being adversarial, may not cooperate with protocol. One particular strategy to deal with adversarial nodes is to add redundancy by dispersing essential activity throughout the network. To that end, there are protocols that replicate computational tasks [1] [2] [8] [9] [14] [15] [37] [40] [41], and storage tasks [3] [5] [7] [13] [35]. However, there is an adversarial counter-strategy that negates the defense offered by redundancy; the adversarial nodes could adopt multiple identities so as to artificially inflate their numbers. This attack (affectionately named “Sybil” after a fictional character suffering from multiple personality disorder) defeats any network that lacks a trusted central authority, regardless of the internal mechanisms employed to self-authenticate member identities.

A protocol strategy, to “counter” this adversarial counter-strategy, requires reliable methods for authenticating the authorship and content of network communication; in other words, encryption. Methods of encryption fall into one of two categories: symmetric, wherein the

users share a single key, or asymmetric wherein the source has a private key and the end-users have a public key. Symmetric encryption is more computationally efficient but is unable to prevent a malicious node that possess the key from forging packets or impersonating a user. Asymmetric encryption does not share this deficiency, but is computationally more complex than its symmetric counterpart by several orders of magnitude.

Encryption is also important because a small number of adversarial nodes can severely disrupt the network by forging or tampering with the control packets in a routing protocol. Some of the symptoms of such tampering include routing loops in which packets traverse a cycle and never reach their destination, black holes in which packets are dropped, or grey holes in which packets are selectively dropped, gratuitous detours which are unnecessarily long and/or circuitous routes when shorter routes are available, and artificial partitions of the network. These types of attacks, in which an adversarial node can disproportionately affect the network operation, are referred to as denial-of-service attacks.

To protect the network from Sybil and other denial-of-service attacks, some protocols incorporate specific encryption mechanisms into their operations. Examples include secure pro-active routing protocols with asymmetric encryption [30] [21] [23], [39], secure on-demand routing protocols with asymmetric encryption [44] [42] [38], secure pro-active routing protocols with symmetric encryption [4] [16] [43], and secure routing protocols with symmetric encryption for limited topologies such as between nodes and a base station [32], or between communicating routers [17].

This variety is partly due to the computational complexity of asymmetric encryption. Networks that use asymmetric encryption, but have limited computational capabilities are also vulnerable to a denial of service attack in which malicious nodes flood the network with bogus encrypted packets that are computationally expensive to decode. The resulting dilemma is whether to use symmetric encryption and accept the risk of packet tampering and forgery, or use asymmetric encryption and expose the network to a denial of service attack.

The TESLA broadcast authentication protocol [31] attempts to resolve this dilemma by using symmetric encryption in conjunction with a loosely synchronized network to periodically release new keys and ensure freshness. The protocol aims to provide the security inherent in asymmetric encryption but retain the smaller complexity of symmetric encryption. Two routing protocols that incorporate TESLA authentication methods are Ariadne [20], a secure on-demand routing protocol, and SEAD [18], a secure pro-active routing protocol.

Finally, some protocols attempt to secure the routing mechanism without recourse to encryption. For example, the Watchdog and Pathrater protocol [29] maintains a blacklist of nodes that appear to misbehave and penalizes them in the algorithm for determining routes. However, the protocol must strike a careful balance when penalizing nodes, because some of the nodes in the blacklist might be victims of false reporting by malicious nodes. To account for this possibility, the protocol also includes a process that allows nodes to be rehabilitated after they have been blacklisted.

Despite the efforts to secure networks via encryption and other means, structural vulnerabilities in certain protocol categories can still escape detection. The rushing attack [19] exposes such a vulnerability in the way route requests are forwarded in most source-routing protocols including those hardened with encryption. The proposed fix, called RAP (Rushing Attack Prevention) uses a modified version of DSR (Dynamic Source Routing) in which route requests from a source are stored and buffered, and one request from the buffer is randomly forwarded while the rest are discarded. The wormhole attack [6] is another example of an attack on a structural vulnerability. In this attack, an adversarial node creates an artificial link between two legitimate nodes in an attempt to control a substantial fraction of the network traffic. A patch for this attack uses what are called “packet leashes” to prevent packets from travelling too far from their transmitter.

It is also possible for an adversarial strategy to combine these attacks so as to increase their destructive effect. For instance, an adversarial node could hypothetically create a wormhole to tunnel a route request, accelerating a rushing attack, while simultaneously flooding

other routes with bogus encrypted packets to slow down the arrival of competing route requests.

There is an underlying trend in this survey of protocols for wireless ad-hoc networks and secure protocols specifically. That trend can be described as an arms race between attacks and fixes; in every protocol is a hidden vulnerability that once discovered leads to the development of a patch for that attack. Over time, the attacks become more sophisticated and the patches more elaborate. At no point in this process is it possible to provide any security guarantees or know whether there are vulnerabilities that have yet to be discovered.

The purpose of this monograph is to develop a system-theoretical approach to secure protocol design that provides provable and comprehensive security guarantees, where any features of a protocol that reduce complexity or maximize the network throughput do not come the expense of security. The desired approach is one of security first and performance second, which, as apparent from the survey, is the reverse of the current approach.

The efforts to design secure wireless ad-hoc networks are part of a deeper and more pervasive engineering challenge; designing complex systems that are secure. A system refers to a set of “interacting or interconnected components that together form an integrated whole”. Wireless ad-hoc networks are but one example of complex systems; other examples include economic systems, military systems, software systems, and power systems.

All of these systems are designed to carry out some system-specific function. In the case of wireless ad-hoc networks, the function is clear; to facilitate the exchange of information between all source-destination node pairs. A secure system, by definition retains its functionality even if some of the individual components of which it is composed, are compromised or attacked. Note that the term “security”, when applied to communication systems, can also be used in another context; to protect the privacy of information exchanged between a source and destination. We will limit our analysis exclusively to the former definition and not the latter. By this definition, a secure wireless ad-hoc network is one in which the legitimate source-destination nodes in the network are

able to reliably exchange data despite the non-cooperative or hostile behavior of the adversarial nodes. Examples of such behavior include all of the attacks described in the survey, in addition to less subtle acts of non-compliance; the adversarial nodes could choose to drop packets, advertise a wrong hop-count, jam, refuse to acknowledge a neighbor, lie about the local topology, in short do anything to undermine the operation of the network.

Naturally, there is a great deal of public interest and research invested in making complex systems secure. The philosophy that best describes the current approach towards secure system design is “defense-in-depth”; compartmentalize the system into more manageable subsystems, and defend each individual subsystem in a way that ensures the system functionality is retained even if other subsystems fail. This approach is understandable, given the near impossibility of comprehensively identifying all possible failure modes in a large complex system.

However, as evident in the application to wireless ad-hoc networks, the defense-in-depth philosophy has some inherent weaknesses. Namely, the possible existence of structural vulnerabilities that underly individual subsystems and the difficulty of completely anticipating, even within a subsystem, all of the ways in which an attack can occur.

In this monograph we propose a system-theoretic approach to secure protocol design in which comprehensive, provable security and performance guarantees can be made, contingent on some underlying model assumptions. As long as the assumptions hold, the guarantees retain their validity; to attack the protocol one must attack the underlying assumptions, and to harden the protocol to these types of attacks, the corresponding assumptions must increase in generality.

In a system-theoretic approach the notions of network “functionality” and “security” (the network functionality “retained” despite adversarial activity), must both be defined quantitatively. We will propose the following game-theoretic framework that addresses these requirements.

First the protocol is announced to all nodes, both legitimate and adversarial, since the adversarial nodes are not known to the protocol or the legitimate nodes a priori. Next, the adversarial nodes choose a

strategy q_p specifically tailored to the protocol p . At time $t = 0$ with respect to some global reference clock, each node turns on and proceeds to execute either p or q_p depending on whether the node is legitimate or adversarial respectively. Over a fixed operating lifetime, the source-destination node pairs exchange some number (possibly zero) of data packets. Let x denote the resulting throughput vector of all source-destination pairs in the network. We will assume that the “functionality” of the network is measured by some utility function $U(x)$ of x , where $U(x)$ reflects the priority that one source-destination pair in x is given relative to the others.

The “payoff” of this zero sum game, $J(p, q_p)$, is a quantitative measure of the “functionality retained” by the system, in this case a wireless ad-hoc network operating according to protocol p , in the aftermath of the adversarial strategy q_p . There are many ways of defining $J(p, q_p)$ but in general, one should expect $J(p, q_p)$ to depend on the legitimate source-destination pairs in x , the effective throughput attained by these source-destination pairs, and the utility function $U(x)$.

The problem of secure protocol design in this context is to choose the protocol \tilde{p} that achieves the max-min payoff in the zero sum game defined by $J(p, q_p)$. In other words, given a priori knowledge of the protocol p , the adversarial nodes, in the worst case, choose a strategy \tilde{q}_p to minimize the payoff $J(p, q_p)$ over all other adversarial strategies. The best any protocol can do, is maximize the worst-case payoff. That is, the best payoff that can be achieved over all adversarial strategies is:

$$\max_{\text{protocols } p} \min_{\text{attacks } q_p} J(p, q_p). \quad (1.1)$$

The max-min payoff in (1.1) is the best that any protocol can achieve no matter how the payoff function $J(p, q_p)$ is defined; the adversarial nodes always know the protocol a priori and can always choose a strategy that minimizes the payoff for a specific protocol. Hence a protocol that achieves (1.1) is proveably secure in the sense that the max-min payoff is guaranteed regardless of what the adversarial nodes do, and no protocol is guaranteed of doing better. For these reasons, the game-theoretic framework addresses the deficiencies in the defense-

in-depth approach to securing complex systems; namely, by allowing a quantitative measure of system security in the form of a payoff function, and the possibility of comprehensive, proveable security guarantees in the form of a protocol that is max-min optimal with respect to the payoff. These guarantees however, are contingent on the model assumptions that decide the payoff $J(p, q_p)$ for a specific choice of protocol p and adversarial strategy q_p .

The challenge then, is to find the protocol p that is max-min optimal with respect to $J(p, q_p)$. This problem is non-trivial because the strategy space of protocols and adversarial strategies is generally undefined. In the chapters to follow we will prove the following main results. First, we will describe a protocol that achieves the max-min payoff in (1.1) for a specific payoff $J(p, q_p)$ and a specific set of model assumptions. Next we will show that the protocol does even better and achieve the following payoff:

$$\min_{\text{attacks } q} \max_{\text{protocols } p} J(p, q). \quad (1.2)$$

Since the min-max is always greater than or equal to the max-min, and the max-min is the best any protocol can achieve, it follows that this game has a Nash equilibrium. In other words, there is a saddle-point in the optimization over protocols and adversarial strategies. The implication of (1.2) is that the attackers gain no advantage (in terms of reducing the payoff) by knowing the protocol beforehand; the payoff in (1.2) is the payoff achieved if the adversarial strategy is announced before the protocol.

Finally, let \mathcal{Q} denote the set of strategies in which the adversarial nodes are restricted to either jamming or conforming to the protocol, actions which no protocol can prevent from occurring (in Chapter 3, we will explain why an adversarial node may choose to cooperate with the protocol). The following payoff is achievable:

$$\min_{\text{attacks } q \in \mathcal{Q}} \max_{\text{protocols } p} J(p, q). \quad (1.3)$$

We will describe a protocol that achieves (1.1), (1.2), and (1.3) for two different sets of model assumptions, increasing in generality, that concern the pre-existing level of clock synchronization in the network.

There are several reasons why clock synchronization specifically, and the notion of time more generally, is relevant to the discussion of wireless ad-hoc network security. First, the notion of time is essential for evaluating the effective throughput vector x , that is, the total number of data bits exchanged for each source-destination pair over the operating lifetime. The effective throughput is an important metric of the functionality, and by extension, security, of a wireless ad-hoc network.

In addition, a common reference time is also useful for distributed half-duplex nodes to coordinate their activity and/or jointly execute a common schedule; recall that in the wireless setting, two active links cause mutual interference if they don't belong to the same independent set. Moreover, even if all other nodes are silent, a pair of adjacent, half-duplex nodes will be prevented from exchanging information if both nodes are always simultaneously in transmit mode; a half-duplex node, by definition, is unable to simultaneously transmit and receive messages. For these reasons, clock synchronization is especially pertinent in the operation of wireless ad-hoc networks. However, achieving clock synchronization in a network with adversarial nodes is no trivial task; secure clock synchronization is an active area of research in distributed systems theory.

To simplify the presentation, we will first describe a protocol that is provably secure under the following assumptions: the nodes turn "on" simultaneously and the local clocks at each node tick at the same rate. The resulting network is *closed* in the sense that the network once formed, does not expand to include new nodes, since no nodes turn "on" after the primordial birth. The nodes are also synchronized, since by assumption the local clocks are not subject to either an offset or a skew. Later, we will modify the protocol and show that similar security guarantees can be made for closed networks with unsynchronized, relatively affine local clocks. These results also appear in [33].

In Chapter 2 we describe the model assumptions on which the results are predicated. The assumptions describe valid modes of communication, topological constraints, encryption capabilities, and the parameters of the local clocks. The goal in this chapter is to chose model

assumptions that impose the minimal necessary conditions needed to provide any security guarantees.

In Chapter 3 we define a specific payoff function $J(p, q_p)$ in terms of any utility function $U(x)$. There are many ways in which $J(p, q_p)$ can be defined, but ultimately the payoff should correspond to the functionality “retained” by the network despite adversarial activity. The results in this monograph apply only for the payoff $J(p, q_p)$ defined in Chapter 3. In Chapter 4 we describe a protocol that operates under the assumption each node is synchronized and the nodes are born simultaneously. We will show that protocol achieves the payoff defined in (1.1), (1.2), and (1.3). In Chapter 5 we consider the model in which the local clocks are relatively affine. We will discuss the difficulties in synchronizing affine clocks when the network is infiltrated with adversarial nodes. Finally, in Chapter 6 we revisit Chapter 4 and show that the results in Chapter 4 hold under the assumption of relatively affine, unsynchronized local clocks. We will offer our conclusions and propose areas for further research in Chapter 7.

References

- [1] William J. Bolosky, John R. Douceur, David Ely, and Marvin Theimer. Feasibility of a serverless distributed file system deployed on an existing set of desktop pcs. In *Proceedings of the 2000 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '00, pages 34–43, New York, NY, USA, 2000. ACM.
- [2] Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, OSDI '99, pages 173–186, Berkeley, CA, USA, 1999. USENIX Association.
- [3] David L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2):84–90, February 1981.
- [4] S. Cheung. An efficient message authentication scheme for link state routing. In *Computer Security Applications Conference, 1997. Proceedings., 13th Annual*, pages 90–98, 1997.
- [5] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. Private information retrieval. *J. ACM*, 45(6):965–981, November 1998.
- [6] Yih chun Hu, Adrian Perrig, and David B. Johnson. Wormhole attacks in wireless networks. *IEEE Journal on Selected Areas in Communications*, 24:370–380, 2006.

- [7] Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In *International Workshop on Designing Privacy Enhancing Technologies: Design Issues in Anonymity and Unobservability*, pages 46–66, New York, NY, USA, 2001. Springer-Verlag New York, Inc.
- [8] Frank Dabek, M. Frans Kaashoek, David Karger, Robert Morris, and Ion Stoica. Wide-area cooperative storage with cfs. In *Proceedings of the Eighteenth ACM Symposium on Operating Systems Principles, SOSP '01*, pages 202–215, New York, NY, USA, 2001. ACM.
- [9] R. Dingledine, M.J. Freedman, and D. Molnar. The Free Haven Project: Distributed Anonymous Storage Service. *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability, Berkeley, CA, USA, July 25-26, 2000: Proceedings*, 2001.
- [10] Danny Dolev, Joe Halpern, and H. Raymond Strong. On the possibility and impossibility of achieving clock synchronization. In *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing, STOC '84*, pages 504–511, New York, NY, USA, 1984. ACM.
- [11] Rui Fan and Nancy Lynch. Gradient clock synchronization. *Distrib. Comput.*, 18(4):255–266, March 2006.
- [12] Nikolaos M. Freris, Scott R. Graham, and P. R. Kumar. Fundamental limits on synchronizing clocks over networks. *IEEE Trans. Automat. Contr.*, 56(6):1352–1364, 2011.
- [13] Yael Gertner, Shafi Goldwasser, and Tal Malkin. A random server model for private information retrieval or how to achieve information theoretic privacy avoiding database replication. In *Proceedings of the Second International Workshop on Randomization and Approximation Techniques in Computer Science, RANDOM '98*, pages 200–217, London, UK, UK, 1998. Springer-Verlag.
- [14] Andrew V. Goldberg and Peter N. Yianilos. Towards an archival inter-memory. In *ADL*, pages 147–156. IEEE Computer Society, 1998.
- [15] J.H. Hartman, I. Murdock, and T. Spalink. The swarm scalable storage system. In *Distributed Computing Systems, 1999. Proceedings. 19th IEEE International Conference on*, pages 74–81, 1999.
- [16] R. Hauser, T. Przygienda, and G. Tsudik. Reducing the cost of security in link-state routing. In *Network and Distributed System Security, 1997. Proceedings., 1997 Symposium on*, pages 93–99, 1997.

- [17] A. Heffernan. Protection of bgp sessions via the tcp md5 signature option, 1998.
- [18] Yih-Chun Hu, David B. Johnson, and Adrian Perrig. Sead: secure efficient distance vector routing for mobile wireless ad hoc networks. *Ad Hoc Networks*, 1:175–192, 2003.
- [19] Yih-Chun Hu, Adrian Perrig, and David B. Johnson. Rushing attacks and defense in wireless ad hoc network routing protocols. In *Proceedings of the 2Nd ACM Workshop on Wireless Security, WiSe '03*, pages 30–40, New York, NY, USA, 2003. ACM.
- [20] Yih-Chun Hu, Adrian Perrig, and David B. Johnson. Ariadne: A secure on-demand routing protocol for ad hoc networks. *Wireless Networks*, 11(1-2):21–38, January 2005.
- [21] Stephen Kent, Charles Lynn, Joanne Mikkelsen, and Karen Seo. Secure border gateway protocol (s-bgp. *IEEE Journal on Selected Areas in Communications*, 18:103–116, 2000.
- [22] Fabian Kuhn, Christoph Lenzen, Thomas Locher, and Rotem Oshman. Optimal gradient clock synchronization in dynamic networks. In *Proceedings of the 29th ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, PODC '10*, pages 430–439, New York, NY, USA, 2010. ACM.
- [23] Brijesh Kumar. Integration of security in network routing protocols. *SIGSAC Rev.*, 11(2):18–25, April 1993.
- [24] Leslie Lamport and P. M. Melliar-Smith. Byzantine clock synchronization. In *Proceedings of the Third Annual ACM Symposium on Principles of Distributed Computing, PODC '84*, pages 68–74, New York, NY, USA, 1984. ACM.
- [25] Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, July 1982.
- [26] Christoph Lenzen, Thomas Locher, and Roger Wattenhofer. Clock synchronization with bounded global and local skew. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science, FOCS '08*, pages 509–518, Washington, DC, USA, 2008. IEEE Computer Society.
- [27] Thomas Locher and Roger Wattenhofer. Oblivious gradient clock synchronization. In *Proceedings of the 20th International Conference on Distributed Computing, DISC'06*, pages 520–533, Berlin, Heidelberg, 2006. Springer-Verlag.

- [28] Jennifer Lundelius and Nancy A. Lynch. An upper and lower bound for clock synchronization. *Information and Control*, 62(2/3):190–204, August/September 1984.
- [29] Sergio Marti, T. J. Giuli, Kevin Lai, and Mary Baker. Mitigating routing misbehavior in mobile ad hoc networks. In *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking*, MobiCom '00, pages 255–265, New York, NY, USA, 2000. ACM.
- [30] Radia Perlman. *Interconnections: Bridges and Routers*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA, 1992.
- [31] Adrian Perrig, Ran Canetti, J.Đ. Tygar, and Dawn Song. The tesla broadcast authentication protocol. *RSA CryptoBytes*, 5, 2002.
- [32] Adrian Perrig, Robert Szcwczyk, J. D. Tygar, Victor Wen, and David E. Culler. Spins: Security protocols for sensor networks. *Wirel. Netw.*, 8(5):521–534, September 2002.
- [33] Jonathan Ponniah, Yih-Chun Hu, and P. R. Kumar. A system-theoretic clean slate approach to provably secure ad hoc wireless networking. *IEEE Transactions on Control of Network Systems*, To appear.
- [34] Jonathan Ponniah, Yih-Chun Hu, and P.R. Kumar. An orthogonal multiple access coding scheme. *Communications in Information and Systems*, 12(1):41–76, 2012.
- [35] Michael K. Reiter and Aviel D. Rubin. Anonymous web transactions with crowds. *Commun. ACM*, 42(2):32–48, February 1999.
- [36] R.S. Robles, J.J. Haas, J.T. Chiang, Yih-Chun Hu, and P.R. Kumar. Secure topology discovery through network-wide clock synchronization. In *Signal Processing and Communications (SPCOM), 2010 International Conference on*, pages 1–5, July 2010.
- [37] Antony Rowstron and Peter Druschel. Storage management and caching in past, a large-scale, persistent peer-to-peer storage utility. In *Proceedings of the Eighteenth ACM Symposium on Operating Systems Principles*, SOSP '01, pages 188–201, New York, NY, USA, 2001. ACM.
- [38] K. Sanzgiri, B. Dahill, B.N. Levine, C. Shields, and E.M. Belding-Royer. A secure routing protocol for ad hoc networks. In *Network Protocols, 2002. Proceedings. 10th IEEE International Conference on*, pages 78–87, 2002.
- [39] B.R. Smith, S. Murthy, and J.J. Garcia-Luna-Aceves. Securing distance-vector routing protocols. In *Network and Distributed System Security, 1997. Proceedings., 1997 Symposium on*, pages 85–92, 1997.

- [40] Marc Waldman, Aviel D. Rubin, and Lorrie Faith Cranor. Publius: A robust, tamper-evident, censorship-resistant web publishing system. In *Proceedings of the 9th Conference on USENIX Security Symposium - Volume 9*, SSYM'00, pages 5–5, Berkeley, CA, USA, 2000. USENIX Association.
- [41] J.J. Wylie, M.W. Bigrigg, J.D. Strunk, G.R. Ganger, H. Kiliccote, and P.K. Khosla. Survivable information storage systems. *Computer*, 33(8):61–68, 2000.
- [42] Manel Guerrero Zapata. Secure ad hoc on-demand distance vector routing. *SIGMOBILE Mob. Comput. Commun. Rev.*, 6(3):106–107, June 2002.
- [43] Kan Zhang. Efficient protocols for signing routing messages, 1998.
- [44] Lidong Zhou and Z.J. Haas. Securing ad hoc networks. *Network, IEEE*, 13(6):24–30, 1999.