Department of Communications and Networking

Security mechanisms in partially isolated networks

Philip Ginzboorg

A!

DOCTORAL DISSERTATIONS

Aalto University publication series
DOCTORAL DISSERTATIONS 71/2014

Security mechanisms in partially isolated networks

Philip Ginzboorg

A doctoral dissertation completed for the degree of Doctor of Science (Technology) to be defended, with the permission of the Aalto University School of Electrical Engineering, at a public examination held at the lecture hall S5 of the school on 27 May 2014 at 13.

Aalto University
School of Electrical Engineering
Department of Communications and Networking
Research Group for Protocol, Services, and Software

Supervising professor
Jörg Ott

Preliminary examiners
Doctor Giovanni Neglia, INRIA - Sophia-Antipolis, France

Doctor Stephen Farrell, Trinity College Dublin, Ireland

Opponent
Professor Gene Tsudik, University of California, Irvine, U.S.A.

Aalto University publication series
DOCTORAL DISSERTATIONS 71/2014

ISBN 978-952-60-5692-0
ISBN 978-952-60-5693-7 (pdf)
ISSN-L 1799-4934
ISSN 1799-4934 (printed)
ISSN 1799-4942 (pdf)
http://urn.fi/URN:ISBN:978-952-60-5693-7

Unigrafia Oy
Helsinki 2014

Finland

Aalto University

Abstract

Aalto University, P.O. Box 11000, FI-00076 Aalto www.aalto.fi

Author
Philip Ginzboorg

Name of the doctoral dissertation
Security mechanisms in partially isolated networks

Publisher School of Electrical Engineering

Unit Department of Communications and Networking

Series Aalto University publication series DOCTORAL DISSERTATIONS 71/2014

Field of research Security, networking technology

Manuscript submitted 28 April 2014 **Date of the defence** 27 May 2014

Permission to publish granted (date) 6 May 2014 **Language** English

☐ **Monograph** ☒ **Article dissertation (summary + original articles)**

Abstract

A challenged network is a network subject to difficult operational constraints, like disrupted links and high delays. An example is a mobile ad hoc network where the nodes communicate without infrastructure support. Specifics of such networks are the consequence of intermittent access to (or complete lack of) infrastructure and partial isolation of the network nodes from each other. Partial isolation of the nodes aggravates resource scarcity. In this work we investigate (1) security and privacy, as well as (2) resource management, in these networks. Our thesis is that in partially isolated networks security and privacy, on the one hand, and resource management, on the other, are closely related and tend to influence each other.

This dissertation includes nine publications. Their summaries are grouped in two chapters. The first, Chapter 2, is about security and privacy issues and also the question of congested buffer management, which we approach via adversarial scenario. In the second, Chapter 3, we consider better utilization of the scarce contact time between the network nodes via message fragmentation when there are intermittent transmission opportunities.

Part of the work comprising this dissertation had industrial impact: I have contributed to the Generic Authentication Architecture standard of 3GPP, described in publication II; and the "Equal Subdomains" technique found in publication IV was included of the Nokia Awarenet platform.

Keywords security, privacy, resource management, fragmentation, DTN

ISBN (printed) 978-952-60-5692-0 **ISBN (pdf)** 978-952-60-5693-7

ISSN-L 1799-4934 **ISSN (printed)** 1799-4934 **ISSN (pdf)** 1799-4942

Location of publisher Helsinki **Location of printing** Helsinki **Year** 2014

Pages 193 **urn** http://urn.fi/URN:ISBN:978-952-60-5693-7

Tekijä
Philip Ginzboorg

Väitöskirjan nimi
Turvallisuusmekanismit osittain eristäytyneissä verkoissa

Julkaisija Sähkötekniikan korkeakoulu

Yksikkö Tietoliikenne- ja tietoverkkotekniikan laitos

Sarja Aalto University publication series DOCTORAL DISSERTATIONS 71/2014

Tutkimusala Turvallisuus, verkkoteknologia

Käsikirjoituksen pvm 28.04.2014 **Väitöspäivä** 27.05.2014

Julkaisuluvan myöntämispäivä 06.05.2014 **Kieli** Englanti

☐ **Monografia** ☒ **Yhdistelmäväitöskirja (yhteenveto-osa + erillisartikkelit)**

Tiivistelmä

Haastavilla verkoilla (Challenged Network) tarkoitetaan verkkoja, joissa esiintyy toiminnallisia rajoituksia, kuten häirittyjä linkkejä ja korkeita viiveitä. Haastaviin verkkoihin lukeutuvat esimerkiksi mobiilit ad hoc-verkot, joissa kommunikointi tapahtuu ilman infrastruktuurin tukea. Tällaisten verkkojen erityispiirteitä ovat solmujen osittainen tai täydellinen eristäytyneisyys toisistaan. Solmujen osittainen eristäytyminen vaikeuttaa vähäisten resurssien höydyntämistä. Tässä väitöskirjassa tutkitaan osittain eristäytyneiden verkkojen turvallisuutta, yksityisyyttä sekä resurssien hallintaa. Työssä osoitetaan, että näissä verkoissa: yhtäältä verkkojen turvallisuuden ja yksityisyyden hallinta, sekä toisaalta resurssien hallinta riippuvat ja vaikuttavat toisiinsa.

Väitöskirja on koostettu yhdeksästä julkaisusta joiden yhteenveto on esitetty kahdessa luvussa. Näistä ensimmäisessä, luvussa kaksi, käydään läpi turvallisuuteen ja yksityisyyteen liittyviä kysymyksiä sekä pohditaan ruuhkaisen puskurin hallintaa erityisesti hyökkääjän näkökulmasta. Jälkimmäisessä luvussa, luvussa 3, käsitellään lyhyiden yhteysaikojen hyödyntämistä verkon solmujen välisessä viestinnässä. Lyhyitä yhteysaikoja voidaan hyödyntää jakamalla viesti pienempiin osiin, jos mahdollisuuksia lyhyisiin yhteyksiin tarjoutuu riittävän usein.

Osa väitöskirjaan tehdystä työstä on vaikuttanut myös teollisuuteen. Julkaisussa kaksi esitettyjä menetelmiä on sisällytetty 3GPP:ssä määriteltyyn "Generic Authentication Architecture" -standardiin. Lisäksi julkaisussa neljä esitetty menetelmä "Equal Subdomains" on mukana Nokian Awarenet-alustassa.

Avainsanat Turvallisuus, yksityisyys, resurssien hallinta, fragmentointi, DTN

ISBN (painettu) 978-952-60-5692-0 **ISBN (pdf)** 978-952-60-5693-7

ISSN-L 1799-4934 **ISSN (painettu)** 1799-4934 **ISSN (pdf)** 1799-4942

Julkaisupaikka Helsinki **Painopaikka** Helsinki **Vuosi** 2014

Sivumäärä 193 **urn** http://urn.fi/URN:ISBN:978-952-60-5693-7

Preface

The publications included here were written collateral to my work on four projects of Nokia Research Center, Finland:

1999-2000: BillNeat. The project addressed the problem of charging and billing for Internet services.

2001-2006: Global Authentication Infrastructure (GAIN). The project developed a generic way to reuse the existing cellular authentication infrastructure for authenticating applications towards service providers. The technology was standardized in 3GPP as Generic Authentication Architecture (GAA) and I have contributed to that standard. (See publication II in this dissertation).

2003-2007: Magic Wand (renamed "Digital Living Easy Setup and Security" in 2006). The project was about initializing security (pairing) between mobile phones and other personal devices. The protocols for device pairing designed in this project were eventually incorporated into standards (Bluetooth Secure Simple Pairing and Wireless USB Association Models).

2008-2012: Awarenet. The project developed a low-power radio technology for ad hoc networking between personal devices based on WLAN hardware. Awarenet supports discovering friends, services and data in the local neighborhood of the personal device. The technology is now being standardized by the WiFi Alliance. The "Equal Subdomains" technique described in publication IV of this dissertation is now part of the Nokia Awarenet platform [1, 2].

Each project was done by a team of dedicated people, and I have experienced friendly atmosphere and entrepreneurial spirit in those teams.

Acknowledgments. I would like to thank emeritus Professor Jorma Virtamo for rekindling my interest in doing a PhD; my supervisor, Professor Jörg Ott: this dissertation would have been relinquished without his help; the pre-examiners Stephen Farrell and Giovanni Neglia, whose comments helped to improve this text; and my coauthors, especially N. Asokan and Valtteri Niemi.

In addition to the above-mentioned projects, my writing has been supported

by a grant from Nokia Foundation; by TEKES as part of the Future Internet program of TIVIT (Finnish Strategic Centre for Science, Technology and Innovation in the field of ICT); and by the European Community's Seventh Framework Programme under grant agreement no. 258414 (SCAMPI).

Otaniemi, Espoo, May 7, 2014,

Philip Ginzboorg

Contents

List of Publications

This thesis consists of an overview and of the following publications which are referred to in the text by their Roman numerals.

I N. Asokan and P. Ginzboorg, "Key Agreement in Ad-Hoc Networks," *Computer Communications*, Volume 23, Issue 17, pages 1627-1637, November, 2000.

II P. Laitinen, P. Ginzboorg, N., Asokan, S. Holtmanns and V. Niemi, "Extending cellular authentication as a service," in *1st IEE International Conference on Commercialising Technology and Innovation*, London, UK, pages 30-35, September 14-15, 2005.

III N. Asokan, K. Kostiainen, P. Ginzboorg, J. Ott and C. Luo, "Applicability of Identity-Based Cryptography for Disruption-Tolerant Networking," in *Proceedings of the 1st international MobiSys workshop on Mobile opportunistic networking ACM/SIGMOBILE MobiOpp 2007*, San Juan, Puerto Rico, pages 52-56, June 11, 2007.

IV J. Solis, N. Asokan, K. Kostiainen, P. Ginzboorg, J. Ott, "Controlling resource hogs in mobile delay-tolerant networks," *Computer Communications*, Volume 33, Issue 1, pages 2-10, January, 2010.

V J. Solis, P. Ginzboorg, N. Asokan, and J. Ott, "Best-Effort Authentication for Opportunistic Networks," in *The 2011 International Workshop on Hot Topics on Wireless Network Security and Privacy (HotWiSec 2011, co-located with IPCCC 2011)*, Orlando, Florida, USA, pages 1-6, November 19, 2010.

VI L. Bindschaedler, M. Jadliwala, I. Bilogrevic, I. Aad, P. Ginzboorg, V. Niemi and J-P. Hubaux, "Track Me If You Can: On the Effectiveness of Context-based Identifier Changes in Deployed Mobile Networks," in *Proceedings of the 19th Annual Network & Distributed System Security Symposium (NDSS 2012)*, San Diego, California, USA, February 5-8, 2012.

VII P. Ginzboorg, V. Niemi and J. Ott, "Message fragmentation for disrupted links," in *Proceedings of European Wireless 2011*, Vienna, Austria, pages 415-424, April 27-29, 2011.

VIII P. Ginzboorg, V. Niemi and J. Ott, "Message fragmentation algorithms for disrupted links," *Computer Communications*, volume 36, issue 3, pages 279-290, February, (2013) `http://dx.doi.org/10.1016/j.comcom.2012.10.001`.

IX P. Ginzboorg, V. Niemi and J. Ott, "Message fragmentation for a chain of disrupted links," in *The 6th IEEE WoWMoM workshop on Authonomic and Opportunistic Communications (AOC 2012)*, San Francisco, California, USA, June 25, 2012.

Author's Contribution

Publication I: "Key Agreement in Ad-Hoc Networks"

Authors: N. Asokan and P. Ginzboorg.
The paper considers how a group of people can initialize security between their devices without supporting infrastructure. When N. Asokan has suggested the topic I've got interested immediately, because at that time I have been grappling with how to charge consumers for unreliable services [3]. This kind of charging needs security associations between the parties and it was natural to turn to the question of how to initialize security. The paper was developed in a series of joint discussions to which both authors have contributed equally.

Publication II: "Extending cellular authentication as a service"

Authors: P. Laitinen, P. Ginzboorg, N., Asokan, S. Holtmanns and V. Niemi.
The 3rd Generation Partnership Project (3GPP) is the organization that standardizes cellular networks and this paper outlines the Generic Authentication Architecture (GAA) standard of 3GPP. This short paper is an outcome of massive work: During the critical three years of standardization, Nokia alone wrote more than 200 contributions that were presented in the 3GPP meetings. I was an originator and the leader of the Generic Authentication Infrastructure (Gain) project in Nokia Research Center, which resulted in the GAA 3GPP standard described in this paper. The project team got the Nokia quality award in the research category in 2005. When the challenged network nodes are mobile phones, the security associations between them could be created using the GAA standard.

Publication III: "Applicability of Identity-Based Cryptography for Disruption-Tolerant Networking"

Authors: N. Asokan, K. Kostiainen, P. Ginzboorg, J. Ott and C. Luo.
The opinion that identity-based cryptography (IBC) is a rather good basis for Delay Tolerant Network (DTN) security was getting traction in 2006 [4], and this publication is our attempt to understand how useful IBC would really be. We do this in the first part of the paper by comparing the IBC solution with the one that uses traditional cryptography. In the second part of the paper we explain how, when the challenged network nodes are mobile phones, the security associations between them could be created using the GAA standard, described in publication II. I have contributed to the formulation of the research question, the idea to compare the identity-based cryptography with traditional cryptography solution and to the analysis of the results.

Publication IV: "Controlling resource hogs in mobile delay-tolerant networks"

Authors: J. Solis, N. Asokan, K. Kostiainen, P. Ginzboorg, J. Ott.
John Solis, who was then with University of California, Irvine, has visited Finland and worked with our group in Nokia Research Center in 2008 and 2009. These internships resulted in two joint publications and this is the first of them. It is about techniques for congested memory management in a mobile delay-tolerant network devices. I have contributed to the formulation of the research question and to the analysis/interpretation of the results in the paper.

Publication V: "Best-Effort Authentication for Opportunistic Networks"

Authors: J. Solis, P. Ginzboorg, N. Asokan, and J. Ott.
This is the second joint publication with John Solis. We define "best effort" authentication as authentication with zero probability of false negative, but large (for instance 0.5) probability of false positive. This definition is inspired by how the authorities spot-check passengers and goods in the airports and border crossings. We show that best-effort authentication may be sufficient for authenticating message fragments by intermediate nodes in a mobile opportunistic network. I have raised the research question and contributed to the solution methods (the spot-checking and the Bloom filter) as well as to the analysis of

the simulation results.

Publication VI: "Track Me If You Can: On the Effectiveness of Context-based Identifier Changes in Deployed Mobile Networks"

Authors: L. Bindschaedler, M. Jadliwala, I. Bilogrevic, I. Aad, P. Ginzboorg, V. Niemi and J-P. Hubaux.
This paper is about location privacy in a mobile ad hoc network. It is based on a four months-long experiment on EPFL campus during the spring semester of 2011: A passive attacker having 37 access points in an area of 66×186 m^2 tracked the location of 80 mobile ad hoc network nodes in that area. The mobile devices, in turn, tried to hinder tracking by changing their pseudonyms (network identifiers) from time to time. They were actually Nokia mobile phones with Nokia Instant Community (now called Awarenet) software [1, 2] in the hands of EPFL students. Laurent Bindschaedler and Murtuza Jadliwala are the primary authors of this paper. I contributed to the formulation of the problem, design of the pseudonym change algorithm in mobile devices, coordination of the trial, and analysis/interpretation of the experimental results.

Publication VII: "Message fragmentation for disrupted links"

Authors: P. Ginzboorg, V. Niemi and J. Ott.
In a challenged network with unreliable transmission links the contact times (i.e. the times when the link between two nodes is in the ON state) can be a very scarce resource. Allowing messages to be fragmented on their way to the destination may utilize these contact times better. This is the first in the series of three publications (VII, VIII and IX) that deal with fragmented message transmission over disruptive links. The topic was suggested by Jörg Ott, who after doing an extensive simulation study of fragmentation methods in DTN [5] thought that also theoretical understanding of the phenomena will be useful. In this paper we develop basic theory of fragmented message transmission over a single disruptive link. The research theme evolved during a series of meetings between the authors to which they contributed equally. I organized these meetings and also wrote most of the text after our discussions. Parts of the Introduction section in this publication are by Jörg Ott. The idea to estimate the mean transmission time by equation (V.4) in the paper comes from Valtteri Niemi. He also drafted the inductive proof of equation (V.20) in Appendix B.

Publication VIII: "Message fragmentation algorithms for disrupted links"

Authors: P. Ginzboorg, V. Niemi and J. Ott.
This is the second publication concerning message fragmentation in this thesis. It is about algorithms for choosing a fragmentation block size f when preparing messages for transmission over a single disrupted link. We divide the algorithms into three classes and test representatives from each class in a large number of simulated environments. As with the other two papers, I wrote most of this text after discussions between the authors. All simulations and graphs were done by Jörg Ott. He also drafted parts of the Introduction and of section 5. The idea of oracle-based algorithm O2 in section 4.1 comes from Valtteri Niemi.

Publication IX: "Message fragmentation for a chain of disrupted links"

Authors: P. Ginzboorg, V. Niemi and J. Ott.
This is the third publication in a series of three about message fragmentation in a challenged network. We derive an approximate formula for the mean transmission time of a single message over a homogeneous chain of dodgy links. As with the other two publications also this paper was developed in joint discussions between the authors. I wrote most of this text after these discussions. The calculations in section IV were done by Valtteri Niemi with my help. Jörg Ott has implemented and run the simulations using which we have verified our theory.

List of Figures

List of Acronyms

2G: Second Generation.

3GPP: Third Generation Partnership Project.

AKA: Authentication and Key Agreement.

BSF: Bootstrapping Function.

CA: Certification Authority.

DH: Diffie-Hellman.

DTN: Delay Tolerant Networking.

DTNRG: DTN Research Group.

EID: End-point Identifier.

EMV: Europay, MasterCard and Visa.

EPFL: École Polytechnique Fédérale de Lausanne.

FIFO: First In First Out.

FQ: Fair Queueing.

GAA: Generic Authentication Architecture.

Gain: Global Authentication Infrastructure.

GBA: Generic Bootstrapping Architecture.

GSM: Groupe Spécial Mobile. (English: Global System for Mobile Communications.)

HTTP: Hypertext Transfer Protocol.

IBC: Identity-based Cryptography.

ID: Identifier.

IdP: Identity Provider.

IETF: Internet Engineering Task Force.

IP: Internet Protocol.

IRTF: Internet Research Task Force.

KDF: Key Derivation Function.

KS: Key Server.

LAN: Local Area Network.

LTP: Licklider Transmission Protocol.

MAC: (1) Message Authentication Code; (2) Media Access Control.

MTU: Maximum Transmission Unit.

ONE: Opportunistic Network Environment.

PC: Personal Computer.

PCA: Pseudonym change algorithm.

PKG: Public Key Generator.

PKI: Public Key Infrastructure.

PP: Public Parameters.

RFC: Request For Comments.

RTT: Round-Trip Time.

SCTP: Stream Control Transmission Protocol.

SIP: Session Initiation Protocol.

SSL: Secure Sockets Layer.

TCP: Transmission Control Protocol.

TLS: Transport Layer Security.

UDP: User Datagram Protocol.

USB: Universal Serial Bus.

WFQ: Weighted Fair Queueing.

WIM: Wireless Identification Module.

WLAN: Wireless Local Area Network.

WPKI: Wireless Public Key Infrastructure.

1. Introduction

This dissertation is mainly about securing data communication between mobile devices with little or no support from infrastructure. The adjective phrase "partially isolated" will be used to describe networks of nodes that communicate in such manner. Our thesis is that in partially isolated networks security and privacy, on the one hand, and resource management, on the other, are closely related and tend to influence each other. For example, security needs impact on resource management whenever we allocate scarce resources to mitigate potential attacks on the system. As another example, access to scarce resources needs to be controlled, which implies some kind of authorization for devices that consume them.

Communication without full infrastructure support is a topic of extensive research. The reasons for this scientific interest are the new possibilities for communication in special conditions and the intellectual challenge in realizing these possibilities.

New systems for wireless device-to-device communications for consumers with limited infrastructure support have been recently announced by Nokia [2], Qualcomm [6], NEC [7], and Peep Wireless [8]. Even though business in that area has not taken up so far, these developments indicate industry's interest. In this kind of systems we need to secure wireless data communication between consumer devices because it may be visible to the public.

In some scenarios the whole network could reside in a physically separated, closed environment with controlled access, like the chromium mine in Kemi, Finland [9]. On these grounds it may be argued that there is no need for additional security measures, like message authentication, inside the mine. But the recent sequence of cyber-attacks on infrastructure in different countries indicate that additional security measures should be at least considered for precisely this kind of environments.

Here is the list of well known incidents [10–12]: In 2010 Stuxnet malware

was detected in Iranian industry sites. It is possible that the malware have initially penetrated the closely-guarded industrial sites on a contaminated USB memory stick. The purpose of this program was to disrupt the Iranian work on enrichment of uranium. Stuxnet have spread also to Russia, India and Indonesia. In 2011 Duqu spyware, that seems to be a relative to Stuxnet based on code analysis, was detected. Duqu gathered information from companies that specialize in industrial automation. It attempts to send the collected data to a remote server and removes itself from the contaminated host after 36 days. In 2012 the Flame spyware was used to gather information from Middle-Eastern data networks. Also in 2012, another spyware Rocra was found. It has been gathering data in tens of countries since 2007.

It seems that the cyber-attacks were planned and prepared in a well organized way. Now, imagine a lab of, say, 30 people somewhere deep inside a government bureaucracy, or in the private sector, who write this kind of software for a living. After delivering last year's successful malware the fellows are probably busy working on the next one. (I cannot think of a reason why such an organization—once it exists and have proven its value by producing Flame or Rocra—will be dissolved.) So the sequence of cyber-attacks is likely to continue and there is a place for data security measures inside mines and other industrial systems, even though these are physically separated, closed environments.

Before going on, let us define the following "special" networks:

(a) a *challenged* network is a network subject to difficult operational constraints, like disrupted links and high delays;

(b) a *delay-tolerant* network (DTN) does not require for its operation (i) small Round-Trip Time (RTT), or (ii) simultaneous end-to-end paths, or (iii) continuous connectivity between nodes;

(c) a *mobile opportunistic* network is a DTN in which some nodes move and contact each other in a way that cannot be predicted precisely;

(d) the devices in an *ad hoc* network communicate without infrastructure support;

(e) a *mobile ad hoc* network consists of mobile devices.

Please note that a DTN is a challenged network by these definitions, while an ad hoc network is not necessarily so.

Potential applications of challenged networks are in the areas were infrastructure needed for good end-to-end connectivity is difficult or inconvenient to deploy. Those areas include communications in industrial environments (mines, factories, shipyards), space, military, emerging markets, and local opportunistic communication between mobile phones, like in Nokia Awarenet [1, 2].

Communication in challenged networks is characterized by (i) intermittent connectivity of links, and (ii) scanty resources in nodes. Those properties cause big variation of communication delay and unstable paths between sender and receiver. As a result, standard networking and security techniques that assume the opposite of (i) or (ii) may be difficult to apply in challenged networks.

Delay Tolerant Networking

Ways of communicating in challenged networks are developed and documented by the research community. We will now outline for reference parts of the work done in the Delay Tolerant Networking Research Group (DTNRG) [13] of Internet Research Task Force (IRTF), because it is encompassing and well-documented. The DTNRG technology is used in space communications, the academy, and several start-up companies.[1] Please note that our results are not restricted only to these challenged networks that follow the DTNRG specifications.

The DTNRG material includes architectural documents and protocol specifications. The architecture [17] is designed so that messages are transmitted from one node to another as the transmission opportunity arises, because end-to-end path across the whole network may not exist.

DTN literature includes many articles and two books: the monograph by S. Farrell and V. Cahill [18], and a collection of papers, edited by A. Vasilakos, Y. Zhang and T. V. Spyropoulos [19]. These books and the surveys [20–23] are good starting points into the DTN literature.

The following notions developed in DTNRG are relevant to us.

Bundle protocol. The bundle protocol [24, 25] offers transport services for applications and allows creation of self-contained messages—called "bundles"—that enable complete application interactions by a single message exchange. The self-contained messages (bundles) may be much larger than IP packets. The messages are transmitted like e-mail: link-by-link in a store-carry-and-forward fashion. Also the responsibility for moving the message to its destination is transferred link-by-link with the message. The protocol includes an option for successful delivery acknowledgment message sent back from the destination to the source node. Bundle has a lifetime, after which it expires and is removed from the network.

Convergence layer. This is the method used for transmitting a bundle over a single link. Convergence layer is immediately below the bundle protocol in

[1] Start-up companies that use (or have used) DTNRG technology include First Miles Solutions [14], Uepaa [15] and Tolerant Networks [16].

the protocol stack. It offers to the bundle protocol reliable transmission of data between neighboring nodes, including indication when the transmission has stopped. Please note that what looks to the bundle protocol as a single link, may be actually a sequence of physical links, hidden under the hood of the convergence layer. The convergence layer can be constructed, for example, from a transport layer protocol, like TCP, SCTP and UDP.[2] Another example is the Licklider Transmission Protocol (LTP) [26, 27], designed for point-to-point communication over a link with very high latency and disruptions.

Naming conventions. DTN nodes are identified by End-point Identifiers (EIDs). Typically an EID refers to a single node. It can also refer to a group of nodes, which is useful for multicasting. The EIDs are at most 1024 bytes long and have the form of Uniform Resource Identifier (URI) [28], i.e. they are similar in structure to web page addresses. The prefix "dtn://" is registered for DTN use. For example, in the DTN of the chromium mine in Kemi, Finlad [9], the EID structure was "dtn://" followed by device's name, like "dtn://server", or "dtn://android-3". But "dtn://" is not the only possible prefix. For instance, "mailto://alice@wonderland.co.uk" has a legitimate EID's form.

Security enablers. The specification supports encryption, digital signatures and time stamps. These techniques enable the basic security services: (1) confidentiality—by encrypting the data with a key that only the intended receivers possess. (2) integrity—the receiver can ensure that no data has been altered in transit; (3) authentication—the receiver can verify that the data really was sent by the claimed sender; (4) freshness—the receiver can verify that data is recent. (The last service relies on time stamps; it needs at least loosely synchronized clocks in network nodes.)

Fragmentation. Two types of fragmentation for DTN are defined: pro-active and reactive. In the former, the source node for the link divides application data into blocks and sends each block in a separate fragment bundle. This is useful, for example, in transmission over satellite links, where the timing of the interruptions is known in advance. In the latter, the data is split only when the transmission between two nodes on any link of the message path is interrupted; resulting in one fragment bundle with data that made it to the receiver and one containing the remainder at the sender. This is useful, for example, in transmission over opportunistically-formed links, where the time of the interruption is not known in advance. The fragmented data is reassembled at its destination, but also an intermediate node can reassemble fragments into

[2]These initialisms abbreviate Transmission Control Protocol, Stream Control Transmission Protocol, and User Datagram Protocol.

a new bundle.

Research questions

In challenged networks the intermittent connectivity between participants, and to the infrastructure services create problems in routing, name resolution, service discovery, and security. Three types of infrastructure services are relevant to us. The first type is the routing infrastructure in the form of fixed routers and stable links between them. The second type is the server infrastructure of on-line servers which provide various services such as name service, directory services, and trusted third party services. The third type is the organizational and administrative support such as registration of users, management of names and network addresses, issuing of certificates, and cross-certification agreements between different user domains.

The DTN architecture and the bundle protocol provide the basic means to communicate without fixed routing infrastructure: like an e-mail, a bundle is a self-contained message forwarded over a chain of intermediate nodes that may each store the bundle for long periods of time. DTN routing, i.e. the question of how the actual chain can be formed, is an area of active research [20, 29–37]. It is, however, not in the scope of our thesis.

Intermediate nodes may have limited memory, computing power and battery charge. In publication IV of this dissertation we study techniques for congested memory management under adversarial scenario. In publication V we show how to reduce the intermediate nodes' computations in an extension of that scenario. Even though reducing the intermediate nodes' computations may increase their battery life, the general question of energy-efficiency in DTN is not in the scope of this dissertation.

If no contact between two nodes is sufficiently long to forward the entire bundle, then there must be some way to fragment the data and forward it one piece at a time. For that reason the bundle protocol includes fragmentation techniques. Securing the bundle pieces implies that the possible boundaries at which the data may be split must be defined by the sending node before transmission. This leads to the question of how to define these fragmentation boundaries at the sender. That question is the topic of our publications VII, VIII and IX.

Also the lack of, or intermittent access to services and to administrative support must be addressed when building an actual network. The issues that need to be solved include initialization of security and more generally key management. We address mainly the security initialization question in publications I,

II and III.

Yet another type of concern is raised by intermediate nodes (or an attacker) peeking into messages of others. To some extent privacy can be addressed by encrypting the message content and thus making it confidential. But the clear-text bundle header still provides the means to track the node.[3] This issue is the subject of our publication VI.

Structure of the dissertation

This dissertation contains nine publications. Their summaries are grouped in two chapters. (A summary of a publication will usually end with some discussion and further analysis points.) The first, Chapter 2, includes six publications dealing with security topics. The second, Chapter 3, includes three publications about fundamentals of fragmented message transmission over disrupted links. This topic is in the area of managing the contact times in challenged networks, but our choice of fragmentation model is motivated by security considerations. The methods used in those publications include thought experiments, architectural studies, mathematical modeling and experiments in simulated as well as in actual environments.

We will now outline the contents of the two chapters. Please note that both chapters also start with comprehensive introductory parts.

Chapter 2. Intermittent access to, or complete lack of infrastructure and partial isolation of nodes from each other introduce security problems that may differ substantially from security problems in networks with full support infrastructure and good connectivity. This chapter deals with a subset of the security problems in partially isolated environments. In particular, we are interested in mobile opportunistic or ad hoc networks of hand-held devices, like mobile phones, that communicate with each other over short-range radio. The topics include initialization of security, authentication techniques, managing congested memory in adversarial scenario, and location privacy. Other security-related topics are not in the scope of this dissertation.[4]

[3] Encrypting also the header to improve privacy will make the routing operation rather heavy due to increased amount of routing-related computations in the intermediate nodes. For this reason we prefer to send the header unencrypted.

[4] Here are three examples of interesting but out-of-scope topics: risk management in automatic interactions between mobile devices over short-range radio connections [38], key updates in disconnected environment (as distinguished from updates during occasional access to the infrastructure), and credit-based incentive schemes for message forwarding [19, ch. 4].

Chapter 3. In a challenged network with unreliable transmission links the connection between the sender and the receiver may be cut before the entire message has been transmitted. For that reason the contact times (i.e. the times when the link between two nodes is in the ON state) can be a very scarce resource. Allowing messages to be fragmented on their way to the destination may help to use these contact times better. In IP networks the maximum size of an IP packet that can be transmitted without fragmentation is typically determined by the path probing technique of RFC 1191 [39]. But it is hard to apply this technique in a challenged network where simultaneous end-to-end path from source to destination is unlikely.

The three publications VII, VIII and IX summarized in this chapter are dedicated to fragmented message transmission over disrupted links. We study message fragmentation in the simplest case of a single disrupted link over which the message needs to be delivered in publications VII and VIII: Methods to estimate the mean transmission time of a fragmented message within a basic system model are developed in publication VII; publication VIII deals with fragmentation algorithms over a single link. In publication IX we study transmission times of fragmented message over multiple disrupted links arranged in a chain. The chapter ends with an overview of our results in section 3.4.

2. Topics in security and privacy

In a typical scenario a person signs a contract on paper that creates (via law and society) mutual obligations and rights between parties. With the help of security techniques those rights and obligations can be transferred into the digital domain and then used to support new transactions. But no matter how impressive those digital credentials are, entering into a transaction with another party still requires trust.

Let us define "trust assumption" of party A about B as the willingness of A to deal with B, even though the subsequent transaction is risky for A. This willingness may increase with the number of successfully completed transactions between the parties.[1]

Security techniques can only help in transforming and transferring already-existing trust assumptions. By themselves they cannot create trust.

For example, with so-called key continuity mechanisms [41], once party A have got a "known-good" key, it can verify a remote entity's identity by verifying that they're still using the same key. Here, the usage of the same key is the security mechanism, while "known-good" means that the past transactions with the party identified by that key have been successful. It is those past successes that may increase A's trust assumption. Indeed, if the party (identified by its usage of the same key) has done harm to A in the past, then A's trust in it is likely to decrease.

Because security techniques by themselves cannot create trust, initialization of security in completely or partially disconnected environments may be one of the thorniest issues. We address this question in publications I and II, and to some extent also in III. In publication I we introduce a scenario where the security is initialized from the common "written on a wall" password that

[1] Please note that while this informal definition should be sufficient for our purposes, an in-depth treatment of trust is not in the scope of this dissertation. One starting point into the topic of trust is O. O'Neill's book [40], where she provides a philosopher's view of trust and deception.

the users in the same location see. This publication is one of the first to define that sort of scenario, where the security requirement is expressed in terms of location. (The other one is Stajano and R. Anderson's paper [42]. In their scenario a newly-initialized device creates a security association over a short-range radio with (only) the first nearby device that it encounters. For example, the newly-initialized device could be a headset and the other device a mobile phone.)

In publication II the we describe the 3GPP Generic Authentication Architecture (GAA) standard. It enables security initialization from the cellular authentication and key agreement. With GAA it is possible to bootstrap ad hoc network security from previously set up associations with the cellular network, provided that (i) the participants have trust in the cellular network operator; and (ii) the devices are (or can be temporarily connected to) mobile phones. The creation of security associations between challenged network devices using GAA is illustrated in the second part of publication III.

Identity-based cryptography (IBC) [43] is a relatively new cryptographic method that enables message encryption and signature verification using the public identifier of the target as a key. This could be an advantage in partially disconnected environments, because the message sender or signature verifier does not need a separate (from the identifier) public key of the target. For example, target's e-mail address could be both his public identifier and his public key. In publication III we compare two security architectures for mobile opportunistic network: the Identity-Based Cryptography security solution, and the solution based on traditional Public Key Infrastructure (PKI). We conclude that (1) IBC has no significant advantage for authentication and integrity compared to traditional cryptography; and (2) IBC can enable better ways of providing confidentiality. But traditional cryptography has more support (solid implementations and installations) than IBC, simply because it has been around for a longer time than IBC. On balance, it is not clear if (2) in itself would be a sufficient reason for preferring IBC in actual DTN implementations.

To cope with temporary isolation of the network nodes the store-carry-and-forward communication method in DTN keeps a forwarded message in an intermediate node until that message is forwarded again, or it is no longer needed. Congestion occurs when there is not enough memory to store messages in intermediate node. In publications IV and V we turn to congested memory management in DTN. We approach this question via an adversarial scenario and test our solutions in simulated environment. The idea underlying our solutions is first, that the management of intermediate node's memory is preemptive: If

there is enough space to accommodate an incoming message, then that message is accepted without fuss; otherwise, we need to make room by dropping either the incoming message, or some of the stored messages. Second, during congestion an intermediate node should allocate its own memory based on who has sent the messages competing for it. Therefore, intermediate nodes should be able to recognize who has sent each message. The "Equal Subdomains" technique found in publication IV is now part of the Nokia Awarenet platform.

The problem we are addressing in publication V is in the area of traffic authorization in DTN: adding message fragmentation to use the contact times better, may substantially increase the amount of computation in intermediary devices, because they will have to authenticate lots of fragments. The solution is inspired, in part, by spot-checking of passengers and goods in the airports. We define "best effort" authentication as authentication with zero probability of false negative, but large (for instance 0.5) probability of false positive,[2] and show that this technique may be good enough for authenticating message fragments by intermediate nodes in our network.

In publication VI, the last to be summarized in this chapter, we address the location privacy issue. In an infrastructure-based network the node's location is typically known in the infrastructure [44]. Even though this is not the case in the infrastructure-less wireless networks, the location of a device can still be determined from its wireless transmissions. In particular, eavesdroppers to short-range wireless communication between mobile personal devices can track devices (and thus people who carry those devices) based on device identifiers in the broadcasted wireless messages. The question is how much the use of transitory device identifiers (pseudonyms) in place of fixed device identifiers helps to evade tracking.

We have approached this question experimentally: During a four-months long experiment in EPFL campus, an attacker having 37 access points in an area of 186 by 66 m^2 tracked daily the location of 80 mobile ad hoc network nodes. The nodes, in turn, tried to evade tracking by changing their network identifiers from time to time. Our experiments indicate that the actual attacker consisting of 37 sniffing stations covering an area of approximately 200×100 m^2 on the EPFL campus can track devices in that area with high probability, even though they employ state-of-the-art techniques; and that reducing the number of sniffing stations hinders tracking. This seems to be the first field-study that evaluated

[2]Here, "false negative" is the event where authentication of a legitimate message fails, and "false positive" is the event where authentication of an illegitimate (attacker's) message succeeds.

context-based identifier-change mechanisms under a practical adversary model with real mobile devices and communication scenarios.

2.1 Key agreement in ad hoc networks: publication I

Consider the following problem in the area of security initialization: a group of people in a meeting room do not have access to public key infrastructure or third party key management service, and they do not share any other prior electronic context. How can they set up a secure session among their devices?

The solution we investigate in this publication is choosing a fresh password and sharing it among those present in the room (e.g., by writing it on a blackboard). If this password is a sufficiently long random string, it can be used directly to set up a security association. In practice, people find it difficult to use long random strings. It is much more user-friendly if the password is a string that people recognize easily, such as a natural language phrase. But natural language phrases are weak as secrets because they are drawn from a rather limited set of possibilities; they are susceptible to dictionary attacks where an attacker records the encrypted traffic and then attempts trial decryptions with candidate passwords until he finds the correct one. Therefore, we need a key agreement protocol to derive a strong shared session key from the weak shared password. We assume an attacker who can insert messages but cannot modify or delete messages sent by others.

In the first part of the publication we examine various alternatives and propose new protocols for password-based multi-party key agreement in this scenario. We start with the generic protocol for password authenticated key exchange of Bellovin and Merrit [45], that combines symmetric and asymmetric cryptography, and examine how it can be extended to multiple parties. Then we look at how password authentication can be added to multi-party Diffie–Hellman (DH) key exchange protocols.

In the second part of the publication we present a fault-tolerant version of a multi-party Diffie–Hellman (DH) key agreement protocol. Becker and Wille showed in [46] how multi-party DH key exchange can be done efficiently, in terms of the number of communication rounds, by arranging the players (nodes) in a d dimensional hypercube. We describe in section 3.3.2 of the publication how authentication based on a single shared password can be incorporated into this protocol.

Then we turn to situations where a player finds that its chosen partner in a protocol round cannot be authenticated with the shared password. Such partner

either does not know the password, or is unreachable for some reason. In these situations the player needs an algorithm to select other potential partners until a non-faulty one is found. An example algorithm for doing this is described in sections 3.3.3 to 3.3.5 of the publication.

Discussion

The network in our scenario is an example of an ad hoc network in which entities construct a communication network with little or no support infrastructure. Lack of infrastructure introduces problems in routing, name resolution, service discovery, and security that are typically not encountered in networks with full support infrastructure.

In the case of security, there is a second source of difficulty. For example, consider the usual form of confidentiality requirement. Participant A requires the following: "only L can read the messages I send", where L is a label in some name space that is meaningful to all participants. A trusts a certification authority CA to correctly certify public encryption keys of entities. Participant A then achieves its confidentiality requirement by encrypting its messages using a public key for L certified by CA. This process started from a certain prior context, consisting of the well defined name space, A's confidentiality requirement, and A's trust in CA.

In providing security services, one always starts from such a context. (Security techniques can only help in transforming and transferring the trust assumptions in the prior context. They cannot create trust.) But in our scenario, we encounter a new type of prior context. The security requirement is expressed in terms of location: "Only people present in this meeting room can read the messages I send."

There are obvious solutions to the problem if additional assumptions are made. Two examples are mentioned in the publication: We can use a trusted third party capable of determining the locations of players. Such support may be available in some ad hoc networks. But it is too strong an assumption in general. Another solution is to use a physically secure channel limited to those present in the room (like an infrared link, or a wired connection that we mention in the publication) to initiate the negotiation of the session key before switching to the insecure wireless channel.

Please note that the location-limited channel (also called "out of band", or "side channel" in the literature) can be implemented with a large variety of options. The channel can be can be one- or bi-directional; clear-text, or confidential; and its usability depends on the actual devices involved. Moreover, in addition to (or

instead of) initiating the negotiation before switching to the insecure wireless channel, the side channel can be used to verify the result of that negotiation.

Key agreement in ad hoc scenarios is both important and tricky. It is important, because personal mobile devices have short-range radio interfaces (like Bluetooth and WLAN) that need to be secured; it is tricky due to the user's limitations (including the difficulties of correctly entering long passwords into the device; and in general of doing a sequence of actions as prescribed by the protocol designer) and the variety of options. For these reasons key agreement in ad hoc scenarios became an active research field. F. Stajano and R. Anderson [42] and the workshop version of our publication I seem to be the first papers to tackle the topic. Here is a sample from the many papers that were later published in that area: E. Uzun, K. Karvonen, and N. Asokan [47], M. Cagalj, N. Saxena, and E. Uzun [48] (usability); N. Asokan and K. Nyberg [49] (protocols); M. Rohs and B. Gfeller [50], N. Saxena, J.-E. Ekberg, K. Kostiainen, and N. Asokan [51] (visual side channel); W. Claycomb and D. Shin [52] (audio side channel); R. Mayrhofer and H. Gellersen [53] (accelerometer as side channel).

I will end the discussion with a pointer to an interesting and relatively recent technique. The basic idea behind it is that an active attacker can easily add energy to the radio channel. But it is very difficult for an attacker to remove energy from the radio channel. So if the parties know that a transmitted message consists of a fixed and equal number of "zeros" (silence periods) and "ones" (activity periods) that the sender randomly mixes before transmission, then attacker's attempts to change the content of that message will be almost surely detected by the receiver. In this way transmission of critical data over the insecure wireless channel is made tamper-evident in the key agreement methods of S. Capkun et al. [54] and S. Gollakota et al. [55].

2.2 Initializing security from the cellular network: publication II

Rather than creating security associations between mobile devices from scratch as is done in publication I, we can bootstrap ad hoc network security from previously set up security associations. The latter could be the security associations with the cellular network, provided that (i) the participants have trust in the cellular network operator; and (ii) the devices are (or can be temporarily connected to) mobile phones.

The 3rd Generation Partnership Project (3GPP) is the organization that standardizes cellular networks and this publication outlines the Generic Authentication Architecture (GAA) standard of 3GPP, which enables cellular operators to

extend cellular authentication to new services. Before GAA, the cellular authentication's reuse was specified separately for each new 3GPP service. Outside of 3GPP, numerous authentication contraptions have been devised based on exchanging short messages between mobile phone and service provider, because such exchange requires cellular authentication. But the short message service was not meant for authentication, for example, because timeliness of message delivery is not guaranteed.

The industry reasoned as follows: Specifying how to reuse cellular authentication separately for each new 3GPP service is not sensible. Defining, instead, a generic authentication method for services offered by the cellular network itself, would, in the long run, save standardization and implementation costs.

The GAA standard

Cellular authentication is based on a master key that is shared between subscriber's smart card application in the mobile phone and her cellular operator's authentication center. This master key is used to authenticate the mobile phone and set up a shared session key(s) between the phone and the network. The phone and the 3G network are mutually authenticated[3] and two session keys are created. These keys are then used to protect the wireless communication between the phone and the network.

To this GAA adds two main procedures. (1) During bootstrapping, GAA credentials are bootstrapped from the cellular authentication and key agreement (AKA), in the sense, that they are derived from the security data of AKA. (2) Then the GAA credentials are used to secure a connection between a client and a server of some network application. These two procedures are enabled by a new bootstrapping function, BSF, in the cellular network. The BSF has four network interfaces which, as is customary in 3GPP standards, are identified with two-letter labels: Ua, Ub, Zh and Zn [56]. The architecture is illustrated in Figure 2.1.

The bootstrapped GAA credentials can be used to authenticate a certificate enrollment request. This variant of the second procedure enables the cellular authentication to bootstrap an operator-specific public key infrastructure (PKI).

[3]We remark that while the mobile phone and the network are mutually authenticated in the 3G system, in the earlier 2G (GSM) system only the mobile phone authenticates itself towards the network. One-sided authentication of 2G enables so-called "false base station" attack, where the attacker masquerades as a legitimate base station towards the mobile phone. To counter this attack, authentication of the network towards mobile phone was added into the authentication and key agreement (AKA) procedure during the design of 3G security.

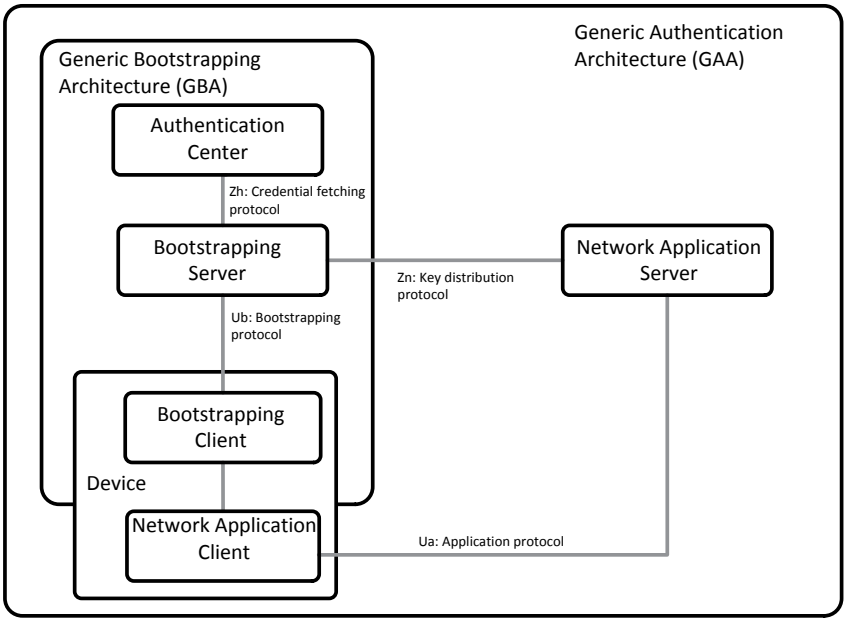

Figure 2.1. Schematic illustration of the 3GPP Generic Authentication Architecture (GAA). This is a general framework that allows the cellular authentication infrastructure used in authorizing subscribers' access to the cellular network to be used in authorizing access to new services. Those services can be provided by the cellular network operators or by third parties that have a business agreement with the network operator. The bootstrapping server is a new function in the cellular network architecture, which enables reuse of cellular authentication in authorizing access to new services.

This ends our summary of publication II. We will return to the bootstrapping theme in the summary of publication III on page 36.

Even though the bootstrapping procedure is simple, there are organizational and technical subtleties in combining existing and new security protocols in one standard authentication and key agreement package. One of the technical subtleties is the possibility of Man-in-the-Middle attack when a legacy client authentication protocol is run within a secure tunnel. This vulnerability arises if the legacy client authentication protocol is used both in tunneled and untunneled forms. The attacker tricks a benign client into authenticating to a fake server, which forwards the authentication protocol messages to a benign server, tricking the latter to authenticate the attacker (now acting as a client) with these messages. The discovery of this attack in the course of GAA work impacted several IETF protocols [57].

Subsequent developments

I shall outline here several GAA-related developments and activities that happened subsequent to publication II.

• *Cellular industry.* The development of the GAA in 3GPP continues. Two notable additions in the standard [56] are the possibility to bootstrap from the older, 2G (i.e. GSM) credentials, and the use of bootstrapped credentials in Session Initiation Protocol (SIP).

• *IETF*. We have already mentioned the discovery during GAA work of the Main-in-the-Middle attack in tunneled authentication protocols. This discovery helped to improve several IETF protocols [57].

Authentication with keys bootstrapped from the cellular AKA in the web browser requires a version of Transport Layer Security (TLS) protocol that works with pre-shared keys, and so the standardization of this TLS feature was initiated in IETF. The pre-shared TLS has been standardized in IETF RFC 4279 [58], and now lives its own, independent from GAA, life.

• *GAA Book*. In September 2008 Wiley has published our book "Cellular Authentication for Mobile and Internet Services" [59]. The book brings together the essential information about GAA in one volume.

• *Academia*. C. Chen's stay in the Royal Holloway University of London have resulted in several GAA-related papers by C. Chen, S. Tan and C. J. Mitchell. In [60] they propose to generate a one-time-password from the bootstrapped key and the long-term password of the user. Subsequent authentication with this one-time-password enjoys better security compared to authenticating with the long-term user's password only. In [61], they propose to make the SSL/TLS connection more secure against man-in-the-middle attack by the server authenticating the terminal with a temporary key derived from the SSL/TLS session identifier and the bootstrapped key. In [60, 62] and [63], two potential alternatives to the cellular security infrastructure in GAA are proposed; these are, respectively, the Trusted Computing [62] and the EMV cards (of MasterCard and Visa) infrastructures.

2.3 Applicability of identity-based cryptography: publication III

Identity-based cryptography (IBC) [43] is a relatively new cryptographic method that enables message encryption and signature verification using the public identifier, such as e-mail address, of the target as a key. An IBC system consists of users (e.g., message senders and recipients) and a commonly trusted third

party called the Private Key Generator (PKG).

A. Seth, U. Hengartner and S. Keshav have proposed in [4] a security architecture for DTNs based on IBC. They argue that traditional Public Key Infrastructure (PKI) is not well suited for disconnected environments, such as DTNs, since access to online servers for fetching public keys and checking certificate revocation lists cannot be assumed.

We've got interested in this topic because DTN with IBC would have been more difficult to include in Nokia mobile phones compared to DTN with traditional PKI: When this publication was written Nokia phones have had a traditional crypto-library for years, but the IBC primitives, though implemented, were less mature.

We examine the proposed IBC security architecture for DTN in the first part of the publication by comparing the IBC solution with the one that uses traditional cryptography.

Before going into our analysis of IBC security architecture for DTN let us outline how the IBC system works.

At system initialization phase the trusted PKG generates system-wide Public Parameters (PP) and a corresponding master secret key S_{PKG}. Using S_{PKG} and a public identifier id_P the PKG can generate a private key S_P for user P. The public identifier id_P could be the user's e-mail address, for example. We illustrate this initialization in Figure 2.2 for two users.

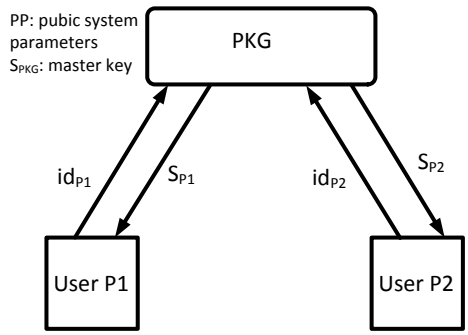

Figure 2.2. Schematic illustration of IBC initialization for two users.

At this point the PKG must verify that the user really is allowed to have this particular id_P, and a confidential communication channel is needed to securely deliver the private key S_P to the user. A message can be encrypted for the user P using PP and id_P. The user P can, in turn, decrypt the resulting ciphertext with the S_P it received from PKG. In similar fashion, P can sign messages with S_P and other users can verify the signature using id_P and system-wide PP.

Analysis of IBC security architecture for DTN

We will now continue with analysis of IBC security architecture for DTN.

Authentication and integrity

Seth et al. argue that certification revocation lists of the traditional PKI are unsuitable for DTNs because updates to these lists can be delayed excessively in a disconnected environment. Instead, they propose the use of IBC. In IBC, revocation is avoided by periodically refreshing the underlying identifiers, and hence the signing keys. Each signing key is valid for a short period (e.g., a day). An underlying identifier is constructed by concatenating the long-lived identifier with a description of the validity period: e.g., "mailto://alice@example.com: 2013-05-03" to refer to the underlying identifier that should be used to encrypt messages for Alice on March 15th, 2013. A verifier can check if the message was signed with a sufficiently recent signing key. Thus, instead of requiring the verifier to receive revocation lists in a timely fashion, IBC-based authentication schemes require the signer to receive fresh signing keys periodically.

But a similar approach can also be used with traditional public key cryptography: the certificates issued to signing keys can be short-lived (e.g., valid for a day). The signer must periodically receive new certificates from the certification authority (CA), but the signing key itself may be long-lived. A verifier can check if the message is correctly signed and is accompanied by a sufficiently recent certificate. Thus we conclude that authentication needs in DTNs can be met without resorting to IBC, but through the use of traditional cryptographic techniques instead. Note that when traditional digital signatures are used for authentication, the sender can compute all the necessary authenticators even when there is no network connectivity.

Adding a certificate (or a chain of certificates verifying all the intermediary CAs from the root CA to the sender) increases the message size by a few kilobytes. However, if messages are relatively large, say, at least hundreds of kilobytes, the overhead introduced by the certificate(s) is not significant.

In sum, IBC has no significant advantage for authentication and integrity compared to traditional cryptography. In both cases, the sender must have been able to receive a message (containing the IBC key or the certificate respectively) from a server (PKG, or CA respectively) sufficiently recently. The receiver can authenticate DTN messages even while disconnected: If the sender is completely disconnected, it can still send a DTN message with a traditional digital signature. Recipients capable of fetching the current certificate for the sender can authenticate this signature.

End-to-end confidentiality

In IBC systems, a sender can encrypt a message for a recipient by just knowing the recipient's identity and common public system parameters. The sender can construct the encryption even when there is no network connectivity. Message decryption requires the recipient to periodically access PKG to get fresh decryption key.[4] This behavior can be imitated with traditional cryptography using a trusted key server (KS) as described in section 3.2.2 of the publication.

First, the sender picks a random symmetric key k and computes a ciphertext c using symmetric encryption. Then he creates an envelope t that holds both the public identity of the receiver and k using traditional public key encryption and public key of KS. He sends envelope t and ciphertext to the receiver. The receiver forwards the envelope t to KS, which decrypts t and constructs a recipient envelope t' that holds k. KS sends t' to the recipient which first recovers k from t' using the key it shares with KS and then uses it to recover the message from the ciphertext.

This design enables the sender to construct the encryption even when there is no network connectivity. Its disadvantage, compared to the IBC system, is first the strict ordering of receiver's operations: (i) receive message(s), (ii) access the trusted server, and (iii) decryption; the recipient needs a connection to the KS *after* he receives a message to decrypt it. In contrast, in IBC system the operation (ii) may happen before (i). Second, this design puts more load on the server, because KS needs to process messages (envelopes) that users send to each other.

Conclusion

To conclude, (1) IBC has no significant advantage for authentication and integrity compared to traditional cryptography; but (2) IBC can enable better ways of providing confidentiality. But traditional cryptography has more support (solid implementations and installations) than IBC, simply because it has been around for a longer time than IBC. All in all, it is not clear if (2) in itself would be a sufficient reason for preferring IBC in actual DTN implementations.

Initializing security with GAA

In the second part of the publication (sections 3.3 and 3.4) we explain how, when the challenged network nodes are mobile phones, the security associations between them could be created using the Generic Authentication Architecture

[4]Please note that the need of the receiver to periodically access the PKG to get fresh decryption key can be problematic in DTN.

(GAA) standard, described in publication II. This is illustrated in Figure 2.3.

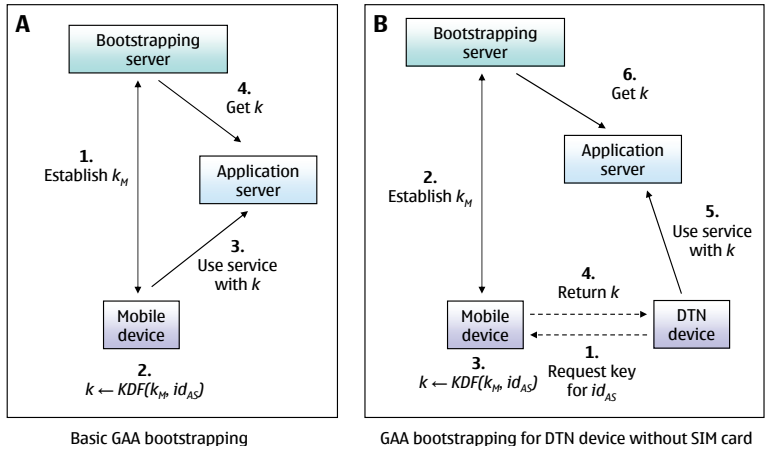

Figure 2.3. Initializing security associations with Generic Authentication Architecture (GAA).

The basic bootstrapping using GAA is illustrated in part A on the left. First, the mobile device and a bootstrapping server of the cellular operator engage in the usual cellular network authentication protocol. As a result, the mobile device and the bootstrapping server share a master session key k_M (step 1). Later, when a mobile device needs a secure connection to an application server, the mobile device can derive with a key derivation function (KDF) a server-specific shared session key k from the master session key and the identifier of the application server (step 2). The mobile device may now use the application server using k (step 3) and the application server may obtain the same key from the bootstrapping server (step 4).

Part B on the right illustrates how creating security credentials in a non-cellular DTN device may happen with GAA. The user first establishes a security association between his non-cellular DTN device and his mobile phone, e.g., by doing Bluetooth pairing. When a DTN device needs a secure connection to an application server, such as CA or PKG, it may send a request to its paired mobile phone. (step 1). The phone may now do the normal cellular authentication (step 2) and derive the application server specific key (step 3). It sends this key back to the DTN device over the short range wireless connection (step 4). The DTN device may now use the service (step 5) and the server may obtain the same key (step 6).

The DTN CA and PKG servers will act as GAA application servers, run either by the cellular operator or independent third parties that have service agreements with the operator. The CA will use GAA to authenticate the enrollment of public keys, and issue short-lived subscriber certificates. The PKG will use

GAA to encrypt IBC private keys for devices. A principal is identified by a well-known identifier (e.g., email address or mobile phone number) which is securely bound to a cellular identifier. Revocation of this identifier will be automatically reflected in the DTN security infrastructure, since the device will no longer be able to receive short-lived certificates or IBC private keys.

To sum up, we show how initialization of security in DTN can be done by bootstrapping it from cellular infrastructure. While the concept of bootstrapping the security of one system from another is well-known [64], bootstrapping DTN security from cellular infrastructure using the 3GPP GAA standard is a contribution of this publication.

2.4 Managing congested memory: publication IV

This publication is about managing congested memory in a mobile opportunistic network where messages are conveyed by intermediaries. We approach this question via an adversarial scenario. The inclusion of an adversary distinguishes our problem setting from other studies of buffer management in DTN, e.g., [65–67].

In our scenario memory and other resources in intermediate nodes can be abused by resource "hogs", i.e., individuals whose message generation rate is much higher than the average. By definition, these nodes are a minority of the population. As an example, a hypothetical leader of resource hogs group may reason like this: "The overall delivery ratio of messages in the network is now 65%. On the one hand, if my hogs send each message twice (disguising the second message so that it seems different from the first), then the network will become more congested. As a result, the overall delivery ratio in the network will drop to, say, 50%. But on the other hand, the chances of our messages delivery will increase: both message copies will be lost with probability $(1-0.5)^2 = 0.25$, and so the delivery ratio that *we* see will be $1-0.25$, or 75%. We will be better off than before."

To generalize this example, suppose that as a result of congestion due to double-sending by hogs the probability of individual message loss in the network increases from p_1 to p_2. The double-sending hogs will experience message loss p_2^2; their strategy will pay off if $p_2^2 < p_1$. And if p_k is the probability of individual message loss when hogs send k copies of each message, then hogs gain advantage (from sending of k copies) if $p_k^k < p_1$.

Solution idea

We assume that each node sets aside a certain amount of memory to buffer the conveyed traffic. The management of that memory is preemptive: If there is enough space in the set aside memory buffer to accommodate an incoming message, then that message is accepted without fuss; otherwise, we need to make room. To do this the node either drops the incoming message, leaving the buffer unchanged, or it deletes some message(s) from those already in the buffer to accommodate the new one.

The idea underlying our solutions is that during congestion an intermediate node should allocate its own memory based on who has sent the messages competing for it. This implies that intermediate nodes should be able to recognize message's origin.

The simulation set-up

The problem caused by resource hogs and our solutions to it were tested in simulated environment using the opportunistic network environment (ONE) simulator [68]: 250 mobile nodes representing pedestrians with handheld devices moved along the streets of Helsinki City downtown for 12 hours. Message sizes generated by these mobile nodes were uniformly distributed between 0.2 and 2 MB. The buffer size of 20 MB in each node was dimensioned to accommodate 20 messages on the average. The benign 225 nodes generated one message per hour per node. The remaining 25 "hogs" generated messages at the rate of ten messages per hour per node. We have experimented with two species of hogs: naive hogs routed other nodes' messages (as any benign node would); malicious hogs accepted and then immediately dropped other nodes' messages.

The nodes routed messages using binary Spray and Wait protocol [32] with the initial number of "copies" parameter in the message header set to six.[5] With this setting the protocol replicates each message at most twice, so the total amount of copies of each message after the replications is at most four.

The trouble caused by the 25 naive hogs in our test-bed is illustrated in Figure 2.4 via the average buffer utilization in benign nodes as a function of time. At time zero the buffers of all nodes are empty. We see, first, that congestion (i.e. the state of average buffer utilization close to one) occurs in our network after a ramp-up time of about 2.6 hours. Second, during congestion the traffic of 25

[5]Please note that we have initially experimented also with two other routing protocols: epidemic [29] and Prophet [69], and found that the message delivery ratios are the best with Spray and Wait. See Figure 1 in the publication.

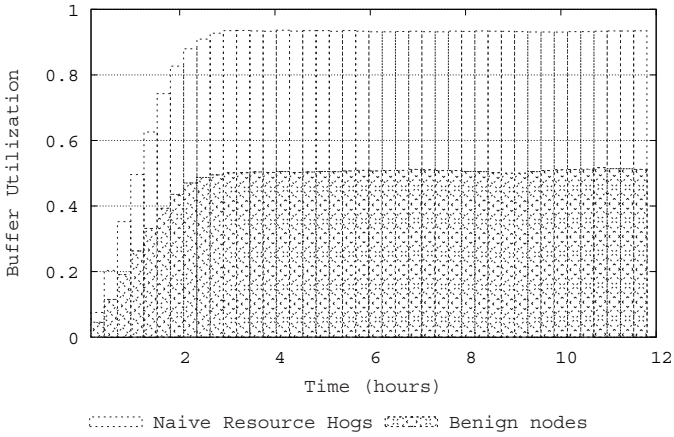

Figure 2.4. Average buffer utilization in benign nodes as a function of time.

hogs (one tenth of the network population) in the top part of the figure occupies about half of the buffer space in benign nodes.

Memory management techniques

I shall describe next two memory management techniques that turned out to be effective in our simulation-based evaluations.

The first technique is called "coarse-grained resource management" in the publication: The nodes in a geographical area organize themselves into an administrative domain and can distinct (based on message authentication) between messages from the domain members and outsiders. Messages from outsiders are carried and forwarded if there is no congestion (i.e. if there is room in the buffer). But during congestion incoming messages from the domain members have priority over messages from outsiders: the latter are dropped to make room for the former. Figure 2.5 (Figure 7 in the publication) shows how domain size effects the normalized delivery ratio of benign nodes. We see that when the domain members constitute a large majority, say, more than 70% of the network nodes, this technique effectively deals with "outsider" hogs in our scenarios. [6]

The second technique is called "fine-grained resource management with equal sub-domains" in the publication. It combats "insider" hogs that are members of the domain. If the buffer is congested after all outsider messages have been

[6]The normalized delivery ratio is above 95% in the "naive outsiders" scenario, where the outsider hogs route other nodes' messages; and it is still above 90% in the "malicious outsiders" scenario, where outsider hogs accept and immediately drop other nodes' messages.

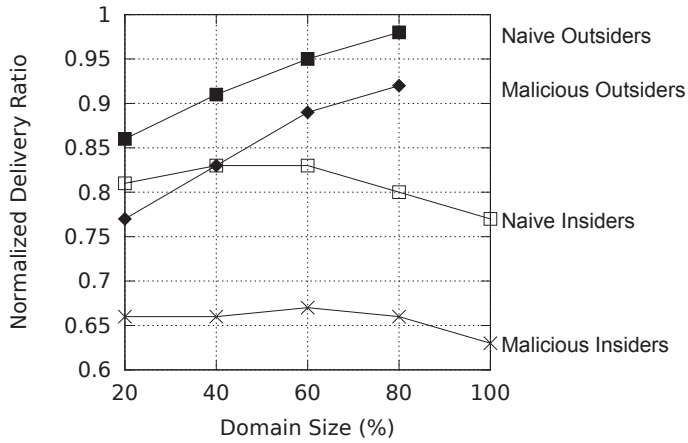

Figure 2.5. Effect of "coarse-grained" buffer management with 10% resource hogs. During congestion the domain members prioritize messages from within the domain over outsiders' messages. We plot the normalized delivery ratio of benign domain members as a function of domain size (shown here as a percent of the total network population).

removed by the basic "coarse-grained" technique above, we apply this method on the remaining insider messages. The idea is that during congestion the buffer of a domain member should be divided equally between members of the domain competing for that space. Therefore, when deciding whose message should be dropped during congestion the node simply chooses the principal with the highest buffer occupancy.

Equal sub-domains combines simplicity of implementation with high delivery ratio. A disadvantage of this method, shown on Figure 10 in the publication, is a pronounced dive in the delivery ratio of large messages.

Discussion

I will conclude with a discussion and analysis points for this work.

Interaction with the routing protocol

Our implementation of preemptive buffer management was generic enough to be used with several different routing protocols. The implementation selects a set of messages from the buffer and invokes the routing protocol to determine which of these to discard. Replication of messages by the routing protocol partly offsets message deletions by the preemptive buffer management.

But using this approach with routing schemes which already define some form of buffer management requires further study. Uncoordinated acceptance or dropping of messages could potentially impact the overall routing efficiency.

Buffer vs. transmission management

As described in section 5.1 of the publication our implementation of Spray and Wait routing protocol first transmits messages destined for the current peer node before sending the remaining in first-in-first-out (FIFO) order. Recall that our simulated network consist of 250 nodes moving randomly with pedestrian speeds in Helsinki city downtown. Node mobility in that environment results in the random "mixing" of messages from senders in the buffers of intermediaries.

If messages from different senders are randomly mixed in the intermediate node's buffer, then FIFO transmission order results in transmission opportunities allocated between buffer users according to their proportional share of buffer space. In particular, if the buffer is equally divided among senders, then FIFO transmission order will also result in an equal share of transmission opportunities. Equal sub-domains buffer management strives to achieve equal sharing of buffer space between domain members during congestion. This also implies equal sharing of transmission opportunities. More generally, independent of mobility model, transmission of messages in random order combined with equal sub-domains buffer management would lead to equal sharing of transmission opportunities between domain members during congestion.

A complementary way to solve the resource hogs problem is by directly managing also transmission opportunities of each principal. This is done by Fair Queuing (FQ) [70] and many subsequent schemes (e.g. [71–74]): Each traffic source is allocated its own queue in an intermediate node and the transmission process of that node takes a packet from non-empty queues, e.g., in a round robin fashion.

The delivery ratios of honest domain members could possibly be increased by intermediate nodes applying FQ on outgoing messages, in addition to the sub-domain buffer management on incoming messages. Verification of this conjecture could be a topic for future work.

Messaging patterns

Another topic for future work is the effect of messaging patterns of applications. Our experiments only use individual messages from a sender to a receiver. Many applications, however, involve a message exchange (request and response, or several of these). A delivered request may be useless if the corresponding response is dropped. How can buffer management schemes be enhanced to take this into account?

For example, a sender may issue a "ticket" for the receiver so that intermediate nodes can assign the same priority to the receiver's response as to the original

request. D. Schürman, J. Ott and L. Wolf have studied this topic further in [75].

"Equal subdomains" in Nokia Awarenet

Awarenet (Awareness network) is an experimental multi-hop ad hoc network between mobile phones that has been developed in Nokia Research Center starting in 2008 [1, 2]. Awarenet is intended mainly for devices discovering local information about people and services without infrastructure support. For that reason Awarenet nodes are designed to broadcast queries and replies regularly and often. Messages are supposed to be small, about 100 bytes on the average, to save energy.

Simulation studies done internally in Nokia have shown that equal subdomains technique effectively deals with buffer congestion in Awarenet: Each Awarenet node targets to divide congested buffer space equally between its one-hop neighbors. As a result, this technique was taken into use in Awarenet platform.

2.5 Best-effort authentication: publication V

This publication is about authentication of message fragments by intermediary nodes in a mobile opportunistic network.

The starting point is our technique for congested memory management in publication IV: intermediary devices belonging to the same administrative domain prioritize messages from their domain over messages from outside of it. During congestion this technique helps domain members to protect themselves from resource "hogs"—i.e. devices that generate disproportionately large amounts of messages—that do not belong to the administrative domain. Domain members need, therefore, to distinguish between in- and out of-domain messages, which implies message authentication.

To authenticate (and integrity protect) messages the sender could compute a hash of a bundle, sign this hash using the sender's private key and attach both the hash and the signature to the end of the original message. Suppose however, that a device is sending a message when the underlying connection breaks. The receiver has not yet received the entire message and thus it cannot verify the hash and authenticate the received piece. The sender cannot just send the remaining message part via another path in the network; it has to either wait for the recovery of the broken link or send the whole bundle again via another path.

One way to authenticate fragments has come up at the DTNRG mailing list

discussions [76]: Quantize the message before transmission and make each piece self-authenticating by attaching to it a signed hash. (Please note that since signatures are typically large and both generating and verifying a signature are computationally extensive operations, a performance overhead is introduced.) A more sophisticated approach from [77] is based on constructing a binary hash tree from the quantized message parts before transmission. With both methods the quantization borders—i.e. the points within the message where it may be fragmented—have to be defined before the actual message transmission, and the natural question is what those borders should be. We will go into this question in Chapter 3.

The problem we are addressing in this publication is that, in general, adding message fragmentation to use the contact times better, may substantially increase the amount of computation in intermediary devices, because they will have to authenticate lots of fragments. [7]

The solution is inspired, in part, by spot-checking of passengers and goods in the airports. We define "best effort" authentication as authentication with negligible (almost zero) probability of false negative, but large (for instance 0.5) probability of false positive, and show that this technique may be good enough for authenticating message fragments by intermediate nodes in our network. In allowing a relatively large probability of false positive, we differ from the typical applications of authentication in networking that strive to make this probability as small as possible.

Spot-checking is perhaps the simplest way to implement best-effort authentication. It involves taking an existing strong authentication scheme and performing verification only on a randomly chosen subset of received fragments. But there are other ways too. For example, in sections 4 and 5 of the publication we describe how Bloom filter[8] could be used for this purpose.

To verify this solution we have added a fragmentation mechanism to the simulation environment from publication V and studied its effect on delivery ratio of benign nodes in several scenarios, with and without fragment authentication by intermediate nodes. Three mobility models were used: the Map-Based from publication V (Helsinki downtown), the Random Waypoint (RWP), and RollerNet [79].

These studies are summarized in Tables 1-4 and Figure 1 of the publication,

[7]The increase in the amount of computation roughly corresponds to the number of pieces into which a message is split on its way to the destination. In hindsight, we could have measured and reported that increase in our simulated environments.

[8]This is a probabilistic data structure for encoding and testing set membership [78]. In our case the set comprises the quantized message's parts.

which shows that the probability of false positive can be as high as 0.6 without delivery ratio decreasing. We reproduce Figure 1 of the publication in Figure 2.6 below.

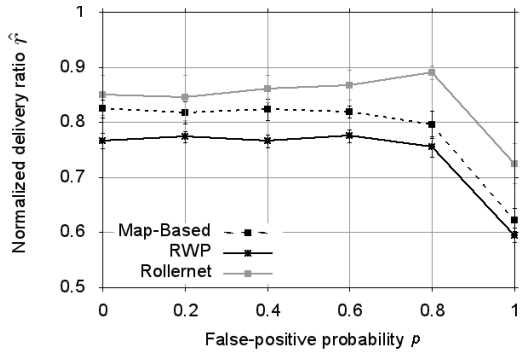

Figure 2.6. Impact of false positive probability p on the normalized delivery ratio $\hat{r}$ in three different scenarios.

The figure shows the overall impact of adopting best effort authentication on delivery ratio in the three mobility models. The Map-Based and the Random Waypoint (RWP) plots behave as we would expect: the delivery ratio does not change much as we increase the probability p of false positive (thus increasing the amount of hogs' traffic carried), until the tipping point somewhere between $p = 0.6$ and $p = 0.8$. But there is a surprise in the RollerNet plot: the RollerNet delivery ratio in Figure 2.6 peaks at $p = 0.8$.

Even though the peak may be just a random effect, there is also another possible explanation: This anomaly may be caused by the relative stability of RollerNet's topology. Like the other two networks, RollerNet's connectivity graph is sparse: a node is connected simultaneously to only a small portion of all other nodes. (This is apparent in Figure 2.7, which reproduces Figure 5 in [79].) But in addition, RollerNet's connectivity graph is also rather stable: the crowd of rollers glides through the streets stretching and shrinking in the forward direction, but there is not much mixing within it.

In this kind of network a hog partially isolates itself by overloading its one-hop neighbors. As a result from this partial isolation, the delivery ratio of benign nodes increases with the hog's traffic rate up to a point.

Conclusion

To conclude, we have shown that a best-effort authentication method, which is easier to attack than generic authentication methods but requires fewer computations for benign nodes, may be enough for certain networking scenarios. The caveat with best-effort authentication is that an authentication strength

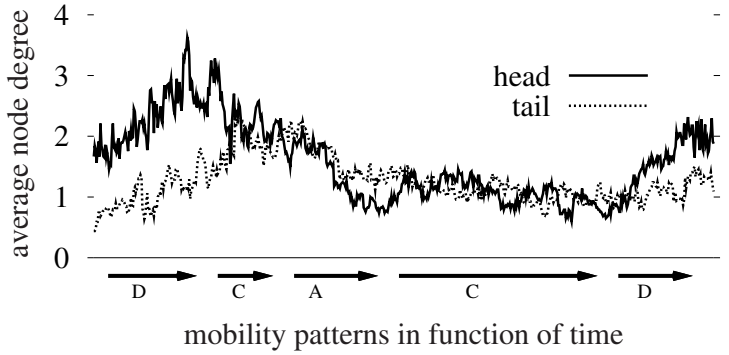

Figure 2.7. Figure 5 in [79] showing the average node degree for head and tail groups of nodes in RollerNet. (A) stands for acceleration period, (C) for constant speed, and (D) for deceleration of the group.

that is enough in one scenario may be unfit for another.

2.6 Location privacy: publication VI

This publication is about location privacy in a mobile ad hoc network. It is based on a four-months long experiment in EPFL campus, where an attacker having 37 access points in an area of 186 by 66 m^2 tracked daily the location of 80 mobile ad hoc network nodes. (See Figure 2.8 below.) The nodes, in turn, tried to evade tracking by changing their network identifiers from time to time. They were actually Nokia mobile phones with Nokia Instant Community (also called Awarenet) software [1, 2] in the hands of EPFL students. We call the transitory device identifiers "pseudonyms" in this publication. The algorithm for changing pseudonyms was developed and implemented in the devices before the trial started. The attacker's algorithms for tracking devices were finalized after the trial, when all the data from the sniffing stations had been collected. In experiments on the collected data we could also make tracking more difficult, e.g., by removing the records collected by some of the sniffing stations from the data set.

Awarenet is a mobile ad hoc network intended mainly for devices discovering local information. Its nodes should broadcast queries and replies regularly and often. Messages are supposed to be small, about 100 bytes on the average, to save energy. Since the constant size of the message header is 20 bytes, the message body is supposed to be only 80 bytes on the average. This restriction on the message size is by design, rather than due the physical limitations of the network. An additional precaution keeps the amount of forwarded message

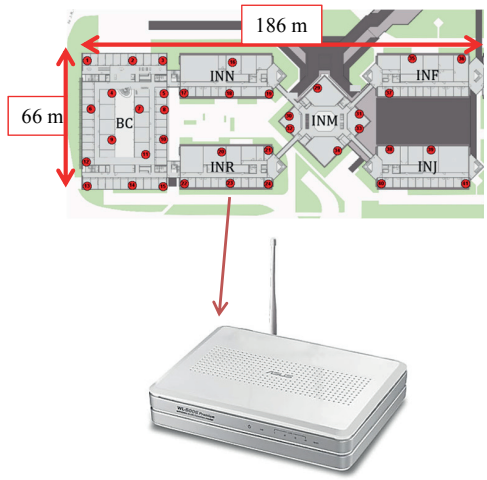

Figure 2.8. Floor plan of the attacked area on EPFL campus. The numbered circles are the attacker's Asus routers.

copies under control. Its idea is that a node will not retransmit a message that it scheduled for sending, if it notices that one of the neighbors happens to retransmit that same message first. Please note that intermediate nodes that retransmit an Awarenet message require access to its header.

Attack and defence

Tracking based on message headers

Each device was configured to broadcast a message once in a few seconds. We have assumed that messages' content is encrypted and the attacker does not know how to decrypt it. This hinders the tracking by looking into messages' content. But message headers in Awarenet are not encrypted and so the attacker can attempt to track user's location by linking messages containing the same sender pseudonym. Even though encrypting message headers would prevent linking, we decided not to employ this technique because decryption / encryption of headers would significantly increase the amount of computations that devices need to do when forwarding messages.

Pseudonym change algorithm (PCA)

To deter location tracking simply by linking messages containing the same sender pseudonym, devices change their pseudonyms from time to time. Some of these changes are unconditional and occur at random times during the day. Others happen only if the current context is considered favorable, like if the

number of neighbors exceeds ten, for example. In addition, the devices implement the concept of mix zones [80, 81]. The idea is that several nearby devices will simultaneously enter a period of radio silence, change their pseudonyms, and each device will start communicating under its new pseudonym after a device-specific random delay. This coordinated change of pseudonyms in several devices is triggered by a mix-request message broadcasted by one of them. We illustrate the mix zone concept in Figure 2.9.

Each pseudonym change and especially radio silence may disrupt ongoing communication. Indeed, powering down your device once and for all will make you invisible to the sniffing stations, but at the cost isolating your device. Thus, there is a trade-off between location privacy and connectivity in setting the algorithm's parameters.

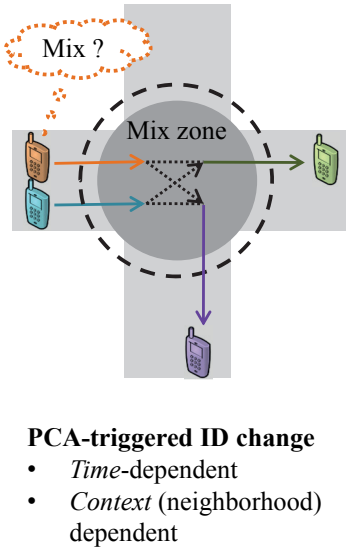

Figure 2.9. Illustration of the mix zone concept, showing silent period and identifier changes.

Analysis of the findings

The trial reported in this publication seems to be the first field-study that evaluated context-based identifier-change mechanisms under a practical adversary model with real mobile devices and communication scenarios. Earlier work on location-privacy preservation with mix zones, like [82] and [83], were mostly about vehicular networks and done using simulations.

Weakening / strengthening the attack

Our implementation of the attacker follows the theoretical attacker model in which the local attacker can eavesdrop on message traffic, but not modify, delete

or inject messages. Our experiments indicate that the actual attacker consisting of 37 sniffing stations covering an area of approximately 200×100 m^2 on the EPFL campus can track devices in that area with high probability, even though they employ the mix-zones technique; and that reducing the number of sniffing stations hinders tracking. The limiting case of weakening the attacker in this way would be a single sniffing station.

Conversely, the limiting case of strengthening the attacker by adding sniffing stations both inside and outside the campus is the global passive attacker model. Please note that this kind of attacker is still weaker than in the Dolev-Yao model [84], where the attacker has full control of the communication channel. In addition to eavesdropping the Dolev-Yao attacker can also modify or drop legitimate messages and send fake ones.

Simplifying/enhancing the defence

We turn now to the limiting cases of our pseudonym change algorithm. The algorithm could be simplified by diminishing the mix-zone component; until, in the limiting case, there is no mix-zone component left. Then the algorithm changes device's pseudonyms at random times without coordination with other devices.

Conversely, the mix-zone component of pseudonym change algorithm could be enhanced. In our experiments the mix-zone component was strengthened essentially by (i) lengthening silence periods—resulting in larger mix zones; and (ii) by reducing the time between pseudonym changes—resulting in more mix-zone instances. The limiting case of enhancing the algorithm in these ways contains all potential improvements to its mix-zones component.

But periods of radio silence in a mix zone, long enough to defer location tracking disrupt communication; and these disruptions seem to be too high for consumer devices. (See section 5.6 in the publication.) And while we could track nodes in the attacked area, we could not follow them outside that area. In practice (and contrary to the global attacker model) it may be reasonable to assume that the attacker is not everywhere.

Conclusions

First, even though the communicating devices in our trial employed state-of-the-art anti-tracking methods, the real-life attacker could follow them with high probability. Future work could look into designing better anti-tracking methods.

Second, recall that the attacker had 37 access points in an area of 186 by 66 m^2, i.e. he had a sniffing station placed about every 18 meters. While dense deployment of sniffing stations is realistic when the attacked area is relatively

small, in large geographical areas it requires rather strong organization.

A weaker but still realistic attacker model, which does not require such strong organization, could be defined as a collection of islands; all messages sent and received inside each island are visible to the attacker. And it could be required that while these islands cover only small part of the total area where devices may communicate, the benign devices do not know where the islands are. Another possibility is to require that significant portion of the communication between benign devices happens outside these islands. Evaluating pseudonym change algorithms against this weaker attacker could be continued in future research.

Third, a very simple anti-tracking algorithm, that changes the device pseudonym at random moments several times per day (and does not even attempt to evade a global attacker), may be in practice on a par with mix zone-based techniques in maintaining a degree of location privacy of hand-held consumer devices. Verifying this could also be a topic of future research.

3. Studies of fragmented message transmission

We study message fragmentation in the simplest case of a single disrupted link over which the message needs to be delivered in publications VII and VIII: Methods to estimate the mean transmission time of a fragmented message within a basic system model are developed in publication VII; publication VIII deals with fragmentation algorithms over a single link. In publication IX we study transmission times of fragmented message over multiple disrupted links.

To enable theoretical understanding we have reduced fragmented message transmission in challenged network in all its complication to a few elementary constituents and interactions. While our single- and multi-link models are interesting and mathematically tractable, none of the practical, actual setups follow exactly all our assumptions. Nevertheless, those models are useful in understanding these practical setups.

The cue to our modeling is the bundle protocol [24] in the Delay Tolerant Networking (DTN) architecture [13, 17]. Recall that the protocol offers transport services for applications and allows creation of self-contained messages that enable complete application interactions by a single message exchange. The self-contained messages (bundles) may be much larger than IP packets. The architecture is designed for networks where an end-to-end path may not exist, i.e., the messages are transmitted link-by-link. Any suitable convergence layer (like TCP or UDP) can be used for transmitting a bundle over a single link. Bundle has a lifetime, after which it expires and is removed from the network.

Main Premises

We discuss next four of the main premises in our system model.

1. Retransmit until it gets through
A message has a time-to-live parameter L assigned to it; that parameter reflects the end of message's usefulness to the receiver. If the link fails during trans-

mission, the sending node will attempt to retransmit this message (fragment) during the next contact time. It will continue to do this until the whole message has been transmitted, or until the lifetime of the message expires. The indication of successful transmission by the node at the other end of the disruptive link, or the recognition by the sender that the link has failed, are assumed to be completely reliable and immediate.

2. ON-OFF link model

We assume in our studies the existence of an underlying protocol stack that offers reliable delivery of messages between neighboring nodes. This protocol stack "hides" many of the bit errors and frame losses occurring at the wireless link layer from the higher layers, which see only a certain net data transfer rate between disruptions. It supports a simple ON-OFF model of the link states, where the link speed during the ON state is constant. Since we abstract the physics of the radio channel between two nodes into a ON-OFF contact pattern we are not concerned with things like the time it takes an electrical wave to traverse the distance between the nodes. The durations of the ON and OFF link states are assumed to have finite mean and variance. We divide all message sizes by the link speed, measuring them in seconds. For instance, if the link speed is 1 Kb/s, we say that the size of a 1 Kb message is one second.[1]

It is assumed that the distribution of the link state durations are known to the sender; in practice, they can be estimated from the past contact and inter-contact times. (One exception are the feedback-based algorithms H2 and H3 in publication VIII, where calculations based on this knowledge are replaced with self-adjusting feedback loop.)

Please note that the average speeds of an actual link seen by the bundle protocol may vary between contacts due to the changing physical conditions. In Appendix A of publication VIII we describe an additional preprocessing step after which methods for estimating transmission time of a message under the assumption that the average link speed is the same during all contacts, may be applied when the average link speeds are different.

An example will clarify this. Suppose that a Wireless LAN link that works according to the 802.11g standard was OFF for a second and then ON for another second. The link's transmission speed during the ON epoch was 24 Mb/s on the average. (So a sender can transmit at most 24 Mb during this on epoch.) The actual data exchange on a 802.11g link will always occur at lower rates than

[1] A natural unit for fragmentation questions is the mean contact (ON) duration. In hindsight, we could have divided all quantities having the dimension of time by the mean contact duration, making them dimensionless.

54 Mb/s. Therefore, we scale the durations of OFF and ON epochs so that the same 24 Mb may be transmitted within the scaled ON epoch at the theoretical maximum speed of 54 Mb/s, while keeping the sum of OFF and ON epoch at two seconds. The scaled ON and OFF epochs are $\frac{24}{54} \cdot 1\,\mathrm{s} = 0.44$ seconds and $1 + (1 - 0.44) = 1.56$ seconds, respectively. This scaling, when applied to past sequence of OFF/ON states having different average speeds during ON, results in a sequence of OFF/ON states having same speed during ON.

Notice that (i) same amount of data may be transmitted in any OFF/ON cycle of the scaled sequence as in the corresponding OFF/ON cycle in the original sequence, and (ii) the size of the individual OFF/ON cycle durations in the scaled sequence is the same as in the original one. As a result, the difference between transmission times of a message over the scaled sequence and over the original sequence is due to the differences in the first and in the last OFF/ON cycles only. (In the first OFF/ON cycle message transmission over the scaled sequence will start later than over the original one; in the last OFF/ON cycle message transmission over the scaled sequence will end later than over the original one.) This allows fragmentation algorithms to work with the scaled sequence of past link state durations (having constant link speed) when estimating message transmission times.

3. The sender quantizes messages before transmission

The sender prepares the message for fragmentation by dividing it into blocks of size f, called "fragmentation unit", and this is done before transmission. The quantized message size x is a function of the original message size m, the fragmentation unit f and the header size h in each block: $x = x(m, f, h)$. The message may be fragmented on its way to the destination only along the borders defined by f.

This is motivated by a combination of (1) security and (2) efficiency considerations.

(1) Firstly, since a contact between two nodes may abruptly end, the forwarding node (be it the originating or an intermediate node) must decide on the fragmentation borders and add the appropriate message authentication codes (MACs) *before* transmitting the message. (2) Secondly, marshaling and assembly of the message pieces at the destination is easier with fixed f.

These considerations support the design where the sender decides on fragmentation borders and creates the MACs needed for validating that message; the intermediate nodes will only validate.

The steps of message packaging (quantization) by the sender, and the subse-

quent transmission of the entire message over a single link are illustrated in Figures 3.1 and 3.2, respectively.

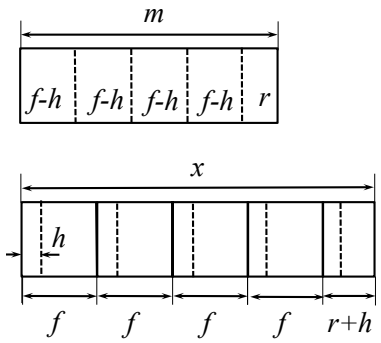

Figure 3.1. After quantization and addition of headers of size h by the sender the message size expands from m to x. The packaged message is ready to be fragmented along the borders defined by the fragmentation unit size f. The size $r+h$ of the last block in the quantized message may be smaller than f.

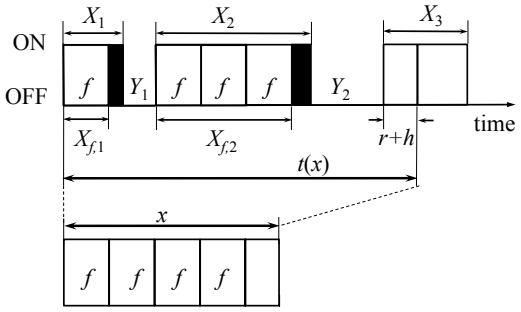

Figure 3.2. Schematic illustration of transmission time $t(x)$ of message x over a single link. The transmission time in that example includes two full ON epochs (contact times), X_1 and X_2; part of a third ON epoch X_3; and two OFF epochs (inter-contact times) Y_1 and Y_2. The shaded parts of X_1 and X_2 could not be utilized due to the message quantization into units of size f. The "useful" unshaded parts of the first two ON epochs, X_1 and X_2, are denoted with $X_{f,1}$ and $X_{f,2}$.

We remark that the possible reassembly of fragments in intermediate nodes does not affect transmission times in our model, as long as the fragmentation unit size f is constant during message transmission.

The possible values of f lie between two extremes: On the one hand, the sending node A has a lower bound on a packet size. This smallest packet size is also the smallest fragmentation unit size in practice, denoted with f_0. (For example, it could be the size of the header h plus the time it takes to transmit one byte.) On the other hand, the largest f that is possible in theory for a given message: $f = x$, results when the packaged message contains only a single block formed by the original message m and the header h. (Since the message is

quantized into a single block it cannot be fragmented.)[2]

It is evident from Figure 3.2 that if we let f and h approach zero, then the transmission time is at its theoretical lower bound: There is no waste of transmission time (no shaded parts in ON epochs) due to quantization of the message into indivisible blocks; and the transmission time is affected by link disruptions only. We derive ways to estimate the average value of this fundamental lower bound in publication VII. The premise that f approaches zero is notated there with "$f = 0$".

4. The less fragments the better

The quantization of the message into blocks before its transmission, enables message fragments consisting of one or more such blocks to be transmitted successfully, even if the connection time is too short for transmitting the entire message. But fragmentation in the actual network may also reduce the chances of delivering the entire message to its destination, due to lost or misrouted fragments. Increasing the block size f will tend to reduce the number of fragments and thus improve the chances of successful delivery. In our simple model we assume that if we decrease f, then the cost of transmitting a message increases.

There is, therefore, a tension in the sender's choice of f: On the one hand, smaller f may help to utilize contact timers better and thus decrease the time needed to deliver a message. On the other hand, larger f may increase the chances of delivering the entire message to its destination. We are thus led to investigate how a sender node A can reduce the number of fragments by setting the size of a fragmentation block f large, while still allowing in-time delivery of the message to the receiver B. Methods for doing this are developed in publication VIII.

3.1 Single link theory: publication VII

In this publication we develop methods to analytically estimate the mean time of fragmented message transmission over a single disrupted link when the distributions of the contact and inter-contact times are independent and identically distributed (i.i.d.). The statistics of the contact times do not need to be the same as the statistics of the inter-contact times. Please note that while the i.i.d. assumption enables us to do mathematics and in that way understand better the fundamentals of fragmented message transmission, we do not claim it to be true

[2] In practice, the largest possible f may be less than x. This happens when x exceeds the link's maximum transmission unit (MTU).

in the real world: Calculations of the mean transmission time made under that assumption may give a good approximation to the actual transmission time only if the circumstances do not change much during transmission of a message.

Our formulas for the mean transmission time $T(x)$ and the mean number of contacts $N(x)$ needed to transmit a message are summarized in Table I of the publication. (The leftmost column in the summary contains the number of the section in which the corresponding expression has been introduced.) There are two ways to estimate $T(x)$ for the case of generally distributed contact and inter-contact times (based on equations V.4 and V.11, respectively), $T(x)$ asymptotic behavior (equations V.12 and V.14), and several ways to compute $N(x)$ in special cases.

We have used those results as a source of approximations in our fragmentation algorithms in publication VIII, and in estimating the mean transmission time over multiple links in publication IX.

Discussion

In the rest of this section we will highlight two qualitative results of publication VII. These are i) asymptotic linearity of the mean transmission time with respect to the message size x, when the fragmentation unit size f is kept constant; and ii) non-linearity of mean transmission time with respect to the fragmentation unit size f, when x is kept constant.

i. When f is kept constant the mean transmission time of a fragmented message over a single link $T(x)$ is an asymptotically linear function of x for a typical distribution of ON and OFF epochs with finite mean and variance. (See equation (V.12) in the publication). We illustrate this fact in Figure 3.3. The slope of the asymptotic line depends on f, the distribution of ON epochs and on the mean size of the OFF epoch. The mean transmission time of large messages is insensitive to the distribution of OFF epochs (inter-contact times) beyond the mean.

ii. But the transmission time behaves in a non-linear way with respect to the fragmentation unit size f. This is due to a number of reasons.

1. An example of the first is that the contact times in challenged networks do not last forever. For instance, if any contact time is at most 5 seconds, then the message transmission time jumps to infinity when $f > 5$s. (If we set f to more than 5 seconds, then there is no chance of transmitting the message.) Generalizing from this example, a non-linearity in the mean transmission time with respect to f may be caused by a discontinuity in the distribution of

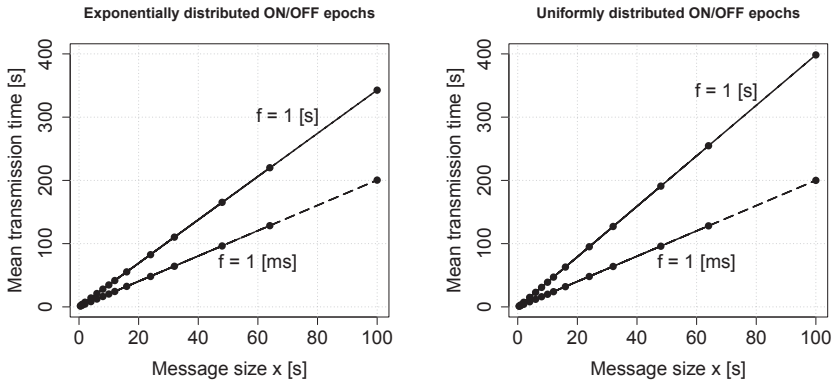

Figure 3.3. Illustration from [85] of the asymptotic linearity of the fragmented message's mean transmission time $T(x)$ over a single link when the ON and OFF periods of the link are exponentially or uniformly distributed and have the mean of one second. The fragmentation unit sizes are $f = 10^{-3}$ seconds and $f = 1$ second. Each dot "•" represents the mean transmission time in simulated environment of 2000 messages having size x.

the contact times.

2. The second is that, as illustrated in Figure 3.1, the quantized (packaged) message size x depends on f in a non-linear way: the growth of the message size due addition of headers to the message's blocks is roughly proportional to $1/(f-h)$, where h is the header's size.

3. After neutralizing the effect of the headers by setting header size h to zero, we have found an additional, deeper cause for non-linearity of transmission time with respect to f: Sometimes, increasing f makes the transmitted message pieces fit better into the forthcoming contact times; then bigger f will speed up the message transmission. This is illustrated in Appendix A of the publication, where we show first, that for a given sequence of the ON times, the transmission time may decrease when f increases. Second, we also show that the mean number of ON epochs needed to transmit a message (and hence the mean transmission time) may decrease when f decreases.[3] Thus the transmission time is not necessarily a monotonic function of f due to occasionally better fit with bigger f.

[3]We remark that this phenomena is not visible in plots that show asymptotic behavior of mean contact time, like those in Figure 3.3. This is because it occurs at smallish message sizes x whose order of magnitude is about the same as the mean contact time; a magnified view of the down-left corner in the figure would be needed to observe this phenomena.

All in all, increasing the fragmentation unit size f slows down the message transmission in our model, because there are less suitable contact times to carry the larger message pieces. But this slowdown is compensated somewhat: first, because there is less data to transmit (less extra headers); and second because the bigger message pieces may pack themselves better into the forthcoming contact times.

The non-linear behavior (including the occasional non-monotonicity) of transmission time with respect to the fragmentation unit size f has algorithmic implications. For example, consider communications in deep space where radio contact times are typically possible to predict: Suppose that—like the oracle-based algorithms in section 4.1 of our publication VIII—you are lucky to know in advance the sequence of future link state durations, and you want to chose the largest f that will still deliver the message in-time, before the deadline expires. Then, a binary search for f that gives the smallest transmission time over the sequence of future link states will not necessarily find the best f. You will have to do exhaustive search in all possible values of f, or design a search strategy that takes the non-linearity of transmission time with respect to f into account. The algorithm O2 in section 4.1 of publication VIII is an example of the latter approach.

3.2 Single link algorithms: publication VIII

Recall that in our model the message may be fragmented on its way to the destination only along the borders defined by the fragmentation unit f, and that errors causing packets drops can be added into our model as link disruptions.

In this publication we show how the sender may choose f in the simplest case of message transmission over a single disrupted link. We first classify the fragmentation algorithms, i.e. the algorithms for choosing f, based on the knowledge available at the sending node and design example algorithms in each class. Then we compare the performance of four fragmentation algorithms in six simulated environments. We have experimented both with constant and with exponentially distributed message sizes in those simulations. The assumption that link state durations are i.i.d., made in publication VII for mathematical simplicity, is waved aside in part of these simulations. The number of simulated scenarios is large because we wanted to understand how robust the perceived difference in performance between the algorithms is to changes in the environment.

We assume that the cost of transmitting a message of size m increases with

the number of fragments into which that message may be split. Our target is to minimize this cost, provided that the message transmission time $t(x)$ does not exceed the deadline L. Therefore, we seek the largest possible fragmentation unit f for a given message size m and the constraint L on the message transmission time.

The question of how to choose f can be approached in two ways: First, what is the largest f, for which the mean transmission time $T(x)$ does not exceed L? This is the problem of in-time delivery in the mean. Second, what is the largest f, for the probability of delivering the message before the deadline L is exceeds $1-\epsilon$? The lower bound $1-\epsilon$ on the probability of in-time delivery could be, e.g., 0.95. This is the problem of in-time delivery in probability. In this publication we give example algorithms designed to solve these problems for a single link. In the example algorithms for in-time delivery in probability the target probability of in-time delivery is achieved using a feedback mechanism.

We divide the algorithms into three classes: Oracle- Distribution- and History-based, according to the information known to the sender. Example algorithms in each class are defined in section 4 of the publication. Our comparative evaluation of four algorithms representing those classes is described in section 5.

An oracle-based algorithm cannot be realized as such in practice because the future is unknown. The concept of an oracle-based algorithm serves two purposes. First, the fragmentation unit size chosen by oracle-based algorithm may serve as a yardstick against which choices by practical algorithms are measured. Second, thinking about how an oracle would solve the in-time delivery problem gives an insight into building practical algorithms. The two practical algorithms that performed best in our simulations are the History-based H3 and the Distribution-based D1.

The first of those, H3, is the algorithm from the paper by Jelenković and Tan [86] that we have augmented with feedback loop to adjust f, as shown in Figure 3.4. We have added the feedback loop because the original Jelenković and Tan's algorithm fails to deliver in time when the message size x is large. This, and the effect of adding the feedback loop are shown on Figure 3.5.

In our evaluations of H3 and other algorithms that incorporate corrective feedback, the target of the feedback loop was that the overall proportion of in-time message deliveries to 0.95. We have found experimentally that when the targeted proportion of in-time delivery is higher, e.g., 0.99, the fragmentation algorithms do deliver messages in-time, but at the cost of splitting the message into many small fragments.

The second algorithm, D1, searches for largest f such that the mean trans-

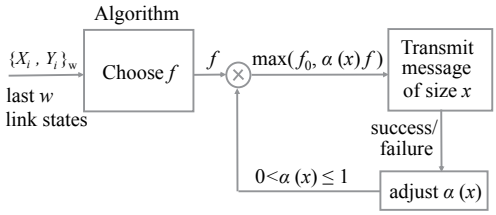

Figure 3.4. Feedback-based adjustment of f. In our implementation the range $[0,\infty)$ of possible message sizes is divided into sixteen bins (intervals): $[B_{i-1},B_i)$. Each bin $[B_{i-1},B_i)$ has a different correction factor α_i to the value of f returned by the algorithm, and that α_i is used when the quantized message size x is in $[B_{i-1},B_i)$. Furthermore, only α_i is adjusted after success or failure of transmitting that message in-time.

mission time $T(x)$ for a given distribution of link disruptions is still less than the deadline L. The distribution of link disruptions is estimated from past observations of link states. The estimate of $T(x)$ as a function of f is based on the results in our publication VII.

Here is the conclusion of our comparative study: If the target proportion of in-time delivery $1-\epsilon \leq 0.95$ is acceptable, then H3 is the best algorithm among those we have tried. It will deliver messages in-time with the least amount of fragments. But if the target proportion should be higher than 0.95, then the algorithm D1 should be used. It will deliver messages in-time with very high probability in most scenarios.

Discussion

In challenged networking we are typically interested more in the proportion of messages that arrive before the deadline expires (the delivery ratio), than in their mean delay. Therefore, in-time delivery in probability would be more interesting than in-time delivery in the mean. But as advocated by Jain, Fall and Patra in [30], optimizing performance with respect to the mean may indirectly improve the probability of delivery, also when improving the probability of delivery directly is hard. This point has been confirmed in our experiments.

End-to-end feedback is widely used to tune the performance of networking algorithms. And indeed, we have confirmed that a feedback loop may improve the performance of a given fragmentation algorithm with respect to the in-time delivery goal over a single link. But applicability of end-to-end feedback over multi-link paths in disrupted networks may be restricted by unreliable

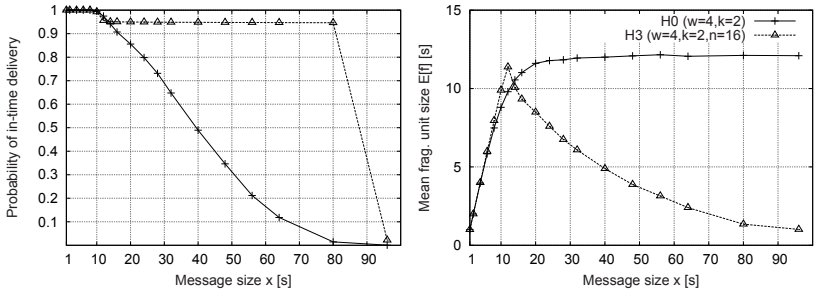

Figure 3.5. Part of Figure 6 (a) in the publication that shows the effect of adding a feedback loop on a performance of an algorithm in simulated environment. On the left is the probability of in-time delivery when the deadline is set to 100 seconds; on the right is the average fragmentation unit size f chosen by the algorithm. As explained in section 4.2.1 of the publication, the parentheses, like "(w=2, k=4)" after the label H0, contain the algorithm's parameters.
We see that the Jelenković and Tan's algorithm—labeled H0($w = 2, k = 4$)—fails to deliver in time when the message size x is large (left), because it does not change the average f when messages are large (right); while the same algorithm augmented with a feedback loop—labeled H3($w = 2, k = 4, n = 16$)—adapts to large messages by reducing f.

transmission of the feedback messages from the destination back to the source node. This restriction motivates further studies of fragmentation algorithms that can work in presence of unreliable multi-link paths.

Fragmentation affects the amount of work (computations and radio transmission time) that a device does per message. Future research could study the precise impact of these fragmentation effects on device's power consumption.

3.3 Multi-link theory: publication IX

This publication is about estimating the mean transmission time of fragmented messages over multiple disrupted links. We define a basic system model for the case where a single message is sent over a chain of n links. The disruptions in these links are identically and independently distributed; the chain is homogeneous in space and time and its links operate independently from each other. This case occurs, e.g., in a static multi-hop wireless network, where link disruptions can be due to physical intervention. We want to estimate the transmission time of a single fragmented message over an empty chain of disrupted links. The message may be rather long. The reason we chose to study this scenario is that it seems to capture one essential aspect of what happens in DTN. In some sense the only steady state in this scenario is the trivial case of the initially empty chain. For that reason, steady-state solutions of queuing networks are not directly applicable here.

We first give the simple upper and lower bounds in inequalities (1)-(3) in the publication. Then we derive an approximation formula for the mean transmission time (equation (5) in the publication), based on number of links, length of fragments and distributions of disruptions. The formula is verified against simulation experiments in the cases of uniform and exponential distributions for disruptions. I stress that while we are using this kind of disruptions to test our formulas, our methods are not restricted to disruptions having exponential or uniform distributions: when the distribution of disruptions has finite mean and variance, the mean transmission times of fragmented messages can be estimated using our methods.

The sending node needs to know the distribution of the time between link disruptions to estimate the mean transmission time with these methods. Assuming that knowledge in the sender is reasonable in the case of homogeneous chain that we address in this publication, because the distribution of the time between disruptions in the first link of the chain can be directly estimated by the sender from accumulated past observations, and the distributions in the other links are the same as in the first.

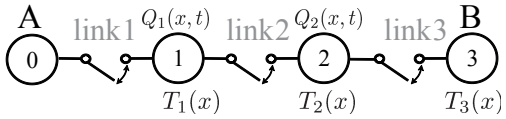

Figure 3.6. Schematic illustration from publication IX of a three disrupted links' chain between the sender A (node 0) and the receiver B (node 3). $T_1(x)$, $T_2(x)$, and $T_3(x)$ are the mean transmission times of a message having size x over one, two, and three links, respectively. $Q_1(x,t)$ and $Q_2(x,t)$ model the average amount of data queued in the intermediate nodes 1 and 2 at time t.

Our numbering of nodes and links in the chain is illustrated on Figure 3.6: we number the nodes starting with zero and the links connecting them starting with one. (For instance, the data comes into node one through link one and it goes out of that node through link two.) The derivation of our approximation publication starts with a simple observation: The transmission time of the whole message into node n is the time it takes to transmit the whole message over the first $n-1$ links, plus the time it takes to transmit over the last, nth link the data queued in the $n-1$ node, when all $n-1$ nodes before it are already empty. In our approximation for the mean transmission time we first replace all variables mentioned in the previous sentence with their respective means.

Second, we replace the average amount of data queued in the penultimate node $n-1$ at time t during transmission of message size x with the average amount of data queued at node $n-1$ at time t when x is very large—much larger

than any message we may wish to send over a chain of links. We illustrate the difference between these two queues in the first intermediate node on Figure 3.7.

An advantage of this substitution is that the always-growing mean queue size when x is considered "infinite" (i.e. very large) is simpler to estimate than when x is finite. And up to the moment when the transmission of the message over the first $n-1$ links ends, the "infinite" x–queue size and the actual, finite x–queue size in the node $n-1$ are close enough to each other for the purposes of our approximation.

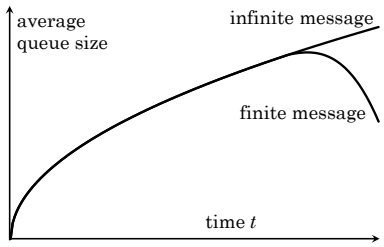

Figure 3.7. Schematic illustration of the average queue sizes in the first intermediate node when the message is finite and "infinite". The "infinite" message is much larger than any actual message that we may wish to send.

The result is the recursive equation (5) in the publication. To compute this recursion we need to know how to estimate (i) the mean transmission time over a single link (for any message size x), and (ii) the mean queue size as a function of time in intermediate nodes when x is "infinite".

The estimate of (i) is based on the results in publication VII in this thesis.

The estimate of (ii) is derived using the concept of random walk in section IV. In brief, we have modeled the length of the queue in node 1 (the first intermediate node of the chain in Figure 3.6) as a random walk with reflecting barrier at zero.[4] This model captures the fluctuations in the data flow through node 1 that are caused by the disruptions of its two adjacent links. (Those are the links numbered 1 and 2.) The model is more complicated farther along the chain, at the distance of $k = 2,3,\ldots,\ n-1$ links from the sender: Similarly to node 1, there will be fluctuations in the data flow through node k that are caused by the disruptions of its two adjacent links. (Those are the links numbered k and $k+1$.) But the disruptions of the preceding, nonadjacent $k-1$ links also affect the flow of data into the kth node; they cause additional gaps in that flow.

[4]This process behaves like an unrestricted random walk whenever it is positive and is pushed up to be equal to zero whenever it tries to become negative.

For that reason we have added negative drift to the random walk model when $k > 1$. The strength of that drift increases with k. Then we have estimated the mean size of the kth queue as a function of time using results from random walk theory and an approximation to the negative drift.

The approximations to the mean transmission time over five links in equation (5) were compared with mean transmission times obtained in simulations with three fragmentation unit sizes, several message sizes, and two distributions of disruptions, exponential and uniform.

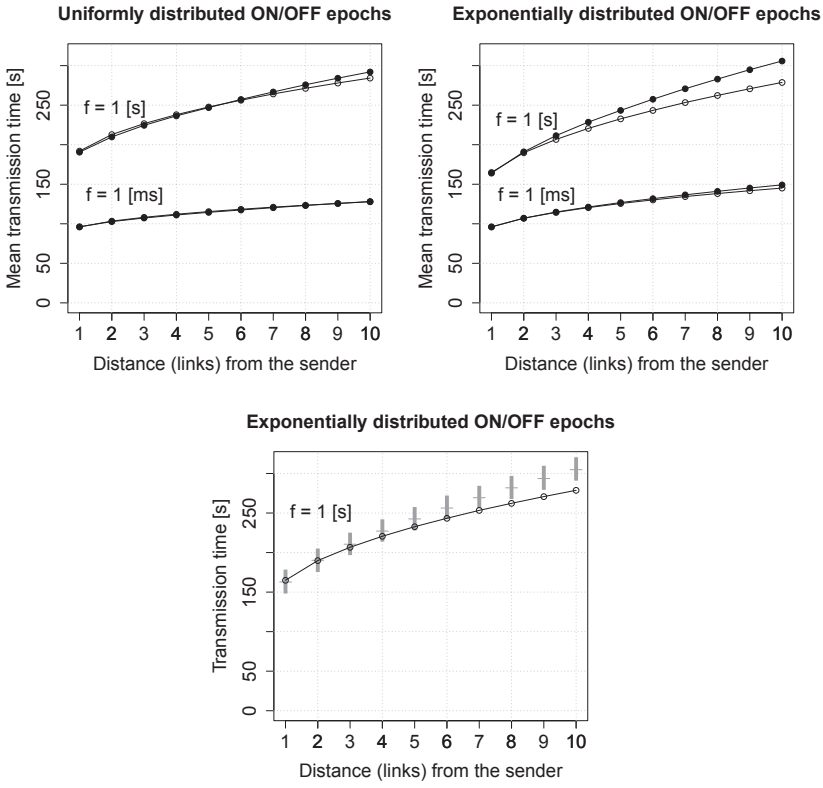

Figure 3.8. The top plots show the actual "•", and the approximated "◦" values of the mean transmission times when message size is 48 s over up to ten links with two fragmentation unit sizes f: 1ms, and 1s. Each actual value "•" is the average of 2000 transmission times. The mean duration of ON and OFF epochs in all links is 1s.
On the bottom plot we repeat the approximation "◦" to the mean transmission time with message size 48 seconds in the case of f = 1s and exponentially distributed disruptions. Among the four cases shown in the top plots, this is the one where our approximation diverges the most from the actual mean. We add to the plot the inter-quartile range and the median of 2000 actual transmission times. The upper and the lower tips of each gray bar indicate, respectfully, the 75% and the 25% quantiles; the bar's mid tick is the position of the median. It can be observed that up to the distance of six links from the sender our approximation is within the interquartile range of the actual transmission times.

In Figure 3.8 we extend the results shown in Figure 2 of the publication from

five to ten links. It is apparent that the mean transmission time is concave with respect to the number of links. From Table 1 in the publication we see that the maximum in absolute value relative error between our approximation and the actual mean transmission time decreases as the message gets larger, on the one hand; and it increases with the number of links, on the other hand.

In sum, our approximation is good for large message sizes that are at least few (say, ten) times bigger than the mean time between disruptions. The main reason for this limitation is that when the transmitted message is small, the queues in the intermediate nodes are mostly empty during the transmission time, and our approximation to their size that is derived assuming very long message does not apply.

We have recently found another formula for messages of the same order of magnitude as the mean of single ON epoch. We have included this result in the extended version [85] of this publication.

3.4 What we have learned

Several recent works study message fragmentation questions that are related to the ones in this dissertation. Jelenković and Tan showed analytically in [86] that the mean transmission time of unfragmented messages over a disrupted link may become arbitrarily large in some cases, and introduced a fragmentation algorithm (labeled H0(w,k) in section 4.2 of our publication VIII) to mitigate this problem. The performance of that algorithm is evaluated using simulations. Their analysis is done under a simplifying assumption that each transmission of a fragmentation unit block starts in the beginning of a contact time: within any contact time at most one block may be transmitted.

Nair, Andrew, Low and Doyle study in [87] the asymptotic behavior of message transmission time over a single link. They also show how to divide the message into blocks so that, under certain restrictions on the distribution of contact times, the mean transmission time of a message is minimized. The optimal choice consists of dividing the message so that all blocks except, perhaps, the last one are of equal size.

Jain, Fall and Patra [30] have identified many important issues in DTN routing, including message fragmentation. Pitkänen, Keränen and Ott [5] conducted extensive study of fragmentation in mobile DTN based on simulations. We have used their insights in our model formulation.

The essence of what can be learned from our publications VII, VIII and IX is as follows. A natural unit for fragmentation questions is the mean contact duration:

a message is small or large if it fits into (can be transmitted in) a typical contact or not. The dynamics of fragmented message transmission for small and large messages are different. This is why small and large messages seem to need different ways to estimate their transmission time over a chain of disrupted links.

Increasing the fragmentation unit size f slows down the message transmission in our model, because there are less suitable contact times to carry the larger message pieces. But this slowdown is compensated somewhat: first, because there is less data to transmit (less extra headers); and second because the bigger message pieces may pack themselves better into the forthcoming contact times.

Fragmented message transmission can be a complex process already in the single link case. Mean transmission time of a message over a single link is asymptotically linear in our model with respect to the message size x, when the fragmentation unit size f is kept constant; and it is non-linear with respect to the fragmentation unit size f, when x is kept constant.

When contact and inter-contact times (i) have finite mean and variance, (ii) are independent, and (iii) have distributions that do not change in time, we know how to estimate the mean transmission time of fragmented message over a single link.

We also know how to estimate the mean transmission time of a fragmented message over a homogeneous chain of n independent links having identical distributions of disruptions (again with finite mean and variance). It is notable that when a message is large this time is concave with respect to the number of links n.

Feedback loop may improve the performance of a given fragmentation algorithm with respect to the in-time delivery goal over a single link. But applicability of end-to-end feedback over multi-link paths in disrupted networks may be restricted by unreliable transmission of the feedback messages from the destination back to the source node.

4. Conclusion

In this work we have investigated (1) security and privacy, as well as (2) resource management, in networks where nodes are partially isolated from from each other and from support infrastructures. Partial isolation of the nodes aggravates resource scarcity. For example, the large pool of remote computation and storage in data centers—the "cloud"—can be accessed only occasionally, or not at all. Recall our thesis that (1) and (2) are closely related and tend to influence each other: On the one hand, security impacts on resource management whenever we allocate scarce resources to mitigate potential attacks on the system. On the other hand, access to scarce resources needs to be controlled, which implies authorization of devices that consume scarce resources. Authorization, in turn, may bring with it other security features, like authentication, data integrity and privacy. Please note that there is a trade-off here: there should be enough resources to support security features, which in the first place were introduced to protect scarce resources.

The close relation between (1) and (2) is evident in the table below, where we classify our publications according to these two lines. The bullet point • indicates the main line of the publication, while the circle sign ◦ indicates a secondary line that strongly impacts the main one.

Publication nr.	I	II	III	IV	V	VI	VII	VIII	IX
Security and privacy	•	•	•	•	•	•	◦	◦	◦
Resource mangement			◦	•	•	◦	•	•	•

In publication III we have put two designs for DTN security side-by-side and compared them based on the amount of resources (computations and the number of messages) required; in publication IV the protection of the benign node's memory from malicious outsiders needs message authentication; in publication V we show that techniques like spot-checking may save a large part of computations expended in intermediate nodes of mobile opportunistic network during

authentication of fragmented messages; in publication VI we conclude that periods of radio silence long enough to defer location tracking in an attacked area significantly disrupt communications by reducing the contract time between devices; and finally, in publications VII, VIII and IX the sender quantizes messages into constant-size fragmentation units to enable the cryptographic validation of message pieces in intermediate nodes.

The conclusion emerging from our work is that security and resource management in challenged networks are likely to impact each other. Thus, in a potentially hostile environment of challenged networks you (the network's designer) should not introduce resource management schemes without examining their impact on the network's security. Conversely, you should not provide security features without carefully considering first the resources they consume.

Bibliography

[1] A. Ahtiainen, K. Kalliojarvi, M. Kasslin, K. Leppanen, A. Richter, P. Ruuska, and C. Wijting, "Awareness networking in wireless environments," *Vehicular Technology Magazine, IEEE*, vol. 4, no. 3, pp. 48–54, sept. 2009.

[2] Nokia conversations, "Nokia instant community gets you social." [Online]. Available: http://conversations.nokia.com/2010/05/25/nokia-instant-community-gets-you-social/

[3] P. Ginzboorg, "Seven comments on charging and billing," *Commun. ACM*, vol. 43, no. 11, pp. 89–92, Nov. 2000. [Online]. Available: http://doi.acm.org/10.1145/353360.353369

[4] A. Seth, U. Hengartner, and S. Keshav, "Practical security for disconnected nodes," in *Proceedings of 1st IEEE ICNP Workshop on Secure Network Protocols (NPSec)*, November 2005, pp. 31 – 36, (the revised, 2006 version of the NPSec paper is at: http://www.cs.waterloo.ca/~a3seth/practical_security_v2.pdf).

[5] M. Pitkänen, A. Keränen, and J. Ott, "Message fragmentation in opportunistic DTNs," in *Proceedings of the Second WoWMoM Workshop on Autonomic and Opportunistic Communications (AOC)*. IEEE, June 23 2008.

[6] M. S. Corson, R. Laroia, J. Li, V. Park, T. Richardson, and G. Tsirtsis, "Toward proximity-aware internetworking," *Wireless Communications, IEEE*, vol. 17, no. 6, pp. 26–33, 2010.

[7] M. Stiemerling and S. Kiesel, "A system for peer-to-peer video streaming in resource constrained mobile environments," in *Proceedings of the 1st ACM workshop on User-provided networking: challenges and opportunities*. ACM, 2009, pp. 25–30.

[8] M. Hachman, "Peep proposes wireless P2P system," *PC Magazine*, January 2011.

[9] P. Ginzboorg, T. Kärkkäinen, A. Ruotsalainen, M. Andersson, and J. Ott, "DTN Communication in a Mine," in *Proceedings of the 2nd Extreme Workshop on Communications*, Dharamsala, India, September 2010. [Online]. Available: http://www.cl.cam.ac.uk/~fb375/extremecom/2010/GinzboorgExtremeCom10.pdf

[10] The Economist, "Cyber-warfare: Seek and hide," June 9 2012.

[11] ——, "Cyber-warfare: Hype and fear," December 8 2012.

[12] Helsinki Sanomat, "Laaja kybervakoilu paljastui," January 15 2013.

[13] Delay Tolerant Networking Research Group, http://www.dtnrg.org.

[14] C. Chyau and J.-F. Raymond, "What works: First Mile Solutions' Daknet takes rural communities online," October 2005. [Online]. Available: http://www.firstmilesolutions.com/documents/FMS_Case_Study.pdf

[15] "Investment case: Uepaa Swiss alpine technology," December 13 2011. [Online]. Available: http://www.equityfair.ch/media/fair/pr%C3%A4sentationen/uepaa.pdf

[16] Tolerant Networks. [Online]. Available: http://tolerantnetworks.com/?q=node/3

[17] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H.Weiss, "Delay-Tolerant Network Architecture," RFC 4838, April 2007. [Online]. Available: http://www.ietf.org/rfc/rfc4838.txt

[18] S. Farrell and V. Cahill, *Delay- and Disruption-Tolerant Networking*. Norwood, MA, USA: Artech House, Inc., 2006.

[19] V. Athanasios, Z. Yan, and T. V. Spyropoulos, Eds., *Delay Tolerant Networks: Protocols and Applications*, ser. Wireless Networks and Mobile Communications. CRC Press, Taylor & Francis Group, 2012. [Online]. Available: http://books.google.fi/books?id=2ERN5lgs3AwC

[20] L. Pelusi, A. Passarella, and M. Conti, "Opportunistic networking: data forwarding in disconnected mobile ad hoc networks," *Communications Magazine, IEEE*, vol. 44, no. 11, pp. 134–141, 2006.

[21] C. Caini, H. Cruickshank, S. Farrell, and M. Marchese, "Delay- and Disruption-Tolerant Networking (DTN): An alternative solution for future satellite networking applications," *Proceedings of the IEEE*, vol. 99, no. 11, pp. 1980–1997, 2011.

[22] M. Khabbaz, C. Assi, and W. Fawaz, "Disruption-tolerant networking: A comprehensive survey on recent developments and persisting challenges," *Communications Surveys & Tutorials, IEEE*, vol. 14, no. 2, pp. 607–640, 2012.

[23] Y. Cao and Z. Sun, "Routing in delay/disruption tolerant networks: A taxonomy, survey and challenges," *Communications Surveys & Tutorials, IEEE*, vol. 15, no. 2, pp. 654–677, 2013.

[24] K. Scott and S. Burleigh, "Bundle Protocol Specification," RFC 5050, November 2007. [Online]. Available: http://www.ietf.org/rfc/rfc5050.txt

[25] S. Farrell, S. Symington, H. Weiss, and P. Lovell, "Bundle Security Protocol Specification," RFC 6257, May 2011. [Online]. Available: http://www.ietf.org/rfc/rfc6257.txt

[26] S. Burleigh, M. Ramadas, and S. Farrell, "Licklider Transmission Protocol - Motivation," RFC 5325 (Informational), Internet Engineering Task Force, September 2008. [Online]. Available: http://www.ietf.org/rfc/rfc5325.txt

[27] M. Ramadas, S. Burleigh, and S. Farrell, "Licklider Transmission Protocol - Specification," RFC 5326 (Experimental), Internet Engineering Task Force, September 2008. [Online]. Available: http://www.ietf.org/rfc/rfc5326.txt

[28] T. Berners-Lee, R. Fielding, and L. Masinter, "Uniform Resource Identifiers (URI)," RFC 2396, August 1998. [Online]. Available: http://www.rfc-archive.org/getrfc.php?rfc=2396

[29] A. Vahdat and D. Becker, “Epidemic routing for partially connected ad hoc networks,” CS-2000-06, Department of Computer Science, Duke University, Tech. Rep., April 2000.

[30] S. Jain, K. Fall, and R. Patra, “Routing in a delay tolerant network,” in *Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, ser. SIGCOMM ’04. New York, NY, USA: ACM, 2004, pp. 145–158. [Online]. Available: http://doi.acm.org/10.1145/1015467.1015484

[31] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, “Single-copy routing in intermittently connected mobile networks,” in *Proceedings of First Annual IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks*, ser. SECON ’04, October 2004, pp. 235–244.

[32] ——, “Spray and Wait: An efficient routing scheme for intermittently connected mobile networks,” in *Proceedings of ACM SIGCOMM Workshop on Delay-Tolerant Networking (WDTN)*, 2005.

[33] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine, “MaxProp: Routing for vehicle-based disruption-tolerant networks,” in *Proceedings of the 25th IEEE international conference on computer communications*, ser. INFOCOM ’06. Barcelona, Spain, April 2006.

[34] Z. Zhang, “Routing in intermittently connected mobile ad hoc networks and delay tolerant networks: Overview and challenges,” *IEEE Communications Surveys and Tutorials*, vol. 8, no. 4, pp. 24–37, January 2006.

[35] A. Balasubramanian, B. Levine, and A. Venkataramani, “DTN routing as a resource allocation problem,” in *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 4. ACM, 2007, pp. 373–384.

[36] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, “Efficient routing in intermittently connected mobile networks: the multiple-copy case,” *IEEE/ACM Trans. Netw.*, vol. 16, no. 1, pp. 77–90, Feb. 2008. [Online]. Available: http://dx.doi.org/10.1109/TNET.2007.897964

[37] Y. Cao and Z. Sun, “Routing in delay/disruption tolerant networks: A taxonomy, survey and challenges,” *Communications Surveys & Tutorials, IEEE*, vol. 15, no. 2, pp. 654–677, 2013.

[38] S. Moloney and P. Ginzboorg, “Security for interactions in pervasive networks: Applicability of recommendation systems,” in *First European Workshop on Security in Ad–hoc and Sensor Networks, (ESAS 2004)*, ser. Lecture Notes in Computer Science, C. Castelluccia, Hartenstein, H., C. Paar, and D. Westhoff, Eds., vol. 3313/2005. Berlin Heidelberg: Springer-Verlag, 2005, pp. 95–106.

[39] J. Mogul and S. Deering, “Path MTU discovery,” RFC 1191 (Draft Standard), Internet Engineering Task Force, Nov. 1990. [Online]. Available: http://www.ietf.org/rfc/rfc1191.txt

[40] O. O’Neill, *A Question of Trust: The BBC Reith Lectures 2002*. Cambridge University Press, 2002, ISBN 978-0521529969. The five lectures comprising this book are available from the BBC web site: http://www.bbc.co.uk/radio4/reith2002/.

[41] P. Gutman, “Key Management through Key Continuity (KCM),” Internet draft, September 2008. [Online]. Available: http://tools.ietf.org/id/draft-gutmann-keycont-01.txt

[42] F. Stajano and R. Anderson, "The resurrecting duckling: Security issues for ad-hoc wireless networks," in *Security Protocols*, ser. Lecture Notes in Computer Science, B. Christianson, B. Crispo, J. Malcolm, and M. Roe, Eds. Springer Berlin Heidelberg, 2000, vol. 1796, pp. 172–182. [Online]. Available: http://dx.doi.org/10.1007/10720107_24

[43] D. Boneh and M. Franklin, "Identity-based encryption from the Weil pairing," *SIAM J. of Computing*, vol. 32, no. 3, pp. 586–615, 2003, extended abstract in Crypto'01.

[44] ZEIT ONLINE, "Tell-all telephone," 2011. [Online]. Available: http://www.zeit.de/datenschutz/malte-spitz-data-retention

[45] S. Bellovin and M. Merritt, "Encrypted key exchange: password-based protocols secure against dictionary attacks," in *Proceedings of 1992 IEEE Computer Society Symposium on Research in Security and Privacy*, ser. SP '92, 1992, pp. 72–84.

[46] K. Becker and U. Wille, "Communication complexity of group key distribution," in *Proceedings of the 5th ACM conference on Computer and communications security*, ser. CCS '98. New York, NY, USA: ACM, 1998, pp. 1–6. [Online]. Available: http://doi.acm.org/10.1145/288090.288094

[47] E. Uzun, K. Karvonen, and N. Asokan, "Usability analysis of secure pairing methods," in *Financial Cryptography and Data Security*, ser. Lecture Notes in Computer Science, S. Dietrich and R. Dhamija, Eds. Springer Berlin Heidelberg, 2007, vol. 4886, pp. 307–324. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-77366-5_29

[48] M. Cagalj, N. Saxena, and E. Uzun, "On the usability of secure association of wireless devices based on distance bounding," in *Cryptology and Network Security*, ser. Lecture Notes in Computer Science, J. Garay, A. Miyaji, and A. Otsuka, Eds. Springer Berlin Heidelberg, 2009, vol. 5888, pp. 443–462. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-10433-6_30

[49] N. Asokan and K. Nyberg, "Security associations for personal devices," in *Security and Privacy in Mobile and Wireless Networking (Emerging Communication and Service Technologies)*, S. Gritzalis, T. Karygiannis, and C. Skianis, Eds. Leicester, UK: Troubador Publishing Ltd, 2009.

[50] M. Rohs and B. Gfeller, "Using camera-equipped mobile phones for interacting with real-world objects," in *Advances in Pervasive Computing*, A. Ferscha, H. Hoertner, and G. Kotsis, Eds. Vienna, Austria: Austrian Computer Society (OCG), 2004, pp. 265–271.

[51] N. Saxena, J.-E. Ekberg, K. Kostiainen, and N. Asokan, "Secure device pairing based on a visual channel (short paper)," in *Proceedings of the 2006 IEEE Symposium on Security and Privacy*, ser. SP '06. Washington, DC, USA: IEEE Computer Society, 2006, pp. 306–313. [Online]. Available: http://dx.doi.org/10.1109/SP.2006.35

[52] W. Claycomb and D. Shin, "Secure device pairing using audio," in *Proceedings of 43rd Annual International Carnahan Conference on Security Technology*, 2009, pp. 77–84.

[53] R. Mayrhofer and H. Gellersen, "Shake well before use: Intuitive and secure pairing of mobile devices," *IEEE Transactions on Mobile Computing*, vol. 8, no. 6, pp. 792–806, 2009.

[54] S. Capkun, M. Cagalj, R. Rengaswamy, I. Tsigkogiannis, J.-P. Hubaux, and M. Srivastava, "Integrity codes: Message integrity protection and authentication over insecure channels," *IEEE Transactions on Dependable and Secure Computing*, vol. 5, no. 4, pp. 208–223, 2008.

[55] S. Gollakota, N. Ahmad, N. Zeldovich, and D. Katabi, "Secure in-band wireless pairing," in *USENIX Security Symposium*, 2011. [Online]. Available: http://static.usenix.org/events/sec11/tech/full_papers/Gollakota.pdf

[56] 3GPP, "TS 33.220, v11.4.0 (2012-09): 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Generic Authentication Architecture (GAA); Generic Bootstrapping Architecture (GBA) (Release 11)."

[57] N. Asokan, V. Niemi, and K. Nyberg, "Man-in-the-middle in tunnelled authentication protocols," in *Security Protocols*, ser. Lecture Notes in Computer Science, B. Christianson, B. Crispo, J. Malcolm, and M. Roe, Eds. Springer Berlin Heidelberg, 2005, vol. 3364, pp. 28–41. [Online]. Available: http://dx.doi.org/10.1007/11542322_6

[58] P. Eronen and H. Tschofenig, "Pre-Shared Key Ciphersuites for Transport Layer Security (TLS)," RFC 4279, Internet Engineering Task Force, December 2005. [Online]. Available: http://www.ietf.org/rfc/rfc4279.txt

[59] S. Holtmanns, V. Niemi, P. Ginzboorg, P. Laitinen, and P. Asokan, *Cellular Authentication for Mobile and Internet Services*. Wiley, 2008, ISBN 978-0-470-72317-3. [Online]. Available: http://books.google.fi/books?id=keVNRsQ94NUC

[60] C. Chen, C. J. Mitchell, and S. Tang, "Ubiquitous one-time password service using the Generic Authentication Architecture," *Mobile Networks and Applications*, 2011. [Online]. Available: http://dx.doi.org/10.1007/s11036-011-0329-z

[61] ——, "SSL/TLS session-aware user authentication using a GAA bootstrapped key," in *Information Security Theory and Practice: Security and Privacy of Mobile Devices in Wireless Communication, 5th IFIP WG 11.2 International Workshop, WISTP 2011*, ser. Lecture Notes in Computer Science, C. Ardagna and J. Zhou, Eds. Berlin Heidelberg: Springer-Verlag, 2011, vol. 6633, pp. 54–68.

[62] ——, "Building general purpose security services on trusted computing," in *Trusted Systems—Third International Conference, INTRUST 2011, Revised Selected Papers*, ser. Lecture Notes in Computer Science, L. Chen, M. Yung, and L. Zhu, Eds. Berlin Heidelberg: Springer-Verlag, 2012, vol. 7222, pp. 16–31. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-32298-3_2

[63] ——, "Building general-purpose security services on EMV payment cards," in *SecureComm 2012: 8th International Conference on Security and Privacy in Communication Networks*, Padua, Italy, September 3-5 2012, (proceedings to be published in the Springer-Verlag LNCST series).

[64] C. Newman, S. Hartman, and K. Raeburn, "RFC 4120: The Kerberos Network Authentication Service (V5)," July 2005. [Online]. Available: http://tools.ietf.org/html/rfc4120

[65] A. Lindgren and K. S. Phanse, "Evaluation of queueing policies and forwarding strategies for routing in intermittently connected networks," in *Proceedings of COMSWARE 2006: First International Conference on Communication System Software and Middleware*. IEEE, 2006.

[66] A. Krifa, C. Baraka, and T. Spyropoulos, "Optimal buffer management policies for delay tolerant networks," in *Proceedings of 5th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks*, ser. SECON '08. IEEE, 2008, pp. 260–268.

[67] Y. Li, L. Zhao, Z. Liu, and Q. Liu, "N-Drop: congestion control strategy under epidemic routing in DTN," in *Proceedings of the 2009 International Conference on Wireless Communications and Mobile Computing: Connecting the World Wirelessly*. ACM, 2009, pp. 457–460.

[68] A. Keränen, J. Ott, and T. Kärkkäinen, "The ONE Simulator for DTN Protocol Evaluation," in *SIMUTools '09: Proceedings of the 2nd International Conference on Simulation Tools and Techniques*. New York, NY, USA: ICST, 2009.

[69] A. Lindgren, A. Doria, and O. Schelen, "Probabilistic routing in intermittently connected networks," *Proceedings of First International Workshop on Service Assurance with Partial and Intermittent Resources (SAPIR)*, pp. 239–254, 2004.

[70] J. Nagle, "RFC 970: On packet switches with infinite storage," Dec. 1985. [Online]. Available: ftp://ftp.internic.net/rfc/rfc970.txt,ftp://ftp.math.utah.edu/pub/rfc/rfc970.txt

[71] H. Luo, P. Medvedev, J. Cheng, and S. Lu, "A self-coordinating approach to distributed fair queueing in ad hoc wireless networks," in *Proceedings of the 22nd IEEE international conference on computer communications*, ser. INFOCOM '01, 2001.

[72] S. Floyd and V. Jacobson, "Link-sharing and resource management models for packet networks," *IEEE/ACM Trans. Netw.*, vol. 3, no. 4, pp. 365–386, 1995.

[73] M. Shreedhar and G. Varghese, "Efficient fair queueing using deficit round-robin," *IEEE/ACM Trans. Netw.*, vol. 4, no. 3, pp. 375–385, 1996.

[74] J.-P. Yang, "Self-configured fair queueing," *Simulation*, vol. 83, no. 2, pp. 189–198, 2007.

[75] D. Schürmann, J. Ott, and L. Wolf, "Authenticated Resource Management in Delay Tolerant Networks using Proxy Signatures," in *Proceedings of Tenth International Conference on Wireless On-Demand Network Systems and Services (WONS 2013)*, Banff, Alberta, Canada, 2013.

[76] DTNRG, DTN-interest mailing list archive, April 2005. [Online]. Available: http://mailman.dtnrg.org/pipermail/dtn-interest/2005-April/

[77] N. Asokan, K. Kostiainen, P. Ginzboorg, J. Ott, and C. Luo, "Applicability of identity-based cryptography for disruption-tolerant networking," Nokia Research Center, Helsinki, Tech. Rep., March 2007, available at: http://research.nokia.com/files/NRC-TR-2007-007.pdf.

[78] B. H. Bloom, "Space/time trade-offs in hash coding with allowable errors," *Commun. ACM*, vol. 13, no. 7, pp. 422–426, July 1970. [Online]. Available: http://doi.acm.org/10.1145/362686.362692

[79] P. U. Tournoux, J. Leguay, F. Benbadis, V. Conan, M. D. de Amorim, and J. Whitbeck, "The Accordion Phenomenon: Analysis, Characterization, and Impact on DTN Routing," in *Proceedings of the 28th IEEE international conference on computer communications*, ser. INFOCOM '09, April 2009, pp. 1116–1124.

[80] A. Beresford and F. Stajano, "Location privacy in pervasive computing," *IEEE Pervasive Computing*, vol. 2, no. 1, pp. 46–55, Jan-Mar 2003.

[81] A. R. Beresford and F. Stajano, "Mix zones: User privacy in location-aware services," in *In Proceedings of the 2nd IEEE Annual Conference on Pervasive Computing and Communications Workshops*, ser. PERCOM '04, March 15-17 2004, pp. 127–131.

[82] L. Buttyán, T. Holczer, and I. Vajda, "On the effectiveness of changing pseudonyms to provide location privacy in vanets," in *Proceedings of the 4th European conference on Security and privacy in ad-hoc and sensor networks*, ser. ESAS'07. Berlin, Heidelberg: Springer-Verlag, 2007, pp. 129–141. [Online]. Available: http://dl.acm.org/citation.cfm?id=1784404.1784417

[83] M. Gerlach and F. Guttler, "Privacy in VANETs using changing pseudonyms – ideal and real," in *IEEE 65th Vehicular Technology Conference. VTC2007-Spring*, 2007, pp. 2521–2525.

[84] D. Dolev and A. C. Yao, "On the security of public key protocols," *IEEE Transactions on Information Theory*, vol. 29, no. 2, pp. 198–208, March 1983.

[85] P. Ginzboorg, V. Niemi, and J. Ott, "Message fragmentation for a chain of disrupted links," *Computer Communications*, 2014. [Online]. Available: http://dx.doi.org/10.1016/j.comcom.2014.03.015

[86] P. R. Jelenković and J. Tan, "Dynamic packet fragmentation for wireless channels with failures," in *Proceedings of the 9th ACM international symposium on Mobile ad hoc networking and computing*, ser. MobiHoc '08. New York, NY, USA: ACM, 2008, pp. 73–82. [Online]. Available: http://doi.acm.org/10.1145/1374618.1374629

[87] J. Nair, M. Andreasson, L. L. H. Andrew, S. H. Low, and J. C. Doyle, "File fragmentation over an unreliable channel," in *Proceedings of the 29th IEEE international conference on computer communications*, ser. INFOCOM '10. Piscataway, NJ, USA: IEEE Press, 2010, pp. 965–973. [Online]. Available: http://dl.acm.org/citation.cfm?id=1833515.1833669

[88] G. Steel, A. Bundy, and M. Maidl, "Attacking a protocol for group key agreement by refuting incorrect inductive conjectures," in *Automated Reasoning*, ser. Lecture Notes in Computer Science, D. Basin and M. Rusinowitch, Eds. Springer Berlin Heidelberg, 2004, vol. 3097, pp. 137–151. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-25984-8_8

Errata

Publication I

Here is the key agreement protocol in the right column, third page of the publication:

(1) $M_n \rightarrow \text{ALL} : M_n, P(E)$

(2) $M_i \rightarrow M_n$: $M_i, P(E(R_i, Si))$, $i = 1, \ldots n-1$

(3) $M_n \rightarrow M_i$: $R_i(\{Sj,\ j = 1, \ldots n\}); i = 1, \ldots n-1$

(4) $M_i \rightarrow M_n$: $M_i, K(S_i, H(S_1, S_2, \ldots S_n))$, for some i

Two modifications can be made to it based on the analysis in Steel, Bundy and Maidl [88]. The first is to include the sender's identifier in the encrypted part of steps (1) and (2):

(1) $M_n \rightarrow \text{ALL} : M_n, P(E, E(M_n))$

(2) $M_i \rightarrow M_n$: $M_i, P(E(M_i, R_i, Si))$, $i = 1, \ldots n-1$

Its goal is to detect attacker's switching the clear-text M_n, or M_i with another identifier.

The second is that step (4) should be addressed to all participants:

(4) $M_i \rightarrow \text{ALL}$: $M_i, K(S_i, H(S_1, S_2, \ldots S_n))$, for some i

The goal of this modification is to reduce somewhat the attacker's advantage when he is the leader M_n.

Publication II

Today I would have written "many", rather than "most" in the sentence "However, for most services cellular authentication would be good enough," in section 1, last paragraph.

Publication IV

1. Definition of SIZE(S) in Figure 5 on the fourth page of the publication as "fraction of buffer occupied by messages in S" implies that the messages in the set S are already in the node's buffer. This was not our intention. SIZE(S) is better defined as "the total size of the messages in the set S divided by the node's buffer size."

2. Third paragraph above the heading of section 5.2, on the fourth page of the publication: "Messages have a time-to-live (TTL) attribute set to 3 hours and their sizes are uniformly distributed between 100 and 2 MB." There should be "KB" after "100".

3. Fifth paragraph of section 6, on the fifth page of the publication: "We can estimate the time it takes to fill buffers by estimating the average buffer size and dividing by the average data generation rate. In our case: 20 MB/(10 messages/h × 1 MB/message), i. e., two hours."

The origin of the "10 messages/h" estimate in the data generation rate is not clear from the text; that value may have been obtained from simulations. Anyway, it seems too high by about 25%: Benign nodes, which constitute 90% of the population, generate one message per hour per node, while resource hogs, which constitute 10% of the population, generate 10 messages per hour per node. On the average, we have $0.9 \cdot 1 + 0.1 \cdot 10 = 1.9$ new messages/h per node. Message sizes were uniformly distributed between 0.2 and 2 MB. So the average message size was about 1 MB. The binary Spray and Wait routing protocol replicates these messages. With our setting of six initial "copies" the protocol replicates each message at most twice, so the total amount of copies of each message after the replications is (at most) four. Taking that as an estimate of the number of copies, we obtain that the data generation rate (including copies) is about $1.9 \times 4 \times 1$ MB = 7.6 MB per hour per node. The time it takes to fill the 20 MB buffer (and reach a steady state) at this rate can be estimated as 20 MB/7.6 MB, i. e., 2.6 hours. This estimate is still less than our three-hours TTL; and it is also more in line than two hours with what we see in Figures 6 and 8 of the publication.

4. The sentence above Figure 11: "At loads of 1000 and 2000 total messages the intermediary node buffers never become congested." should be "At offered loads of 1000 and 2000 total messages circulating in the network the intermediary node buffers never become congested." Similarly, the horizontal axis in Figure 12 should have the "Offered load" label.

Publication VII

The last sentence in the first paragraph on the right column of the Introduction section: “In both cases the fragmented data is re-assembled only at its destination,” is wrong: according to the RFC 5050 [24] also an intermediate node can re-assemble fragments into a new bundle.

Publication VIII

1. The last sentence in the third paragraph on the right column of the Introduction section: “In both cases the fragmented data is re-assembled only at its destination,” is wrong: according to the RFC 5050 [24] also an intermediate node can re-assemble fragments into a new bundle.

2. In the last sentence of the column below Figure 3, on the third page of the publication: “Please note that in reality the sender may also need to quantize a set of several messages; and we outline a way to do this for in-time delivery in the mean in B,” the words “in B” should be “in Appendix B”.

3. The description of the algorithm O2 on the fourth page, right column, continues until the bottom of the page. Hence the square sign “□” indicating the end of the algorithm’s description should be the last thing on the page, after Step 2.2.

Publication IX

In the caption of Figure 1, section II: “Schematic illustration of a chain of three disrupted links’ chain,” the words “of a chain” should be deleted.

In the last column of section III, first paragraph from the top, replace the second sentence with “For the limiting case where the fragmentation unit f is the smallest possible, we get a simple estimate by $T(x) = 2x$ because the link is in the ON state half of the time (in average) and it is possible to utilize the whole of each ON epoch for transmitting message fragments. This simple approximation ignores the well-known residual life paradox: when transmission starts during an OFF epoch, this first OFF epoch is expected to be longer than the average length of OFF epochs. But especially for longer messages, the effect of this error seems to be rather small.”

In section IV, fifth paragraph below Table I: the sentence “The relative errors were computed as the difference between the actual (simulated) value of $T_k(x)$ and our approximation to $T_k(x)$, divided by the actual value of $T_k(x)$.” should be

"The relative errors were computed as the difference between our approximation to $T_k(x)$ and the actual (simulated) value of $T_k(x)$, divided by the actual value of $T_k(x)$."

Aalto-DD 71/2014

ISBN 978-952-60-5692-0
ISBN 978-952-60-5693-7 (pdf)
ISSN-L 1799-4934
ISSN 1799-4934
ISSN 1799-4942 (pdf)

Aalto University
School of Electrical Engineering
Department of Communications and Networking
www.aalto.fi

BUSINESS + ECONOMY

ART + DESIGN + ARCHITECTURE

SCIENCE + TECHNOLOGY

CROSSOVER

DOCTORAL DISSERTATIONS