


FOR THE INTERNET OF THINGS

EDITED BY TYSON T. BROOKS



IEEE PRESS

WILEY

CYBER-ASSURANCE FOR THE INTERNET OF THINGS

Edited by

TYSON T. BROOKS

School of Information Studies
Syracuse University, Syracuse, NY, USA


IEEE PRESS

WILEY

Copyright © 2017 by The Institute of Electrical and Electronics Engineers, Inc.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey. All rights reserved
Published simultaneously in Canada

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permission>.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data is available.

ISBN: 978-1-119-19386-9

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

CYBER-ASSURANCE FOR THE INTERNET OF THINGS

IEEE Press
445 Hoes Lane
Piscataway, NJ 08854

IEEE Press Editorial Board
Tariq Samad, *Editor in Chief*

| | | |
|-------------------|-------------------|--------------|
| George W. Arnold | Xiaoou Li | Ray Perez |
| Giancarlo Fortino | Vladimir Lumelsky | Linda Shafer |
| Dmitry Goldgof | Pui-In Mak | Zidong Wang |
| Ekram Hossain | Jeffrey Nanzer | MengChu Zhou |

To Jesus Christ
and my parents

FOREWORD

Effective Cyber-Assurance Will be Essential for the Internet of Things

ZEAL ZIRING

Information Assurance Technical Director, National Security Agency,
Fort Meade, MD, USA

Our society has become substantially dependent upon the Internet, on the ability to access and use cyberspace, in a wide variety of ways. The Internet has given us amazing capabilities to exchange information, conduct commerce, enlighten, and entertain. But for all of the development and growth of the Internet, the virtual world and the physical world were at most lightly connected, often through the actions of people. The domain of packets and protocols was always separate from the world of fields, roads, and buildings. No longer the virtual world and the physical are becoming increasingly intertwined. The interposition has profound potential for benefits and for harm. This revolution-in-progress has been dubbed the Internet of Things (IoT), and cyber-physical systems (CPS), and various other names. It is a complex trend, founded on technology advances, but with economic and social drivers. It is already well underway, though we are feeling only modest effects so far.

As IoT technologies and capabilities become more prevalent, and eventually ubiquitous, many aspects of the physical world will become more visible from cyberspace. In some cases, processes in cyberspace will influence or control physical objects and environments. Points of contact between the physical world and the virtual will proliferate. There have been many estimates of how many connected “things” will be dispersed through our physical environment during the growth of IoT, from 10 to 50, to even 200 billion. As a result of the greatly increased integration between physical and virtual worlds, our dependence upon the Internet and associated technologies will increase.

There have been many books and articles written about the technologies driving the IoT, and the wonderful benefits we will realize from it. But those benefits are not certain. As the physical world becomes more dependent on the virtual, threats that are today confined to cyberspace will expand and transform. The benefits we hope

to enjoy will be at risk, subject to attacks mediated and scaled by cyberspace. This book is about understanding those risks: why they arise, how they differ from cyber risks that we face today, and especially how to address them.

SOME HISTORY

There have been many histories written about the Internet, focused on technology or people or other factors. One way of looking at the Internet is how it grew from convergence of previously independent systems and domains. This is relevant to understanding the IoT and the importance of its cyber-assurance, because it represents the biggest convergence yet.

From the beginnings of telephony and radio, military and civilian communications were distinct and separated. From the time of World War II, they used different technologies and different means of protection. Military communications were usually encrypted, used different frequency bands protocols, and infrastructure from their civilian counterparts. Since its creation in 1952, the National Security Agency (NSA) designed and codified the security necessary for national security communications, including military. Working there, it became clear around 1990 that convergence was inevitable. Over the course of two decades, from the mid-1990s to today, military and civilian (commercial) communications have become much closer: common technologies, protocols, infrastructure, and standards underpin both. Levels of cryptographic strength that were first envisioned for safeguarding national security are now used to protect both strategic intelligence and social media. Tactical military operations still use specialized radios, but they also use commercial smartphones and cellular standards. From the military side, convergence has been driven mainly by the greater functionality and capability available from the commercial products. From the commercial side, adoption of security mechanisms formerly confined to national security applications has been driven by the need for assurance and privacy for business conducted online.

Another convergence is still underway, though nearly complete: convergence of voice telephony and data networks. Voice telephony networks came first, of course, and by the time computing began to grow in the 1960s, national and international telephone networks were already well-established. In fact, the telephone network was so large and reliable that early digital communications used it as an infrastructure, converting digital data from serial lines into modulated audio signals, transferring them over the telephone networks, and then converting them back into bits at the other end. But over the course of the 1970s and 1980s, the telephone network itself became digital, and the same switching networks were used to carry voice calls and dedicated digital links (so-called “leased lines”). Some of the earliest wide-area data exchanges, such as bulletin boards and Usenet, employed these technologies. But at the same time, the foundations of packet networking were being created in universities and companies and the U.S. Department of Defense (DoD).

By the early 1980s, many of the key technologies were in place for the Internet to begin exponential growth. But the telephony network was still built around static

trunks lines and circuit switching. Over the course of the 1990s and 2000s, the core technologies of packet switching and Internet protocols were integrated into global telephony networks, and voice became just another kind of digital traffic on packet networks. Today, the global network fabric is entirely packet-based, and the distinction between voice service and data service is visible mainly for cellular systems. But the convergence of formerly independent voice and data networks has had security consequences. Voice telephony services can be attacked over data networks, but assurances built into modern networks can help protect both voice and data services.

One more convergence is also underway, and is especially relevant here: the convergence of industrial networks and public data networks. Computer control of industrial systems began in the 1960s with direct digital control (DDC) systems. The first programmable logic controller (PLC) system was built in 1968. By the late 1970s, PLCs were being connected using modems, serial links, and proprietary protocols. Standards for interoperability and transport of industrial control protocols over Transmission Control Protocol/Internet Protocol (TCP/IP) emerged in the early 1990s, but control systems were still connected and managed over dedicated links or leased lines. But since about 2000, controlling industrial systems over the Internet has been growing rapidly. There are several drivers for this convergence: reduced cost, greater operational flexibility, and especially integration of industrial control and monitoring systems with business systems. The benefits are substantial, but exposing industrial systems to direct or indirect access from the Internet imposes substantial risks. Control system components are generally designed for reliability, simplicity, and economy. Repeated tests by government, academic, and commercial labs have identified numerous vulnerabilities, consistently across the industry, for well over a decade. The trend toward connecting industrial control systems to the Internet, and integrating them with other Internet systems, is sometimes called “the industrial Internet,” as if it were a separate network – it is not.

Along with the convergence history sketched above, there is a parallel history of malicious activities directed at computers and data networks. That history is documented in multiple books and papers, only a few highlights are necessary to illustrate the growth of the threat. In the pre-Internet years, computers and networks were certainly subject to malicious acts, but they were relatively narrow in scope. Some early personal computer (PC) viruses propagated fairly widely, but were confined to a very narrow range of operating systems and applications. Military networks were subject to passive collection by nation-state actors, but that was expected and the risks posed were manageable – risks from passive collection can be managed with effective encryption.

In the early years of the global Internet, there were many large-scale malicious events, beginning with the Morris Worm in 1988, and continuing through the 1990s and into the early 2000s. While these infections garnered headlines, there was also a quiet growth of more sophisticated malware and capabilities for espionage. Also, with the growth of the World Wide Web (www), there was a corresponding growth in web defacement attacks. During much of this period, the value of information stored and business conducted on the Internet was modest. Many malicious actors were

motivated by notoriety: releasing a virus that spread worldwide garnered acclaim from peers. Web sites were important to the image of a company or government agency, and thus web defacers attacked the element of value that was most accessible to them. The primary suppliers of computer and network technology also began to take security much more seriously during this period. As an example, in 1992, Microsoft's flagship product was Windows 3.1, which shipped with effectively no security; by 2000, their flagship Windows 2000 product included a broad array of security features.

In the most recent decade, convergence has driven large portions of our economy, government, and society onto the Internet. The increase of value and diversity of connected systems and services has driven a corresponding growth and diversification of malicious activities. For example, greater use of Internet services for banking was quickly followed by Internet crime targeting bank accounts and transactions. Similarly, as national governments and economies became more dependent on the Internet, governments around the world have increased their use of the Internet as a domain for collecting intelligence and pressuring rivals. Many nations, including the United States, have incorporated cyberspace operations into their military doctrine.

We have also seen the first Internet-borne attacks where effects have extended beyond cyberspace into the physical world. Most of the early ones were accidental, denials of service by PC malware infecting PCs used to manage industrial controls. But by 2008, it was clear that some actors were deliberately targeting power utilities to conduct extortion. In 2010, the Stuxnet worm was discovered; it appeared to have been targeted at a particular industrial installation, propagated over the Internet and other networks, and caused physical damage to that installation (as well as disruption elsewhere).

The clear message from history is this: attacks follow value. The more value and dependence we place on the Internet, the greater motivation malicious actors, criminals, and hostile regimes will have to operate there. We are in the early stages of the biggest convergence yet, and the assurance we will require will be commensurately great.

THE BREADTH AND DIVERSITY OF THE INTERNET OF THINGS

The IoT is a very broad phenomenon, ranging across nearly every sector of industry, many different technology standards, and geographic scales. It encompasses both the connected "things" and the various data analysis, management, and infrastructure services with which they interact. The data and interaction are the foundation for the benefits we expect to gain – a single car with an Internet connection might help one driver navigate to their destination, but when a majority of cars are connected, analytics and active management will keep traffic flowing efficiently across a city. Innovative companies are devising new models for analyzing data and acting on it in sectors like housing, transportation, manufacturing, healthcare, public safety, energy, retail, and more.

The standards landscape for IoT is complicated, and in many areas, standards are still emerging or evolving rapidly. Standards are essential for IoT because they foster interoperability, stability, and innovation. There are many areas where standards will be essential, but four are particularly relevant to IoT cyber-security.

- Cellular communication – the radio spectrum is a finite, precious resource. As more devices join the Internet, managing the availability of that resource for all of them will be critical.
- Personal Area Networks (PAN) – standards for very short-range data exchange among wearable and nearby devices are still evolving to support all the capabilities and assurances we will need.
- Security and cryptography – most existing secure protocols, credential schemes, and other standards were designed for the world of desktop computers and enterprise servers. Standards will be needed to provide basic security services to large numbers of small, constrained devices. These services include identity and credential management, authorization, data protection, and more. As discussed below, IoT will impose new requirements in provisioning, efficiency, and scale.
- Sensing and data management – some of IoT's greatest benefits will flow from sensing aspects of the physical world, and exposing that data for analysis and fusion in cyberspace. Standards will be needed for representing and managing vast amounts of sensor data.

IoT devices will use a variety of modalities in connecting to the Internet. Some will be accessible only when activated by something else, such as a radio frequency identification (RFID) tag reader. Others will have periodic interaction, delivering data or accepting commands, but otherwise quiet (e.g., an implanted medical device, a weather sensor). Many devices will expect continuous connectivity to deliver data or allow remote entities to exert real-time control (e.g., a smart TV, an electrical substation monitor) and still others will act as local gateways, supporting local interaction and providing Internet connectivity for other devices within their scope (e.g., a smart car, bus, or train).

As described above, IoT will offer us enormous benefits, but most of those benefits will depend on some form of trust. We will need confidence enough in the operation of IoT devices and supporting services to entrust them with control of physical systems and environments. We will need confidence that the data delivered from sensors is accurate in order to rely on them when making personal, business, and even military decisions. Establishing and maintaining necessary trust will be challenging in many ways. Complete and comprehensive trust is not usually possible, even for narrowly scoped traditional computers. Instead, we will need to build systems that can deliver specific kinds of trust. We will need trust management and associated assurances at several levels for IoT systems: individual devices, populations of devices, users, services, and infrastructure.

WHAT IS CYBER-ASSURANCE FOR THE INTERNET OF THINGS, AND WHY DOES IT MATTER?

At the highest level, assurance for the IoT is just like assurance for other elements of cyberspace. But the scale and constraints of IoT, and the potential impacts of assurance failures, will mean that current strategies for achieving assurance will not be sufficient.

The five basic assurance properties are:

1. **Authenticity** – assurance that an entity claiming an identity does possess the right to use it. Assigning and authenticating identities will be challenging for IoT.
2. **Integrity** – assurance that information is created, modified, and deleted only by entities with the rights to do so.
3. **Confidentiality** – assurance that information is accessible or readable only by entities with requisite rights.
4. **Availability** – assurance that information or services are available or accessible under all conditions that it is supposed to be.
5. **Non-repudiation** – assurance that an action can be irrefutably bound to an accountable entity.

These assurances are primitives. By using and combining them, systems can offer higher order properties, such as privacy, legal compliance, or resilience. All of them will be important to the secure operation of IoT devices and the services they will support.

In addition to direct security risks to devices, IoT will have profound effects on the risk posture of traditional systems and networks to which they are attached. Connecting a broad range of IoT devices to conventional networks will expand the attack surface for those networks. To support the devices, conventional networks will have to support a broader set of protocols and data formats, adding new potential for exploitable vulnerabilities. Finally, many IoT use cases bridge traditional trust boundaries, or require system owners to establish new trust relationships. Build assurance into IoT devices and systems will be essential for managing these risks too.

Achieving the basic assurance properties for conventional networks has proven extremely difficult – recent security incidents have shown us that our technical measures and practices are not sufficient to prevent adverse impacts from cyber-attacks. Achieving the basic properties will be even more difficult for IoT systems. Why? First, the scale and diversity of IoT will require approaches and standards that span a very wide range. Device capabilities vary along several axes, such as computation speed, data storage, and communication bandwidth. For connected devices, some of these capabilities will vary over six orders of magnitude or more, from tiny tags and sensors to smart vehicles and buildings.

Another challenge for supporting assurance for connected devices and service is their diversity of security needs. Some devices will need very tight security rights – for

example, an implanted medical device will have very high integrity requirements, and should deliver data only to the patient and authorized doctors; in contrast, a weather sensor might offer data to any requester. Longevity will also present a challenge for assuring some IoT devices. Some devices will have the power and bandwidth to accept frequent security updates, but others will not. Some types of sensors, for example, will have to operate for years, and cannot be expected to receive any software updates or trust anchor updates in that time. This means that the security mechanisms built into such devices will need to be exceptionally simple and robust.

Finally, there will be many assurance challenges for IoT based on the relative immaturity of the law, policy, and practices for assuring IoT device data and access. Consider a smart building – what parties should be authorized to read the sensor data from the building’s systems? The building owner? The tenants? The local fire department? Maintenance workers, such as plumbers or electricians? Each of these stakeholders has a good rationale for accessing portions of the building’s data or adjusting aspects of the building’s operation. But neither the technical controls, legal precedents, nor accepted practices are ready to support them.

The IoT will let us use the flexibility and power of information technology to sense, understand, manage, and optimize many aspects of the physical world, from wearables on a single person, to a retail store, to a highway system. We can only depend on IoT to do these things for us, and enjoy the corresponding benefits, if we have certain essential assurances. The list below is based on the fundamental properties, but is tuned to be actionable for designers and builder of IoT systems:

- Assurance that collected data are valid (i.e., values reported are values sensed).
- Assurance that access to collected data is appropriately constrained.
- Assurance that control over devices is exercised only by authorized parties, and that those parties can be held accountable.
- Assurance that applicable laws, regulations, and policies are enforced.
- Assurance that the interactions between IoT systems and other cyber systems can be monitored and controlled.
- Assurance that overall security properties continue to hold as individual devices or components are updated or replaced.

The most important security properties for IoT will be system properties, assurances that are offered by, qualified over, and dependent upon multiple layers of hardware and software, service providers, data aggregation middleware, and presentation systems.

EXAMPLES

The examples below examine the assurance challenges for four different IoT scenarios.

Example 1 – A medical implant with connection to the Internet can offer faster detection of health problems, more nuanced responses, and better overall health monitoring. The devices themselves are subject to serious limitations on size, power consumption, and connectivity. There are immediate risks to use of such a device – a cyber-attack against it might pose direct threat to the user's health and safety. But an attack that alters data reported by a device may also pose such a threat, because medical treatment might be based on it. There are also strong privacy concerns around the collected data. Assurance for data access will be complex, because there are multiple stakeholders: the patient, their doctors, hospitals, first responders, insurance companies, the device manufacturer, etc. Also, medical devices and health data are subject to a complex regulatory regime that is still adapting to cyber threats.

Example 2 – A connected car will support a wide variety of use cases, from simple collision avoidance to entertainment to maintenance to full autonomous operation. There are large potential benefits for transportation safety and efficiency. Such a complex system will also have a complicated authorization model, with different rights for the driver, the mechanic, the manufacturer, highway systems, and network infrastructure. Some operations will be subject to hard real-time constraints, while others involve communication with the global Internet. Interactions between vehicles and smart highway systems are still being defined, but imply a very close trust relationship. Recent vulnerability demonstrations from researchers have shown that current vehicle telematics systems do not enforce trust boundaries effectively, that will have to change. Lastly, connected cars will connect to a wide variety of other networks, in owner's homes, at maintenance facilities, and while on the highway. There will need to be very specific and bounded trust relationships between each car and these networks.

Example 3 – Smart buildings will contain a wide variety of sensors, actuators, and control systems for a wide variety of purposes: lighting, safety, heating and cooling, entry control, and more. Many of these systems are installed to improve the cost efficiency of a building, or make it more hospitable to users. There will be some privacy or confidentiality concerns for the collected data. But the primary risks will be based on control: abuse of the control systems within a building can make it uninhabitable or even damage it. Control integrity and authorization will be key assurance concerns for smart buildings, but as noted above, the set of authorized users for such buildings will be large and diverse. In addition to the exposure from connection to the Internet, many building automation technologies employ wireless networks, using standards such as Wi-Fi, ZigBee, and Bluetooth. These can leave the network of a building exposed to anyone with physical proximity.

Example 4 – Sensor networks offer the potential from monitoring physical conditions across many different environments and locales. An ocean sensor network, for example, might be composed of sensor buoys, communication relays, and other floating and anchored elements. The components of the network will be widely distributed and subject to harsh conditions and uncertain connectivity.

The components may be power-constrained, expected to operate for long periods on stored power. The data collected from such sensors may be public, but its integrity may be critical for ocean navigation and weather prediction. Data from the sensor network will be fused with other sources in analytic systems, where there is likely to be much greater value to attract threat actors. This implies a need to manage the trust between the sensor network and the analysis systems, to prevent compromise of a sensor propagating upward.

These four examples show several common elements. First, integrity is a crucial concern for most IoT use cases – integrity of reported data, and integrity of control. Second, many of the suppliers that produce components for various IoT sectors have not, historically, had to worry about cyber-assurance for their products – it is only now that their products are exposed to such threats. Third, there is no simple model or universal model for trust relationships in these use cases. Each of them includes a variety of stakeholders with different roles and rights. Finally, none of the connected devices in these use cases operate independently, they all interact with other infrastructures and systems, and both inherit risks from and impose risks on those systems.

KEY ELEMENTS OF CYBER-ASSURANCE FOR IoT

Researchers, academics, professionals, and science-practitioners have a lot of work ahead to create an assured and trustworthy IoT. Research is already underway and needs to continue. Standards bodies and consortia have taken up the challenge of building security into many of the standards required. The next step is for the broader community, manufacturers, service providers, data aggregators, to build assurance into their offerings, and for users to demand it. We do not yet know all the assurances and security features that IoT will require, but we know some that will be essential. That kind of partial knowledge, and learning while building, had been a feature of every major convergence leading to today's Internet environment. We can learn as we build, but we must build in the essentials at every step. Some of those essentials are listed below, and explored more fully in the chapters of this book.

Basic security properties, the fundamentals, must be designed in to IoT devices, infrastructures, and back-end analysis systems. The security designs must reflect IoT requirements and constraints, and must enable high-level assurance as end-to-end guarantees. Chapters 1 and 2 explore general facets of designing cyber-assurance for IoT. Provisioning identities for IoT devices and services, and managing credentials, attributes, and rights associated with those identities, will be critical for supporting high-level assurance properties like privacy and access control. Several chapters touch on this area. IoT devices must be able to integrate securely into existing network services and enterprise IT environments – this will require certain security features in the devices themselves and substantial evolution in the way enterprises handle trust boundaries in which Chapter 3 explores this very challenging area. Establishing and maintaining assurance for IoT systems will depend on trust management services,

which will have to extend from individual devices to high-level data analysis services which Chapters 4 and 5 examine. Chapter 6 reviews the privacy and security concerns of wearable computing while Chapter 7 focuses on the vulnerabilities of industrial control systems. Chapter 8 approaches to leverage Big Data techniques to enhance IoT provenance, which is itself only one of multiple measures needed to improve cyber-assurance. Assurance is not something that can be established once and then forgotten – it must be actively managed, measured, and maintained and Chapter 9 explores the more general challenge of assessing security mechanisms. Chapter 10 researches the future artificial intelligence aspect of cyber-assurance and Chapter 11 explores the threats toward cyber physical systems for the IoT.

To ensure that the essential assurance elements are built into the devices and systems that will comprise the Internet of Things, it is necessary to raise awareness about the challenges and possible solutions. This book is one step in that direction. By raising tough issues, and presenting potential solutions, it will encourage discussion and debate, expose engineers and designers to new strategies and emerging standards, and promote active development of cyber-assurance. With those assurances, we will be able to take full advantage of the potential benefits of the IoT.

PREFACE

The Internet of Things (IoT) has resulted in the widespread deployment of a relatively immature technology. There are, however, many significant challenges faced by the programmers, designers, and implementers of IoT technologies in ensuring that the level of security afforded is appropriate. As innovative technologies using the IoT will focus more on wireless technologies, there are numerous complex considerations which must be taken into account when deploying wireless infrastructures and without adequate forethought their use may be ill-advised. Researchers and commercial organizations are predicting that there will be 50 billion devices connected to the Internet by 2020¹ and the potential economic impact – including consumer surplus – of as much as \$11.1 trillion per year in 2025 for IoT applications.² IoT networks will become popular because they can be deployed quickly with very little equipment infrastructures. These networks also lend themselves well to environments with populations of transient users. The possible applications of the IoT are almost limitless and organizations throughout the world have been quick to realize its potential.

The heavy utilization of wireless equipment and technologies renders the IoT operation very complicated. At the same time, the pace of data-in-transit and data-in-storage processing is significantly accelerated with the focus of how the data are delivered to the IoT systems, whereas the quick shifting of the focus will inevitably bring about swift and constant changes in the tactics of information security. Under such a highly complex and ever-changing environment, organizations must pay attention to the use of information security tools and techniques with a view to defeating

¹ <http://blogs.cisco.com/news/cisco-connections-counter>

² http://www.mckinsey.de/sites/mck_files/files/unlocking_the_potential_of_the_internet_of_things_full_report.pdf