Chapter 1

IoT Systems Introduction

Internet of Things (IoT) systems are rapidly changing the world around us and will continue to do so. These systems offer substantial potential benefit in terms of social value and business value in the institutions and cities in which they are deployed. However, *how* IoT systems are selected, implemented, and operated – and the ease or difficulty of implementation and operation – has significant consequences for the success of IoT systems implementation in institutions and cities.

All systems, whether natural, technical, or social, experience systems loss. IoT systems, these sociotechnical systems deployed and geographically distributed throughout the environments around us in our cities and institutions can have elements of all three systems – natural, social, and technical. The technical part, the technology part, of IoT systems is only a fraction of the overall system. Systems losses occur within each of these different types of systems and between these systems. All of the resources that appear to be available to implement and operate the IoT system do not get converted into actual positive, value-added output. A considerable portion of these ostensible resources – staffing, time, funding, technologies, and others – simply get converted to waste. As we will see, these systemic losses involving IoT systems will manifest themselves in lost Return on Investment (ROI) and degraded cybersecurity capabilities and cyber risk profiles for institutions and cities.

This systems waste, this loss of resources, is felt particularly fully in resource-constrained environments – which almost all institutions and cities are. This waste, and thus this impact on IoT systems success, can be mitigated substantially by paying attention to the *manageability* of the IoT system. This manageability is not just technical aspects. It presents itself in the deployed environment, supporting technical infrastructures, and most importantly supporting social and organizational environments within the city or institution.

Given the rapidly changing and dynamic aspects of IoT systems and the increasingly complex and resource-constrained environments in which they are deployed – and the number of variables that are outside of an institution's or city's control – the manageability of an IoT system or systems becomes critical for systems success (or failure).

The Potential Benefits of IoT Systems

The potential benefits of appropriately selected, procured, implemented, and managed IoT systems are substantial. Universities and institutions can benefit from IoT systems such as traditional building automation systems (e.g., heating ventilation and air conditioning (HVAC)), energy management and conservation systems, building and space access systems, environmental control systems for large research environments, academic learning systems, and safety systems for students, faculty, staff, and the public. Cities also benefit from IoT systems supporting public safety (e.g., surveillance of high crime areas), air quality monitoring by sector, transportation control systems, city accessibility guidance and support, and many others.[1]

In automotive and transportation systems, IoT can enable health checks of automotive components, Global Positioning System-based location monitoring, route optimization, crash prevention, car-to-car communication, and real-time traffic analysis. City governments and institutions can use traffic data for more effective city planning. In health systems, whether lifestyle, recreational, or patient monitoring for critical functions such as blood pressure, glucose levels, heart rate, or others, IoT devices and supporting systems can monitor, analyze, report health data, and even directly provide appropriately dosed medicine to patients.

Sensor-based analysis in retail spaces can provide business owners valuable analysis of customer behavior and buying patterns, reducing waste, and enhancing profitability.[2] Institutions and cities can use arrays of IoT devices, sensors, and actuators to monitor and analyze buildings, campuses, and spaces for energy usage. With this data, opportunities for increased energy efficiency can be identified. Regulatory and compliance requirements such as carbon emissions requirements can be measured, reported, and enforced. Further, aspirational objectives around carbon emissions, other air quality measures, water contaminant levels, and others can be recorded, studied, and reported.

Systems Loss

> Oh, ye seekers after perpetual motion, how many vain chimeras have you pursued? Go and take your place with the alchemists.
>
> **Leonardo Da Vinci**[3]

As with systems in nature, in social/societal organizations, and particularly in sociotechnical organizations – of which most modern societies are – there is always systems loss. Sociotechnical systems can have many components, facets, and attributes – there can be plans, intentions, resources, alignment, conflict, lag, cause and effect, uncertainty, and surprises. One thing is certain though – there is *always* system loss. Nothing is free. In the excitement, novelty, complexity, and hype around IoT and IoT systems, often what is not realized by cities and institutions is that there is still systems loss. Worse, not only is there systems loss, this loss is substantial and will directly and negatively impact the opportunity for successful implementation of the IoT system.

Systems Loss in Nature

The formalization of the study of systems loss has a rich history of study and publication. It is hard to imagine that the advances in science, medicine, technology, and society that we witness today would have been possible without the discovery, formalization, and documentation of systems loss.

From a scientific viewpoint, the quintessential study of the development and application of systems loss can be found in thermodynamics and, particularly, the second law of thermodynamics.[4]

In the early 19th century, French military engineer, Sadi Carnot, built on some of the work of his father, Lazare Carnot, and introduced the idea of an idealized heat engine. (Among other things, Lazare Carnot is also known for appointing Napoleon as the general-in-chief of the Army of Italy, subsequently being named Minister of War by Napoleon and later as Minister of the Interior by Napoleon).[5]

In his book, *Reflections on the Motive Power of Fire*,[6] (Sadi) Carnot abstracted out the core components of steam engines of the day into an idealized system so that consistent math could be performed in the context of these "heat engines (Figure 1.1)."

In the course of this, Carnot introduced the idea of a heat transfer pattern cycle, subsequently named the Carnot cycle.[7] In this abstraction, he showed that *there is always system loss* – that even when disregarding the effects of friction and machine imperfections (which also cause loss), Carnot proved that there is a maximum efficiency, well less than 100%, of *any* engine. That is, regardless of the machine (engine) or the type of fluid on which the engine runs – whether steam, gas, or others – *a portion of the energy added to the engine will not be converted to work.* That is, *a portion of the energy added will always be lost*. (This is also consistent with a concept that Leonardo Da Vinci had introduced that a perpetual motion machine is impossible.)

Rudolf Clausius,[8] a physics professor at the Artillery and Engineering School in Berlin in the mid-19th century, extended upon Carnot's contributions by formulating the second law of thermodynamics. In his 1850 paper, *On the Moving Force of Heat*, he introduced the early concepts of the second law of thermodynamics.

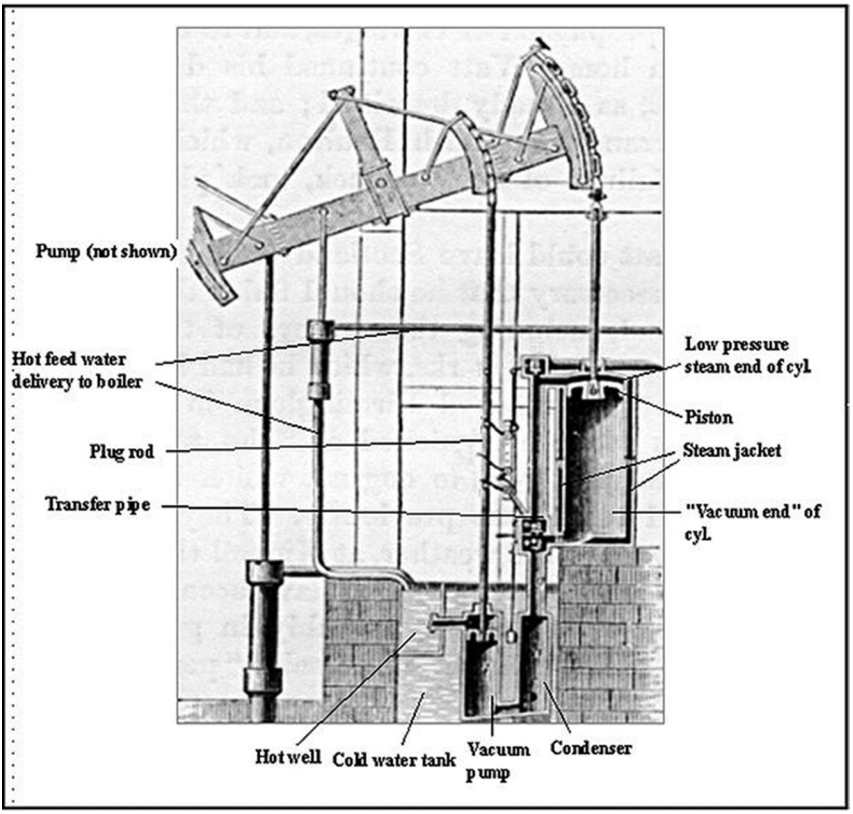

Figure 1.1 A James Watt steam engine, similar to those studied in Sadi Carnot's Reflections on Motive Power of Fire. (From Thurston, Robert H. English: A Schematic of Watt's Steam Engine Printed in a 1878 Book. 1878. Thurston, Robert H. History of the Growth of the Steam Engine. D. Appleton & Co. 1878. https://commons.wikimedia.org/wiki/File:Watt_steam_pumping_engine.JPG.)[11,12]

In 1865, Clausius gave this irreversible heat loss a name – entropy. The broad concepts are that, left on their own, in systems involving heat (and all do), everything gets cooler, and more generally in all systems, everything tends towards disorder. There is always loss in the system – not everything that goes into the system produces useful output or work.

Systems Loss in Societal Systems – Warfare

An example of systems loss within complex social systems is that of warfare. Because of its complexity, lessons learned in millennia of warfighting can offer some clues to planning for, implementing, and managing complex sociotechnical systems in complex societal groups such as cities and institutions.

> Everything in war is simple, but the simplest thing is difficult. The difficulties accumulate and end by producing a kind of friction that is inconceivable unless one has experienced war.
>
> **Carl Von Clausewitz**, *On War*

Carl von Clausewitz, the famous Prussian war theorist of the 19th century, introduced the concept of friction in war. He describes friction as "the force that makes the apparently easy so difficult."

This notion of friction being things that are "apparently easy" but that are actually difficult in reality can apply to large IoT systems implemented in institutions and cities. At face value, deploying one sensor in a location, routing data over a network, aggregating that data for subsequent processing, analysis, distribution, and consumption should be straightforward and easy. However, doing that 100, 1,000, or 10,000 times on a network (or networks) that is not has homogenous and predictable as originally thought, with resources less than originally anticipated, with vendor support that is not quite what was promised when the deal closed, and surprises, both small and large, will start to reveal cracks. These cracks, in turn, aggregate to "friction," with its detrimental effects on system implementation and management success. Continuing some of the metaphors established in the study of warfare,

> Friction may be self-induced, caused by such factors as lack of a clearly defined goal, lack of coordination, unclear or complicated plans, complex task organizations or command relationships, or complicated technologies.[9]

There are strong metaphors and parallels from which to learn in the study of warfare to implementing these complex technology systems in our complex societal organizations such as cities and institutions.

Another relevant metaphor to draw upon, per Clausewitz and modern military doctrine, is that in warfighting it is desirable to reduce friction, but there is no expectation that it will be eliminated. In fact, *effective military organizations plan and train to operate in the presence of friction, knowing that it cannot be entirely eliminated*.

Systems Loss in Management Systems

Another example appears in management theory. In introducing the idea of systems thinking in analyzing business operations, tactics, and strategy, Peter Senge says in *The Fifth Discipline*, "The irony is that, today, *the primary threats to our survival, both of our organizations, and of our societies, come not from sudden events but from slow, gradual processes*."[10] This is yet another example of nondescript systems loss that quietly aggregates and ultimately directly impacts the chances of IoT systems implementation success in terms of both ROI and cybersecurity posture.

Senge goes on to describe a "learning horizon" and our ability to observe and develop internal models of cause and effect.

"But what happens when we can no longer observe the consequences of our actions? What happens if the primary consequences of our actions are in the distant future or are in a distant part of the larger organization in which we operate? … When our actions have consequences beyond our learning horizon, it becomes impossible to learn from direct experience."[10] IoT systems also stretch across these large organizations – institutions and cities – and the comprehensive impact is never felt in any one place.

"We learn best from experience but we never directly experience the consequences of many of our most important decisions."[10] Similarly, because the full impact of IoT systems across cities and institutions is not comprehensively felt in any one place, learning from any particular IoT systems implementation is challenging.

This has strong parallels to where we are as institutions and cities. To add value and be competitive with other institutions and cities, we are rapidly buying and deploying IoT systems of increasing scale, and yet we really have limited basis for making good decisions about them. We're pressured into making decisions about deploying our nth IoT system, but we have little feedback from our early IoT systems deployments that can help influence that decision.

Systems Loss in IoT Systems

In IoT systems, loss occurs at the device level in IoT systems through unexpected installation costs, failures, misconfiguration costs, and others. Loss occurs at the network and network segment level through challenges of device enumeration and management, miscommunications between network managers (e.g., the central IT organization) and the IoT system owner around network segmentation and management, and others. Other losses occur in the organizational coordination level. Loss also occurs in resourcing. At each individual point or region, the loss may be small and not obvious. However, those losses aggregate and become substantial. Because of the rapid proliferation of IoT devices, the network segments supporting them, and the complexity of coordination of developing and maintaining organizational support resources, that system loss becomes significantly larger.

This systemic loss is insidious. It doesn't show up and knock on one's door advertising itself as loss or as a problem. Rather, it quietly, and initially imperceptibly, aggregates until system performance is substantially degraded, resourcing estimates prove themselves to be inadequate, uncertainty is significantly increased, cyber risk increases – perhaps greatly, and system manageability is degraded to the point of lost utility or failure. The scale and rate of growth of IoT devices, network segments, and systems greatly exacerbate this loss.

These analogies in other complex societal endeavors, such as warfighting, natural physical systems, and management organizations, can act as strong reminders and possibly even, in part, provide the basis for building partial models to help us frame the implementation of complex sociotechnical systems such as IoT in a substrate of complex societal organizations.

IoT Systems Manageability

Because systems loss is inevitable – not the least of which in complex IoT systems in complex societal organizations – developing approaches and methods to manage this loss, and most importantly, to be aware of this loss and learning to operate with that loss are critical to the success of IoT systems in cities and institutions. In turn, this is critical to the success of the smart cities or smart campus concept.

As leaders and administrators in institutions and cities, we can tend to get lost in the glitter and bling of potential promises of technology and not fully grasp the challenges of administering these complex systems.

Ironically, this is also true of some technologies that are offered to ostensibly mitigate IoT systems risk. The reason that these ostensible risk mitigation technologies are often ineffective is that they tend to vastly oversimplify the problem of IoT systems risk mitigation. These risk mitigation technologies often present themselves in neat technology packages and appear very convenient. While there are definitely technology components that can help with IoT systems risk mitigation, proposals or sales pitches that a single technology solution can address all IoT risk mitigation issues can be the equivalent of selling snake oil.

However, there are things that institutions and cities can do to mitigate this systems loss and increase the likelihood of successful IoT system implementation and operation. While there are many aspects to choosing, implementing, and operating IoT systems, a recurring thread and theme is that of *systems manageability*. By being aware of and demanding high degrees of systems manageability in the IoT systems that institutions and cities acquire, deploy, and operate, the opportunity for positive ROI's and non-degraded institutional cyber risk profiles are possible.

> Heeding and developing institutional maturity for identifying and demanding highly manageable IoT systems is critical for success.

It is one of the most important factors that cities and institutions can control in this rapidly increasing complexity of our societal and institutional environments, the exponential growth in the number of IoT devices, the systems that support and integrate them – all combined with the limited human, technical, and fiscal resources that almost all institutions and cities face that overwhelm our traditional approaches to enumerability, inventory, classification and categorization, and generally risk mitigation and management.

Systems manageability is perhaps the primary factor in addressing the slow bleed – that systems loss, the entropy – of the complex sociotechnical environments in which we live. Our societal/social systems are already complicated, but the integration and embeddedness of technology only expand that complexity. The manageability of an IoT system has direct and substantial implications on these limited human, technical, and fiscal resources.

There are many components and aspects to IoT systems manageability. The first step in discerning an IoT system's manageability is discerning the system's "knowability." What do we know about these systems that we are deploying in and around our institutions and cities? Do we just open the doors to our physical spaces and networks and let a third party install whatever, wherever? Or do we seek to know what is entering our physical and network environments? Discerning a system's knowability is a prerequisite to determining a system's manageability.

Some examples of these systems' knowability components and attributes include:

- Does the system have a name shared by all stakeholders?
- Who are the stakeholders in the consuming institution or city?
- Is there a primary point of contact within the institution/city for that system?
  - A coordinator at the vendor side?
  - A coordinator at the city/institution side?
- What are the expectations of data produced by the IoT system by *all of the stakeholders – more likely than not that there will be different expectations of data?*
- How well known is the system?
  - New vendor? Known vendor?
  - Degree of trust with vendor?
  - Documentation?
  - User and system support training?
  - How many endpoints, e.g., sensors and actuators, are there in the system?
    - Where are they? Do we know the location?
    - What is the IP address? The MAC address?
    - What is the current firmware version?
    - How are the devices updated/patched?
    - Are they updated/patched?
  - What are indications of health of these endpoint devices?
    - How does a healthy device present itself?
  - What is the central managing/controlling/aggregating application of these devices?
    - Does it sit on an on-premise server? Software-as-a-Service (SaaS)? Other?
    - What are the requirements of this application and supporting server?
    - How many servers?

- What are the desktop client applications involved with this application?
  - What are the requirements?
  - How many are there?
- Is there a risk agreement between the provider and the institutional consumer?
  - What is the mutual ownership between system success and failure?
  - Is there mutual ownership? Or is the client on their own?
- Others.

This is just a subset of IoT systems *knowability*. How much can we know about a complex IoT system in our complex institutional and city sociotechnical environments?

Capturing this information (and other information not enumerated here) is involved and resource-intensive.

We'd like to capture as much as we can within the bounds of our limited resources. But, importantly, if – when – we can't capture it all, *we want to know – and admit to ourselves – that we haven't captured it all and that capturing it all may indeed be impossible*. We acknowledge that we will be working with incomplete information and plan for this and work with this.

The institutional internal view and reflection that – though we seek to capture as much information as possible but acknowledge that we will always be working with partial and/or incomplete information – is a much better approach than extending the fantasy that we have complete knowledge of our systems. Making the assumption that we can or have captured, counted, and enumerated all aspects of the IoT system is a critically flawed basis for developing management and mitigation strategies for the institution or city.

We have to admit to ourselves as institution and city leaders that it is probable that many of *our sociotechnical IoT systems have become feasibly non-enumerable.* That is, we really don't know what's there. And, further, if we keep building management and risk mitigation approaches on this assumption that we can count everything, that we know where everything is, and we know what everything does – then we've got problems. Again, we seek to know as much as we can about our systems, but we must come to terms with the fact that we will not capture everything.

Given this tough state of affairs regarding just knowing the systems that we are deploying in our cities, institutions, and corporate campuses, discerning systems manageability – which has systems knowability as a requisite – is even harder.

Towards IoT systems *manageability*, within our institutions and cities – we want to use knowledge about systems, knowability, as a strong basis and then we want to look at internal resources available, skill sets available, market availability of skill sets, contract agreements, projected and actual vendor support, communication and coordination between and within institutional organizations and departments and others. This is no small task. Some components of IoT systems manageability include:

- Do the stakeholders, and their respective staffs, in city and institutional IoT system(s) have capacity to participate in the oversight, management, risk mitigation of these systems?
  - What current tasks and roles will they give up to participate in this work?
  - What skill sets do the staffs of the stakeholders need to support the stakeholders?
  - What skill sets are available in the institution or city to interpret the data so that it is actionable and usable in the context in which it is produced and consumed?
- What staff time is available to manage the performance of IoT systems vendor contracts?
  - What will these staff need to give up to manage these vendor system contracts?
- For the hundreds, thousands, or more of endpoints – devices deployed, who is going to support those devices?
  - What organization or department will support those devices?
  - Does that organization or department know that they are tasked with supporting these devices?
  - What is the existing budget and effort capacity to do this work?
    - Is it planned or is this a surprise?
    - What other work does the supporting organization give up to do this work?
    - Where is the resourcing – staff, funding, and scheduling – coming from to get this work done?
- Regarding systems applications
  - Who supports the server applications or SaaS applications?
  - Who establishes, adjusts, and maintains the configurability?
  - Who supports the client (e.g., desktop) applications?
  - Who trains the users?

Identifying *IoT systems manageability* as one of the top priorities, if not the top priority, for IoT systems selection, procurement, implementation, management, and even systems retirement – is essential. It is critical that systems-consuming institutions and cities see this and that *they create demand from IoT systems providers for more manageable IoT systems* and importantly that the provider/vendor shares the effectiveness of that objective with the consumer.

Ecosystem and Market – IoT Systems Consumers and Providers

IoT systems institutional and city consumers and their provider partners make up an ecosystem of IoT systems. While they are distinct entities in one sense, much of their operational lines are increasingly blurred between institutional/city

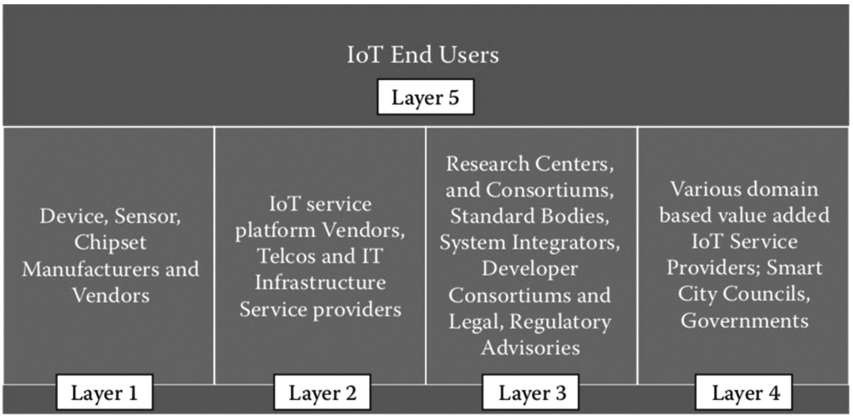

Figure 1.2 Complex IoT ecosystems. (Courtesy of Abhik Chaudhuri, Internet of Things, for Things, and by Things.)

consumers and their technology systems providers to include network boundaries, support agreements and roles, risk sharing, performance contracts, and others (Figure 1.2).

Because of this, IoT systems vendors must also be motivated by developing not just IoT systems, but also IoT systems services and support mechanisms. Selecting, deploying, and managing IoT systems are still nascent for most cities and institutions. Mistakes will be made. One of the key places where mistakes are being made is estimating and planning for the resources and skill sets required to deploy and operate these systems. While this is problematic, over time, these institutional consumers will learn from these mistakes, mature in their IoT systems acquisition approach, and begin demanding better systems. In particular, the demand for better IoT systems will drive the demand for better IoT systems manageability.

Organizations seek to be aware of, manage, and reduce this systems loss by selecting for more manageable systems – in both the technical and the sociotechnical senses – and will see much greater value add of their IoT systems than those of their peers or competitors.

Similarly, those IoT systems vendors and providers acknowledge this guaranteed systems loss within a consumer city or organization and help them reduce that loss through increased systems manageability – while acknowledging that there will always be some loss – *will outperform and outlive their competitors.*

Because of the unbridled potential of new IoT systems coupled with our lack of experience as city and institutional consumers for establishing performance expectations, vetting criteria, and IoT systems deployment and operational experience, there are many, many IoT systems on the market that have very limited value and, in fact, can have negative value by causing lost investment and degrading a city's

or institution's cybersecurity posture. In a nutshell, there is a lot of IoT systems garbage on the market.

In time, some of these ill-conceived, poorly defined, and poorly supported systems will shake out as being not useful and possibly harmful. But this will take some years, and along the way, damage will be done in terms of lost ROI and degraded cyber risk profiles.

Those IoT systems providers that can grasp the complexity of matching the institutional consumer need for the system, along with comprehending the institutional consumer's complex technical and social/societal substrate (often including working with and within a large bureaucracy) – into/onto which that IoT system is deployed – will have a substantial market advantage and will bring greater impact over time.

Approach

This book will focus on the relationship between this technology and that of major societal organizations such as cities and universities. These major societal organizations have a duty and obligation to serve, protect, and enhance the lives of the people that live and work within them. As such, one of their obligations is to identify and seek to manage and mitigate risk to their constituencies and organizational structures.

The book is composed of chapters that cover various aspects of IoT systems and risk mitigation and cybersecurity around the same. Within each chapter, a particular aspect or phenomena of IoT systems in cities and institutions will be discussed. The intent is to provide some language and conceptual frameworks for the issue. Our shared language about these IoT systems must evolve and do so quickly, if we are to successfully manage these systems and manage and mitigate risk around the same. Some chapters will also include proposed mitigation steps and actions. None of these approaches are written in stone, and there are many ways to accomplish the objectives, but these will be worthy of consideration for your own city or institution. Similarly, commercial providers of IoT systems products and services can create competitive advantage by helping institutions and cities solve these complex challenges.

Chapter 2, "Differences between IoT and Traditional IT Systems," discusses how IoT systems are different from traditional enterprise IT systems within institutions and cities.

Chapter 3, "Defining IoT Systems Implementation Success," provides criteria for analyzing the success (or not) of IoT systems implementations.

Chapter 4, "Systems of Systems and Sociotechnical Systems," introduces IoT systems as both systems of systems and sociotechnical systems.

Chapter 5, "Systems Seams, Boundaries, and the IoT Ecosystem," discusses systems losses at organizational boundaries and how those losses aggregate in

resource-constrained environments, such as most institutions and cities, to directly impact the likelihood of IoT systems implementation success.

Chapter 6, "IoT Systems Manageability," describes systems manageability in more detail, and the positive impact of strong systems manageability can have mitigating systems losses in resource-constrained environments in institutions and cities.

Chapter 7, "IoT Systems Vendor Relations & Vendor Management," covers the critical relationship of the city or institution and the IoT systems vendor or provider.

Chapter 8, "Templates for Institutional & City IoT Systems Planning & Operations," offers some templates for planning and implementing IoT systems.

Chapter 9, "Strategy Implementation," presents a high-level strategy for selecting, deploying, and managing IoT systems within the institution or city.

Suggested Reading

1. Cleveland, Robin, and Carolyn Bartholomew. Hearing on China, the United States, and next generation connectivity. U.S.-CHINA ECONOMIC AND SECURITY REVIEW COMMISSION, 2018, 132.
2. Chaudhuri, Abhik. *Internet of Things, for Things, and by Things*. Boca Raton, FL: CRC Press/Taylor & Francis Group, 2019.
3. sataksig. Perpetual Motion Machines and the Search of Free Energy. Earth Buddies (blog), November 26, 2017. https://earthbuddies.net/perpetual-motion-machines-and-the-search-of-free-energy/.
4. Second Law of Thermodynamics. In *Wikipedia*, December 25, 2018. https://en.wikipedia.org/w/index.php?title=Second_law_of_thermodynamics&oldid=875356570.
5. Carnot, Lazare. In *Wikipedia*, November 20, 2018. https://en.wikipedia.org/w/index.php?title=Lazare_Carnot&oldid=869870374.
6. Carnot, Sadi. "Heat Engine" *Reflections on the Motive Power of Fire: And Other Papers on the Second Law of Thermodynamics*, by É. Clapeyron and R. Clausius. Mineola, NY: Dover Publ., 2009.
7. Carnot Cycle. Accessed January 3, 2019. http://hyperphysics.phy-astr.gsu.edu/hbase/thermo/carnot.html; https://en.wikipedia.org/wiki/Carnot_cycle.
8. Clausius, Rudolf. In *Wikipedia*, November 13, 2018. https://en.wikipedia.org/w/index.php?title=Rudolf_Clausius&oldid=868579650.
9. Warfighting, Mcdp1.Pdf. Accessed January 3, 2019. https://clausewitz.com/readings/mcdp1.pdf.
10. Senge, Peter M. *The Fifth Discipline: The Art and Practice of the Learning Organization*. Rev. and updated. New York: Doubleday/Currency, 2006.
11. Thurston, Robert H. *English: A Schematic of Watt's Steam Engine Printed in a 1878 Book*. 1878. https://commons.wikimedia.org/wiki/File:Watt_steam_pumping_engine.JPG
12. Thurston, Robert H. *History of the Growth of the Steam Engine*. D. Appleton & Co. 1878. https://commons.wikimedia.org/wiki/File:Watt_steam_pumping_engine.JPG.