Networks of ‘Things’

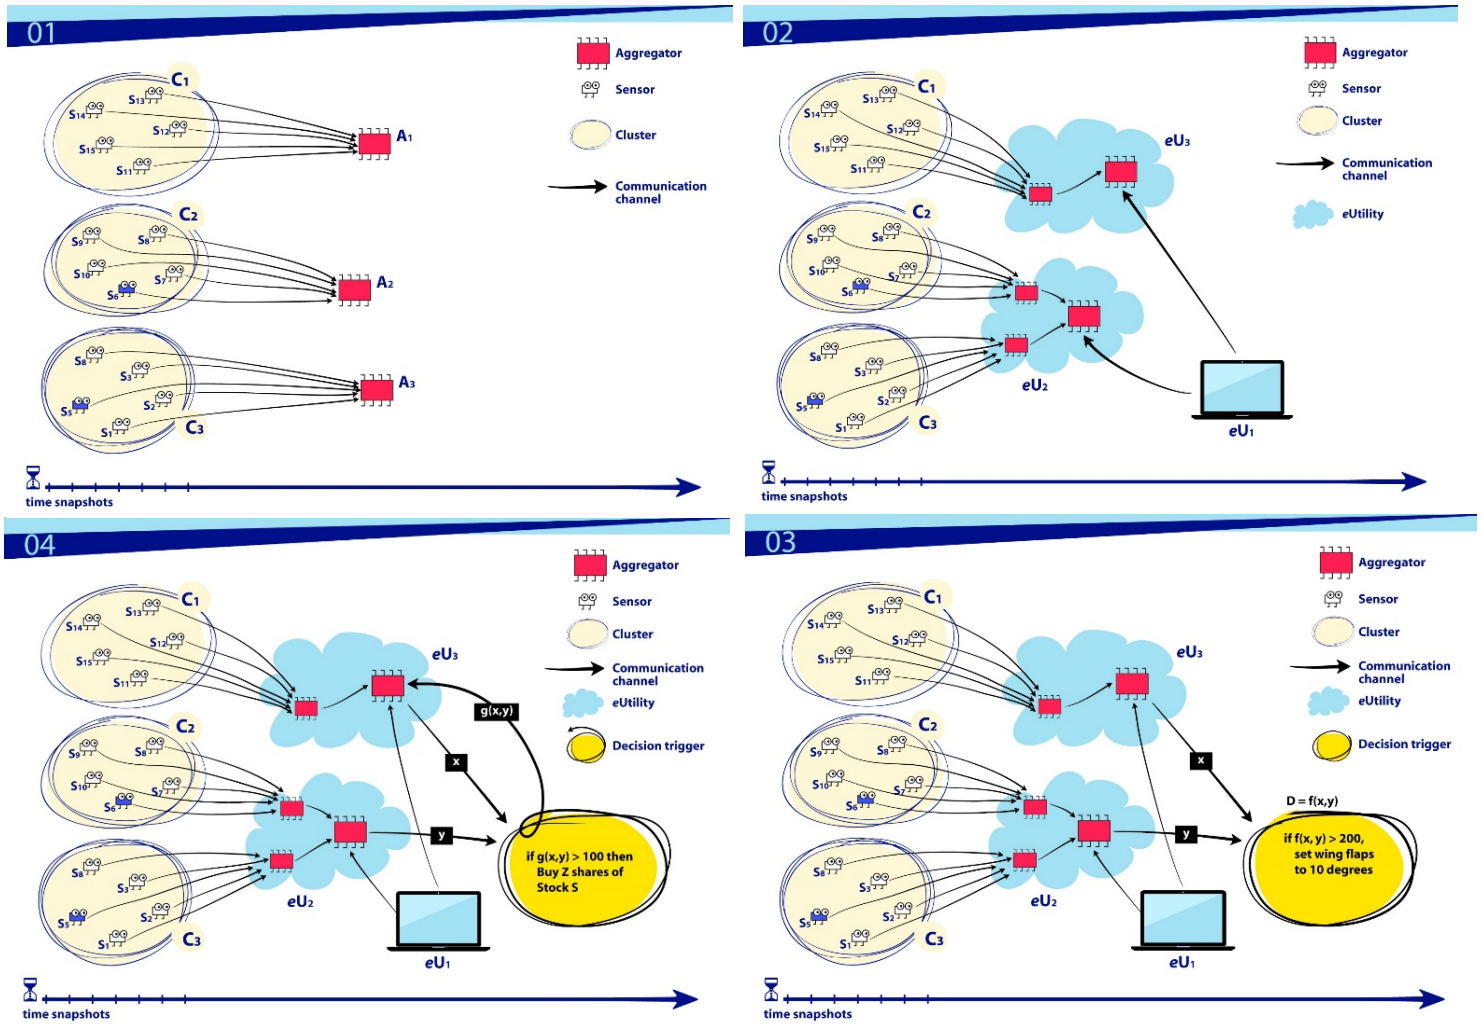

J. Voas

Computer Scientist

National Institute of Standards and Technology

18 Months Ago We Asked

What is IoT*?*

"***The term Internet of Things generally refers to scenarios where network connectivity and computing capability extends to objects, sensors and everyday items not normally considered computers, allowing these devices to generate, exchange and consume data with minimal human intervention.***"

The Internet of Things (IoT): An Overview, Karen Rose, et.al. The Internet Society, October 2015. p. 5.

"***Although there is no single definition for the Internet of Things, competing visions agree that it relates to the integration of the physical world with the virtual world – with any object having the potential to be connected to the Internet via short-range wireless technologies, such as radio frequency identification (RFID), near field communication (NFC), or wireless sensor networks (WSNs). This merging of the physical and virtual worlds is intended to increase instrumentation, tracking, and measurement of both natural and social processes.***"

"Algorithmic Discrimination: Big Data Analytics and the Future of the Internet", Jenifer Winter. In: *The Future Internet: Alternative Visions*. Jenifer Winter and Ryota Ono, eds. Springer, December 2015. p. 127.

*"**Industrial Internet of Things (IOT) is a distributed network of smart sensors that enables precise control and monitoring of complex processes over arbitrary distances**."*

"Ensuring trust and security in the industrial IoT", Bernardo A. Huberman. *Ubiquity: An ACM Publication*, January 2016, p. 1.

"*The concept of Internet of Things (IOT) … is that every object in the Internet infrastructure is interconnected into a global dynamic expanding network.*"

"An efficient user authentication and key agreement scheme for heterogeneous wireless sensor network tailored for the Internet of Things environment", Mohammad Sabzinejad Farasha, et.al. *Ad Hoc Networks* 36(1), January 2016. p. Abstract

"***In what's called the Internet of Things, sensors and actuators embedded in physical objects—from roadways to pacemakers—are linked through wired and wireless networks, often using the same Internet Protocol (IP) that connects the Internet.*"**

"The Internet of things", M. Chui, M. Löffler, and R. Roberts. McKinsey Quarterly, Sept. 23, 2015. As cited in: "Control Systems and the Internet of Things", Tariq Samad. *IEEE Control Systems Magazine*, 36(1), February 2016. p. 14.

"The main idea behind the IoT is to bridge the gap between the physical world of humans and the virtual world of electronics via smart objects. These smart objects allow the interactions between humans and their environment by providing, processing, and delivering any sort of information or command. Sensors and actuators will be integrated in buildings, vehicles, and common environments and can tell us about them, their state, or their surroundings."

"Using an Epidemiological Approach to Maximize Data Survival in the Internet of Things", Abdallah Makhoul, et.al. *ACM Transactions on Internet Technology*, 16(1), February 2016. p. 5.

"We must first define what we mean by 'things.' It could be very simple objects or complex objects. Things do not need to be connected directly to the public Internet, but they must be connectable via a network (which could be a LAN, PAN, body area network, etc.). The IoT is the network of physical objects that contain embedded technology to communicate and interact with the external environment. The IoT encompasses hardware (the 'things' themselves), embedded software (software running on, and enabling, the connected capabilities of the things), connectivity/communications services, and information services associated with the things (including services based on analysis of usage patterns and sensor or actuator data). An IoT solution is a product (or set of products) combined with a service either a one-to-one or a one-to-many relation. Meaning one service is combined with one (set of) product(s), or one service is combined with multiple (sets of) products."

"Internet of Things in Energy Efficiency", Francois Jammes. *Ubiquity: An ACM Publication*, February 2016, p. 2

"At the very high level of abstraction, the Internet of Things (IoT) can be modeled as the hyper-scale, hyper-complex cyber-physical system."

"On resilience of IoT systems: the internet of things". Kemal A. Delic. *Ubiquity: An ACM Publication*, February 2016, pp. 1.

"The Internet of Things (IoT) paradigm is based on intelligent and self-configuring nodes (things) interconnected in a dynamic and global network infrastructure."

"Integration of Cloud Computing and Internet of Things: A Survey", Alessio Botta, et.al. *Future Generation Computer Systems*, Vol. 56, March 2016, p. 2.

"The Internet of Things (IoT)…connecting everyday objects to the Internet and facilitating machine-to-human and machine-to-machine communication with the physical world."

"When things matter: A survey on data-centric internet of things", Yongrui Qin, et.al. *Journal of Network and Computer Applications*, Vol. 64, April 2016. p. Abstract

"Whilst the definition of 'Internet of Things' is elusive in general, the use of the term refers to the use of sensors and data communications technology built into physical objects in order to track, coordinate or control the functioning of those objects based on data over the network or the Internet."

"National Security in a Hyper-Connected World: Global Interdependence and National Security", Christian O. Fjader. In: *Exploring the Security Landscape: Non-Traditional Security Challenges*. Anthony J. Masys, Ed. Springer, 2016. p. 33

"The internet of things is a new paradigm in which every device is digitally connected, regardless of their function, and can communicate with other devices and people over communication protocols. "

"Sensorization to Promote the well-being of people and the betterment of health organizations", Fabio Silva and Cesar Analide. In: *Applying Business Intelligence to Clinical and Healthcare Organizations*. José Machado and António Abelha, Eds. Medical Information Science Reference, 2016. p. 117

"The Internet of Things is a term used to describe the ever-growing number of devices connecting to a network, including televisions and appliances."

Web Design with HTML & CSS3: Comprehensive. Jessica Minnick and Lisa Friedrichsen. Course Technology, 8th edition, 2016. p. HTML 4

"…the interconnectness of all systems through the internet [is known as] ***'the internet of things'."***

"Risk management and cyber risk in the financial services sector", Ruth Taplin. In: *Managing Cyber Risk in the Financial Sector: Lessons from Asia, Europe, and the USA*. Ruth Taplin, Ed. Routledge, 2016. p. 15

"The Internet of Things (IoT) envisions a world where smart objects connected to the Internet, share their data, exchange their services and cooperate together to provide value-added services that none of these objects could provide individually."

"Event-Aware Framework for Dynamic Services Discovery and Selection in the Context of Ambient Intelligence and Internet of Things" A. Yachir, *IEEE Transactions on Automation Science and Engineering*, 13(1), 2016. p. 85.

"Although many standardization groups such as IEEE, ITU, 3GPP, and IETF have presented various definitions, in its broadest sense, Internet of the Things means 'technology through which additional values can be provided to users by linking things or devices to the Internet.'"

"A Study on Actual Cases & Meanings for Internet of Things", Dong-Woo Lee.

International Journal of Software Engineering and Its Applications, 10(1), 2016, p. 287.

Organizational Definitions

Definition / Text	Link
IEEE IoT Initiative: Towards a definition of the Internet of Things (IoT), SEE Chapter 5 for IEEE IoT definition.	http://iot.ieee.org/images/files/pdf/IEEE_IoT_Towards_Definition_Internet_of_Things_Revision1_27MAY15.pdf
European Research Cluster on IoT (IERC): "A dynamic global network infrastructure with self-configuring capabilities based on standard and interoperable communication protocols where physical and virtual "things" have identities, physical attributes, and virtual personalities and use intelligent interfaces, and are seamlessly integrated into the information network."	http://www.internet-of-things-research.eu/about_iot.htm
ITU: The Internet of Things (IoT) has been defined in Recommendation ITU-T Y.2060 (06/2012) as "a global infrastructure for the information society, enabling advanced services by interconnecting (physical and virtual) things based on existing and evolving interoperable information and communication technologies."	http://www.itu.int/en/ITU-T/gsi/iot/Pages/default.aspx
OASIS: "System where the Internet is connected to the physical world via ubiquitous sensors." OASIS describes the ubiquity of sensors as existing in "every mobile, every auto, every door, every room, every part, on every parts list, every sensor in every device in every bed, chair or bracelet in every home, office, building or hospital room in every city and village on Earth."	https://www.oasis-open.org/presentations/open-protocols-and-internet-of-things-oasis.ppt

Definition / Text	Link
W3C: "The Web of Things includes sensors and actuators, physical objects and locations, and even people. The Web of Things is essentially about the role of Web technologies to facilitate the development of applications and services for things and their virtual representation. Some relevant Web technologies include HTTP for accessing RESTful services, and for naming objects as a basis for linked data and rich descriptions, and JavaScript APIs for virtual objects acting as proxies for real-world objects."	https://www.w3.org/community/wot/wiki/Main_Page (link in IEEE document is not working)
Wikipedia: The Internet of Things (IoT) is the network of physical objects—devices, vehicles, buildings and other items embedded with electronics, software, sensors, and network connectivity—that enables these objects to collect and exchange data.[1] The Internet of Things allows objects to be sensed and controlled remotely across existing network infrastructure,[2] creating opportunities for more direct integration of the physical world into computer-based systems, and resulting in improved efficiency, accuracy and economic benefit;[3][4][5][6][7][8] when IoT is augmented with sensors and actuators, the technology becomes an instance of the more general class of cyber-physical systems, which also encompasses technologies such as smart grids, smart homes, intelligent transportation and smart cities. Each thing is uniquely identifiable through its embedded computing system but is able to interoperate within the existing Internet infrastructure. Experts estimate that the IoT will consist of almost 50 billion objects by 2020.[9]	https://en.wikipedia.org/wiki/Internet_of_Things
Gartner: The Internet of Things (IoT) is the network of physical objects that contain embedded technology to communicate and sense or interact with their internal states or the external environment.	http://www.gartner.com/it-glossary/internet-of-things/

Definition / Text	Link
Webopedia: The Internet of Things (IoT) refers to the ever-growing network of physical objects that feature an IP address for internet connectivity, and the communication that occurs between these objects and other Internet-enabled devices and systems. IoT Extends Internet Connectivity: The Internet of Things extends internet connectivity beyond traditional devices like desktop and laptop computers, smartphones and tablets to a diverse range of devices and everyday things that utilize embedded technology to communicate and interact with the external environment, all via the Internet.	http://www.webopedia.com/TERM/I/internet_of_things.html
Techopedia: The Internet of Things (IoT) is a computing concept that describes a future where everyday physical objects will be connected to the Internet and be able to identify themselves to other devices. The term is closely identified with RFID as the method of communication, although it also may include other sensor technologies, wireless technologies or QR codes. The IoT is significant because an object that can represent itself digitally becomes something greater than the object by itself. No longer does the object relate just to you, but is now connected to surrounding objects and database data. When many objects act in unison, they are known as having "ambient intelligence."	https://www.techopedia.com/definition/28247/internet-of-things-iot
The Internet Engineering Task Force (IETF): IETF provides its own description of IoT, along with definitions for "Internet" and "thing" (IETF, "Internet of Things," 2010): "The basic idea is that IoT will connect objects around us (electronic, electrical, non-electrical) to provide seamless communication and contextual services provided by them. Development of RFID tags, sensors, actuators, mobile phones make it possible to materialize IoT which interact and co-operate each other to make the service better and accessible anytime, from anywhere." IETF's definition of "Internet": "The original 'Internet' is based on the TCP/IP protocol suite but any network based on the TCP/IP protocol suite cannot belong to the Internet because private networks and telecommunication networks are not part of the Internet even though they are based on the TCP/IP protocol suite. In the viewpoint of IoT, the 'Internet' considers the TCP/IP suite and nonO TCP/IP suite at the same time." IETF's definition of "things": "In the vision of IoT, 'things' are very various such as computers, sensors, people, actuators, refrigerators, TVs, vehicles, mobile phones, clothes, food, medicines, books, etc. These things are classified as three scopes: people, machine (for example, sensor, actuator, etc.) and information (for example, clothes, food, medicine, books, etc.). These 'things' should be identified at least by one unique way of identification for the capability of addressing and communicating with each other and verifying their identities. In here, if the 'thing' is identified, we call it the 'object.'"	Could not find the source – the following is from the IEEE document http://iot.ieee.org/images/files/pdf/IEEE_IoT_Towards_Definition_Internet_of_Things_Revision1_27MAY15.pdf

So?

Economics

- “IDC expects the worldwide market for IoT solutions to grow at a 20% CAGR from $1.9 trillion in 2013 to $7.1 trillion in 2020.”
- “By 2025, Internet of things applications could have $11 trillion impact”
- “Gartner Says By 2020, More Than Half of Major New Business Processes and Systems Will Incorporate Some Element of the Internet of Things”

Economics

IT Spending Cuts Spare the Internet of Things
http://blogs.wsj.com/cio/2016/02/09/it-spending-cuts-spare-the-internet-of-things/

"Facing a downturn in information-technology spending, some enterprise-technology firms are betting on growth in the emerging Internet-of-Things market to cushion declines elsewhere, industry analysts say."

Reports on Economics

Source	Link
Forbes, Internet Of Things (IoT) Predictions From Forrester, Machina Research, WEF, Gartner, IDC, January, 2016	http://www.forbes.com/sites/gilpress/2016/01/27/internet-of-things-iot-predictions-from-forrester-machina-research-wef-gartner-idc/#4b1601546be6
IDC, IDC FutureScape: Worldwide Internet of Things 2016 Predictions, November 2015	https://www.idc.com/research/viewtoc.jsp?containerId=259856
IDC Report, 2014	http://www.business.att.com/content/article/IoT-worldwide_regional_2014-2020-forecast.pdf
Gartner Says By 2020, More Than Half of Major New Business Processes and Systems Will Incorporate Some Element of the Internet of Things, January 14, 2016	http://www.gartner.com/newsroom/id/3185623
McKinsey Global Institute, Preparing IT systems and organizations for the Internet of Things, November, 2015	http://www.mckinsey.com/industries/high-tech/our-insights/preparing-it-systems-and-organizations-for-the-internet-of-things
The Internet of Things: Five critical questions August 2015	http://www.mckinsey.com/industries/high-tech/our-insights/the-internet-of-things-five-critical-questions
By 2025, Internet of things applications could have $11 trillion impact, July 22, 2015	http://www.mckinsey.com/mgi/overview/in-the-news/by-2025-internet-of-things-applications-could-have-11-trillion-impact
Unlocking the potential of the Internet of Things, June 2015	http://www.mckinsey.com/business-functions/business-technology/our-insights/the-internet-of-things-the-value-of-digitizing-the-physical-world
Wall Street Journal, Measuring the Economic Potential of the Internet of Things, July, 2015	http://blogs.wsj.com/cio/2015/07/17/measuring-the-economic-potential-of-the-internet-of-things/

Reality

No universally-accepted and actionable definition has existed to the question, “What is IoT?”

This presentation is NIST’s attempt

National Institute of Standards and Technology

Draft NISTIR 8063

Primitives and Elements of Internet of Things (IoT) Trustworthiness

Jeffrey Voas
Computer Security Division
Information Technology Laboratory

February 2016

U.S. Department of Commerce
Penny Pritzker, Secretary

National Institute of Standards and Technology
Willie May, Under Secretary of Commerce for Standards and Technology and Director

NIST Special Publication 800-183

Networks of 'Things'

Jeffrey Voas

This publication is available free of charge from:
http://dx.doi.org/10.6028/NIST.SP.800-183

COMPUTER SECURITY

NIST
National Institute of
Standards and Technology
U.S. Department of Commerce

http://dx.doi.org/10.6028/NIST.SP.800-183

Opening Statement

‘This technology’ (IoT) employs a mixture of *sensing*, *communication*, *computation, actuation*.

We stepped back from the acronym IoT

Network of Things (NoT)

IoT *vs.* NoT

We use two acronyms, IoT and NoT (Network of Things), extensively and interchangeably—the relationship between NoT and IoT is subtle. IoT is an instantiation of a NoT, more specifically, IoT has its 'things' tethered to the Internet. A different type of NoT could be a Local Area Network (LAN), with none of its 'things' connected to the Internet. Social media networks, sensor networks, and the Industrial Internet are all variants of NoTs. This differentiation in terminology provides ease in separating out use cases from varying vertical and quality domains (e.g., transportation, medical, financial, agricultural, safety-critical, security-critical, performance-critical, high assurance, to name a few). That is useful since there is no singular IoT, and it is meaningless to speak of comparing one IoT to another.

Primitives

1. **Sensor** A *sensor* is an electronic utility that measures physical properties such as temperature, acceleration, weight, sound, location, presence, identity, etc. All sensors employ mechanical, electrical, chemical, optical, or other effects at an interface to a controlled process or open environment
2. **Aggregator** An *aggregator* is a software implementation based on mathematical function(s) that transforms groups of *raw* data into *intermediate, aggregated* data. Raw data can come from any source. Aggregators address 'big' data challenges.
3. **Communication channel** A *communication channel* is a medium by which data is transmitted (e.g., physical via USB, wireless, wired, verbal, etc.).
4. ***e*Utility** An *e*Utility (external utility*)* is a software or hardware product or service.
5. **Decision trigger** A *decision* trigger creates the final result(s) needed to satisfy the purpose, specification, and requirements of a specific NoT. This is typically actuation or a transaction.

For Each Primitive

Basic properties, assumptions, recommendations, and general statements about Primitive x include:

Sensor (10 of 29)

1. Sensors are physical; some may have an Internet access capability.
2. A sensor may also transmit device identification information, such as via RFID
3. Sensors may be heterogeneous, from different manufacturers, and collect data, with varying levels of data integrity.
4. Sensors may be associated with fixed geographic locations or may be mobile.
5. Sensors may have an owner(s) who will have control of the data their sensors collect, who is allowed to access it, and when.
6. Sensors will have pedigree – geographic locations of origin and manufacturers. Pedigree may be unknown, and suspect. This has ties to Supply Chain Risk Management (SCRM).
7. Sensors may be cheap, disposable, and susceptible to wear-out over time.
8. There will differentials in sensor security, safety, and reliability, e.g., between consumer grade, military grade, industrial grade, etc.
9. Sensors may return no data, totally flawed data, partially flawed data, or correct/acceptable data. Sensors may fail completely or intermittently. They may lose sensitivity or calibration. Sensors may have their data encrypted.
10. Security is a concern for sensors if they or their data is tampered with, stolen, deleted, dropped, or transmitted insecurely so it can be accessed by unauthorized parties. Building security into specific sensors may or may not be cost effective.

Aggregator (5 of 11)

1. Aggregators may be virtual due the benefit of changing implementations quickly and increased malleability. A situation may exist where aggregators are physically manufactured, e.g., a field-programmable gate array (FPGA) or hard-coded aggregator that is not programmable. Aggregators may also act in a similar way as *n*-version voters.
2. Intermediate, aggregated data may suffer from some level of *information loss*. Proper care in the aggregation process should be given to significant digits, rounding, averaging, and other arithmetic operations to avoid unnecessary loss of precision.
3. Aggregators are: (1) executed at a specific time and for a fixed time interval, or (2) event-driven.
4. Security is a concern for aggregators (malware or general defects) and for the sensitivity of their aggregated data. Further, aggregators could be attacked, e.g., by denying them the ability to operate/execute or by feeding them bogus data.
5. Aggregators have two actors for consolidating large volumes of data into lesser amounts: Clusters and Weights. This is the only primitive with actors.

Actor: Cluster (6 of 7)

1. Clusters are abstractions of a set of sensors along with the data they output—clusters may be created in an *ad hoc* manner or organized according to fixed rules.
2. Clusters are not inherently physical.
3. C_i may share one or more sensors with C_k, where i ≠ k, or with other NoTs. This is somewhat important, because competing organizations may be receiving data that they believe to be unique and purposed only for them to receive, and not realizing a competitor is also receiving the same sensor data.
4. Clusters are malleable and can change their collection of sensors and their data at any time.
5. *Continuous-binding* of a sensor to a cluster may result in little ability to mitigate trustworthiness concerns of a real-time NoT if the binding occurs *late*
6. The composition of clusters is dependent on what mechanism is employed to aggregate the data, which ultimately impacts the purpose and requirements of a specific NoT.

Actor: Weight (6 of 9)

1. A weight may be hardwired or modified on-the-fly.
2. A weight may be based on a sensor's perceived trustworthiness, e.g., based on who is the sensor's owner, manufacturer, geographic location of manufacture, geographic location where the sensor is operating, sensor age or version, previous failures or partial failures of sensor, sensor tampering, sensor delays in returning data, etc. A weight may also be based on the value of the data, uniqueness, relation to mission goals, etc.
3. Different NoTs may leverage the same sensor data and re-calibrate the weights per the purpose of a specific NoT.
4. Aggregators may employ artificial intelligence techniques to modify their clusters and weights on-the-fly.
5. Weights will affect the degree of information loss during the creation of intermediate data.
6. Security concerns for weights is related to possible tampering of the weights.

Communication Channel (8 of 12)

1. Communication channels move data between computing, sensing, and actuation.
2. Since data is the "blood" of a NoT, communication channels are the "veins" and "arteries", as data moves to and from intermediate events at different snapshots in time.
3. Communication channels will have a physical or virtual aspect to them, or both. Protocols and associated implementations provide a virtual dimension, cables provide a physical dimension.
4. Communication channels may be wireless
5. Communication channel dataflow may be unidirectional or bi-directional. There are a number of conditions where an aggregator might query more advanced sensors, or potentially recalibrate them in some way (e.g., request more observations per time interval).
6. No standardized communication channel protocol is assumed; a specific NoT may have multiple communication protocols between different entities.
7. Communication channels are prone to disturbances and interruptions.
8. *Redundancy* can improve communication channel reliability. There may be more than one distinct communication channel between a computing primitive and a sensing primitive.

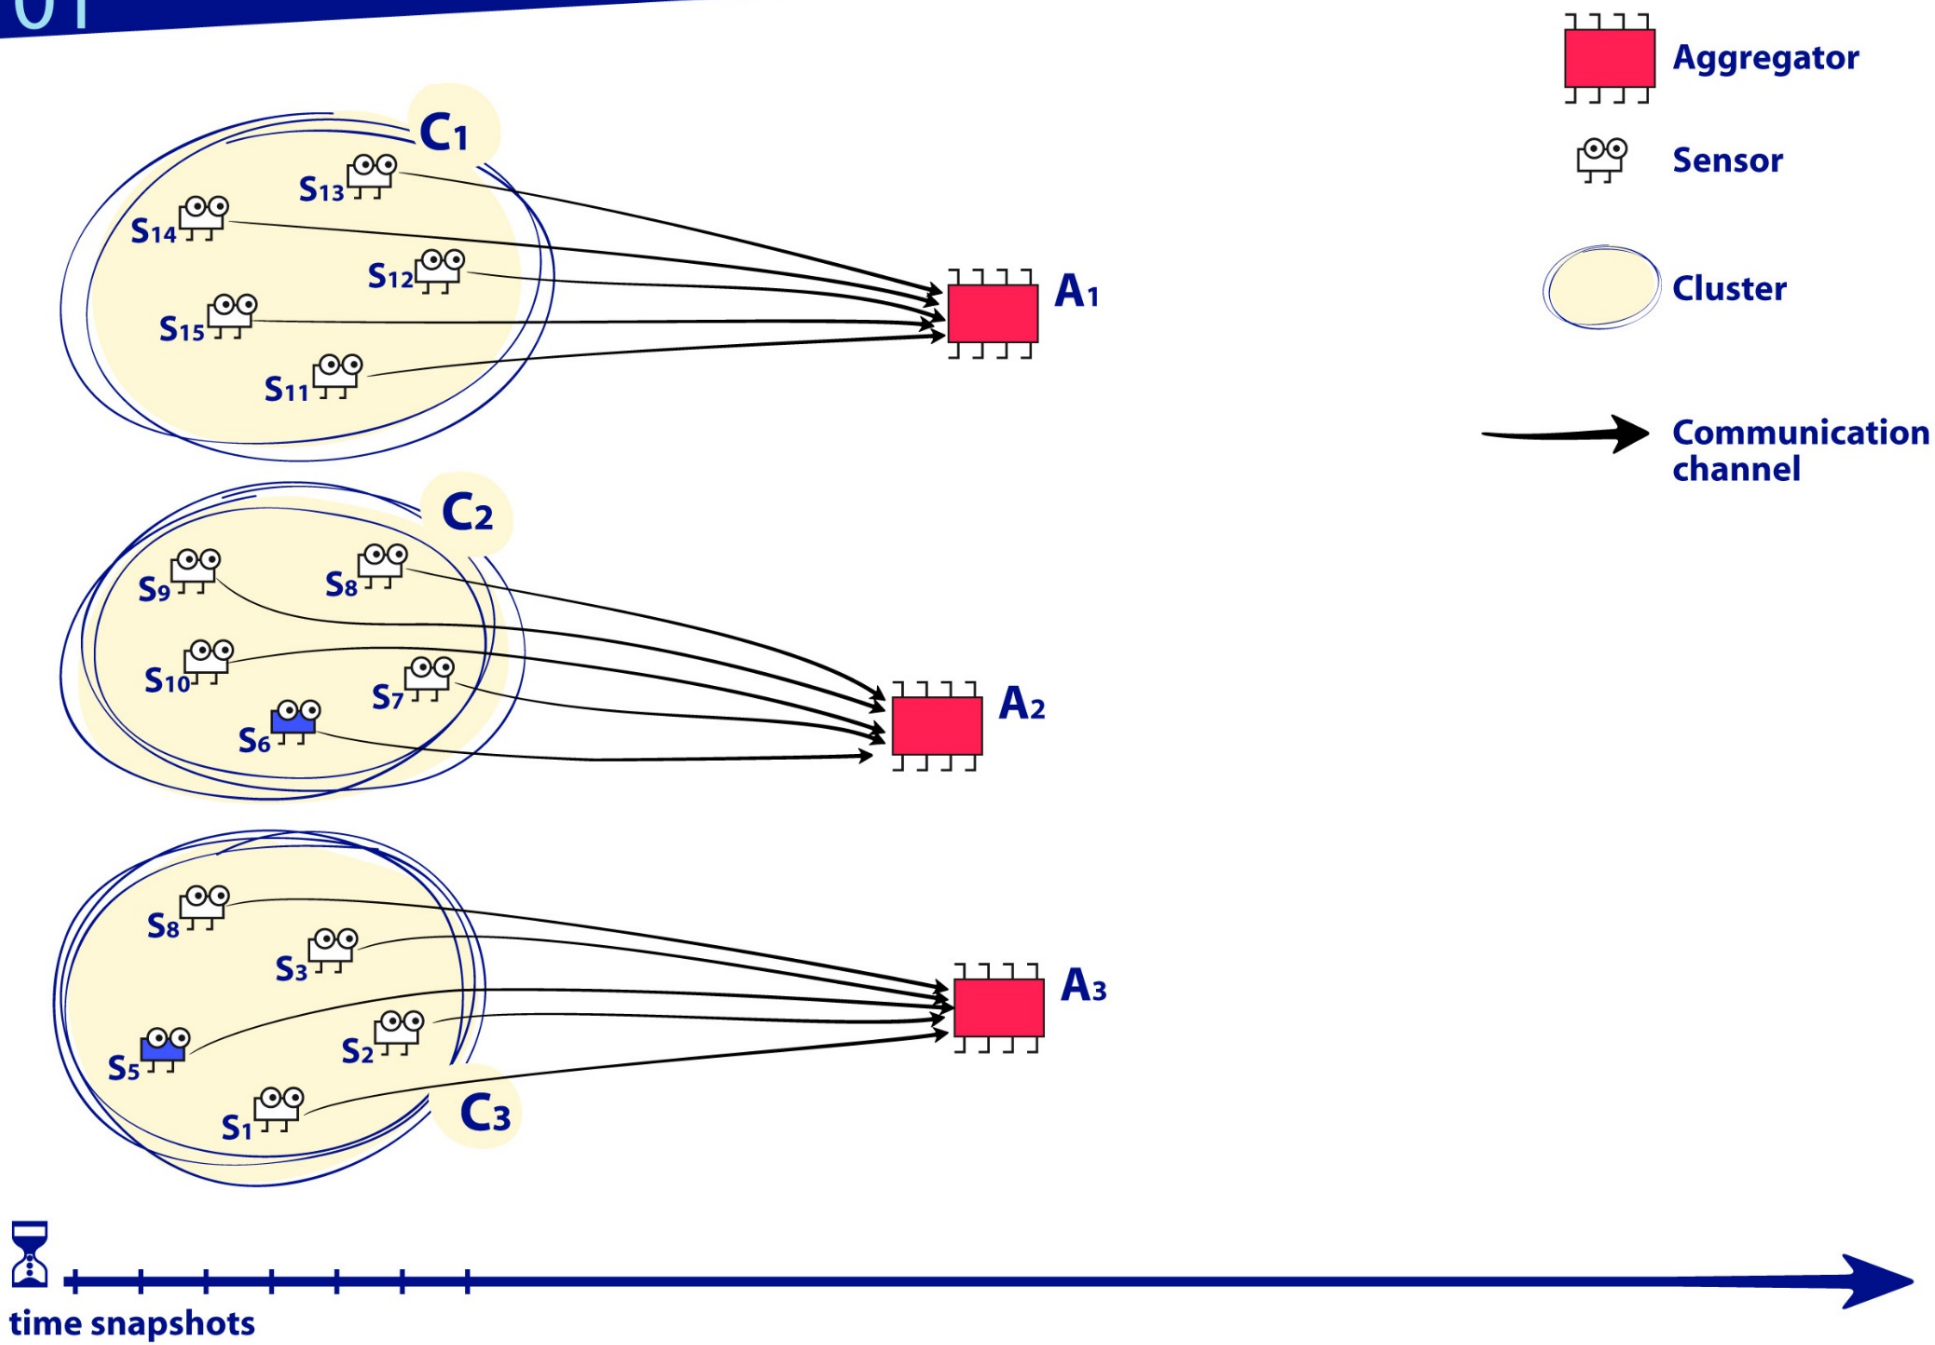
Aggregator
Sensor
Cluster
Communication channel
C1
S13
S14
S12
S15
S11
A1
C2
S9
S8
S10
S7
S6
A2
S8
S3
S5
S2
S1
C3
A3
time snapshots

*e*Utility (6 of 9)

1. *e*Utilities may include databases, mobile devices, misc. software or hardware systems, clouds, computers, CPUs, etc. The *e*Utility primitive can be subdivided, and probably should be decomposed to make this primitive less abstract.
2. *e*Utilities execute processes or feed data into the overall workflow of a NoT.
3. *e*Utilities will likely be acquired off-the-shelf from 3rd parties.
4. *e*Utilities, such as clouds, provide computing power that aggregators may not have.
5. A human may be viewed as a *e*Utility. A human is sometimes referred to as a 'thing' in public IoT discourse.
6. Non-human *e*Utilities may have Device_IDs; Device_IDs may be crucial for identification and *authentication*.

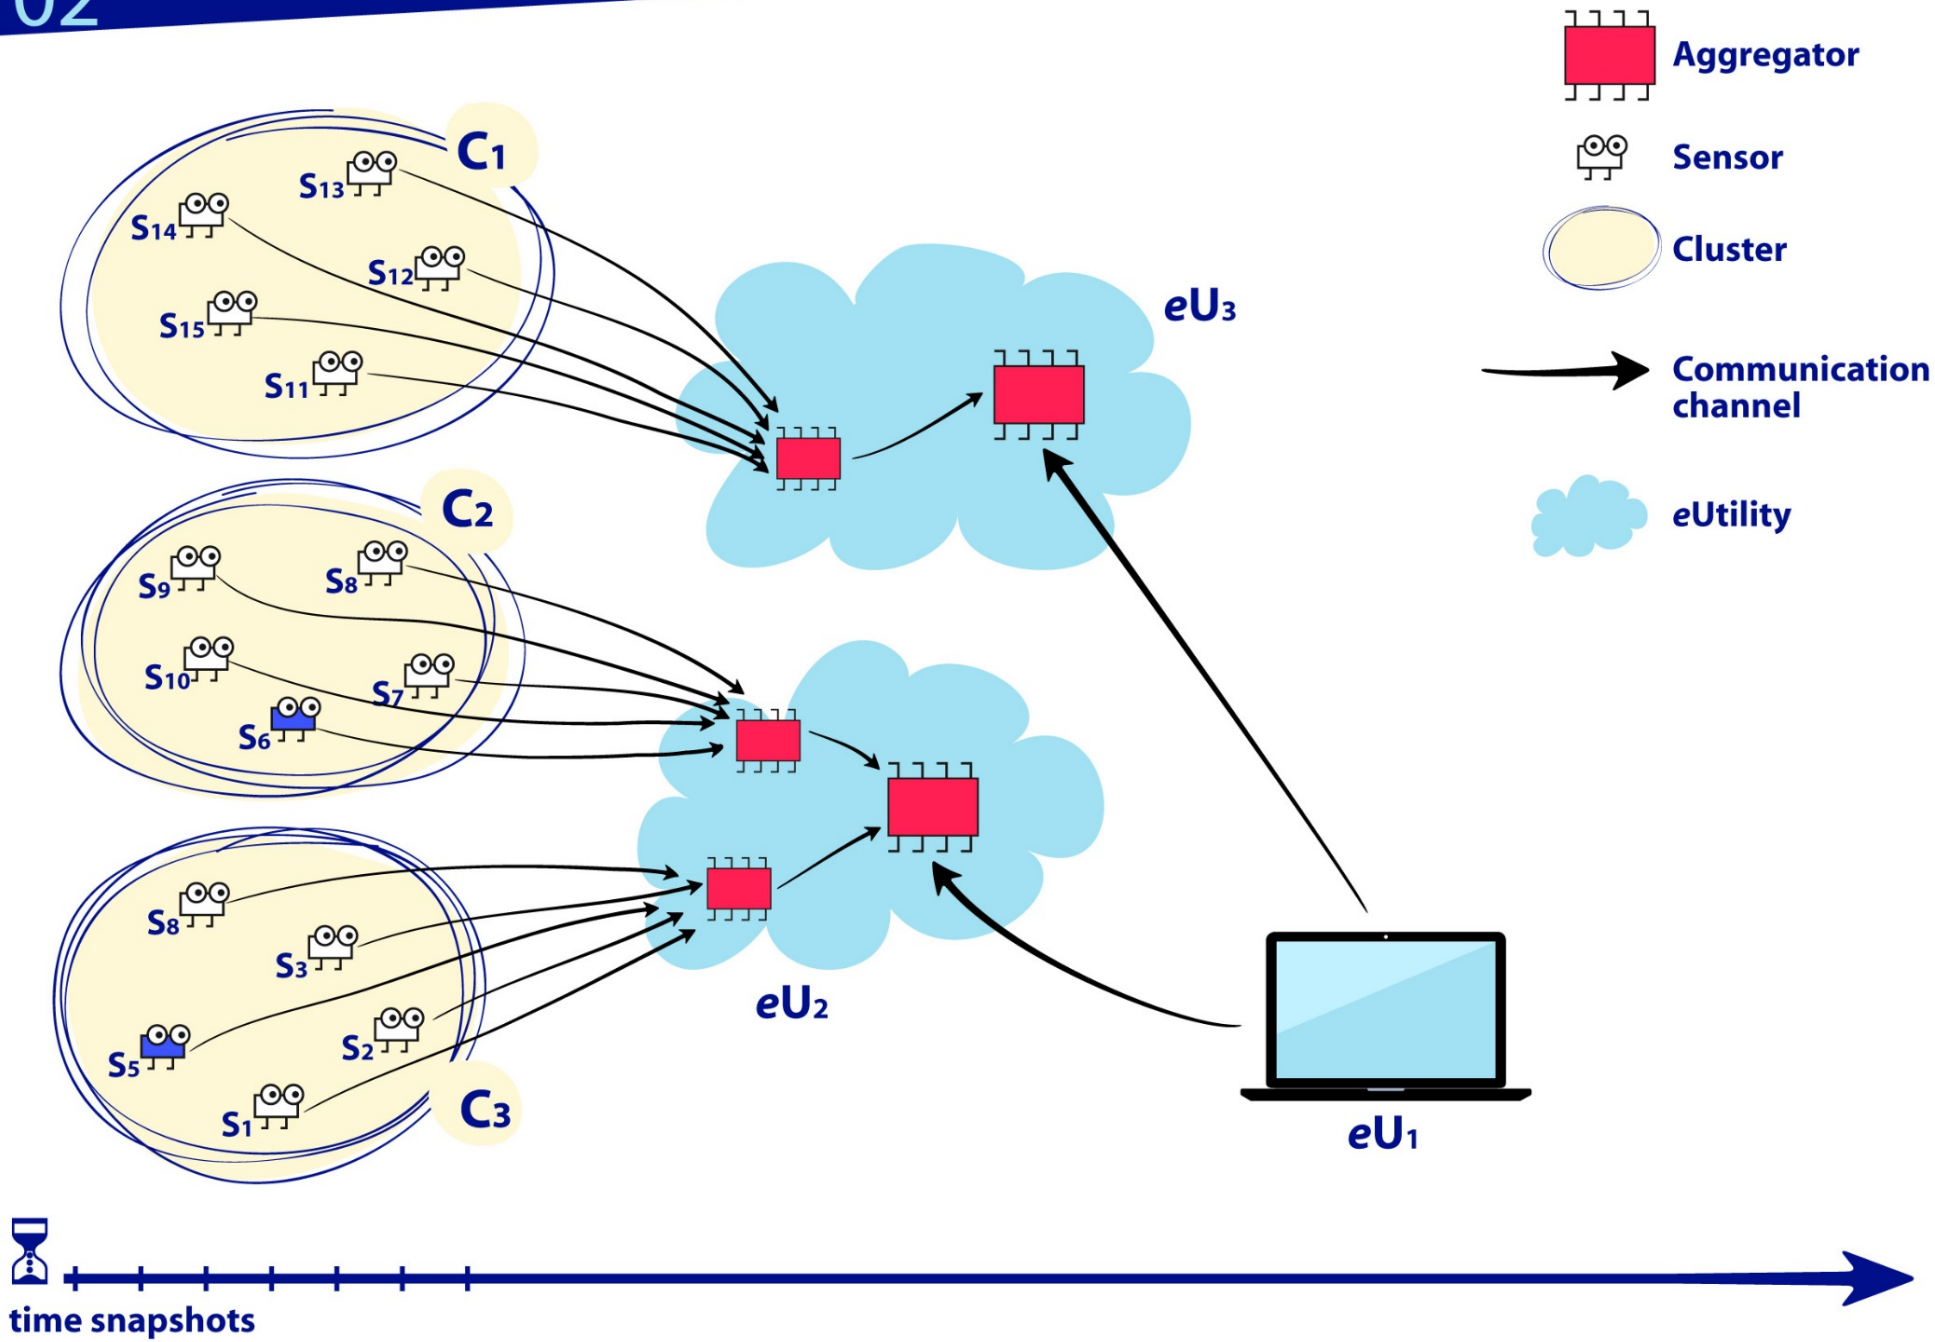
Aggregator
Sensor
Cluster
Communication channel
eUtility
C1
C2
C3
S13
S14
S12
S15
S11
S9
S8
S10
S7
S6
S8
S3
S5
S2
S1
eU3
eU2
eU1
time snapshots

Decision trigger (10 of 19)

1. A decision trigger is a conditional expression that triggers an action. A decision trigger's outputs control actuators and transactions. Decision triggers abstractly define the end purpose of a NoT.
2. A decision trigger will likely have a corresponding virtual implementation.
3. A decision trigger may have a unique owner.
4. Decision triggers may be acquired off-the-shelf or homegrown.
5. Decision triggers are executed at specific times and may execute continuously as new data becomes available.
6. It is fair to consider a decision trigger as an *if-then* rule.
7. Failure to execute decision triggers at time t_x may occur due to tardy data collection, inhibited sensors or *e*Utilities, inhibited communication channels, low performance aggregators, and a variety of other subsystem failure modes
8. Economics and costs play a role in the quality of the decision trigger's output.
9. There may be intermediate decision triggers at any point in a NoT's workflow.
10. Security is a concern for decision triggers (malware or general defects). Other possibilities here might be indirect manipulation of input values to the trigger by tampering with or restricting the input values.

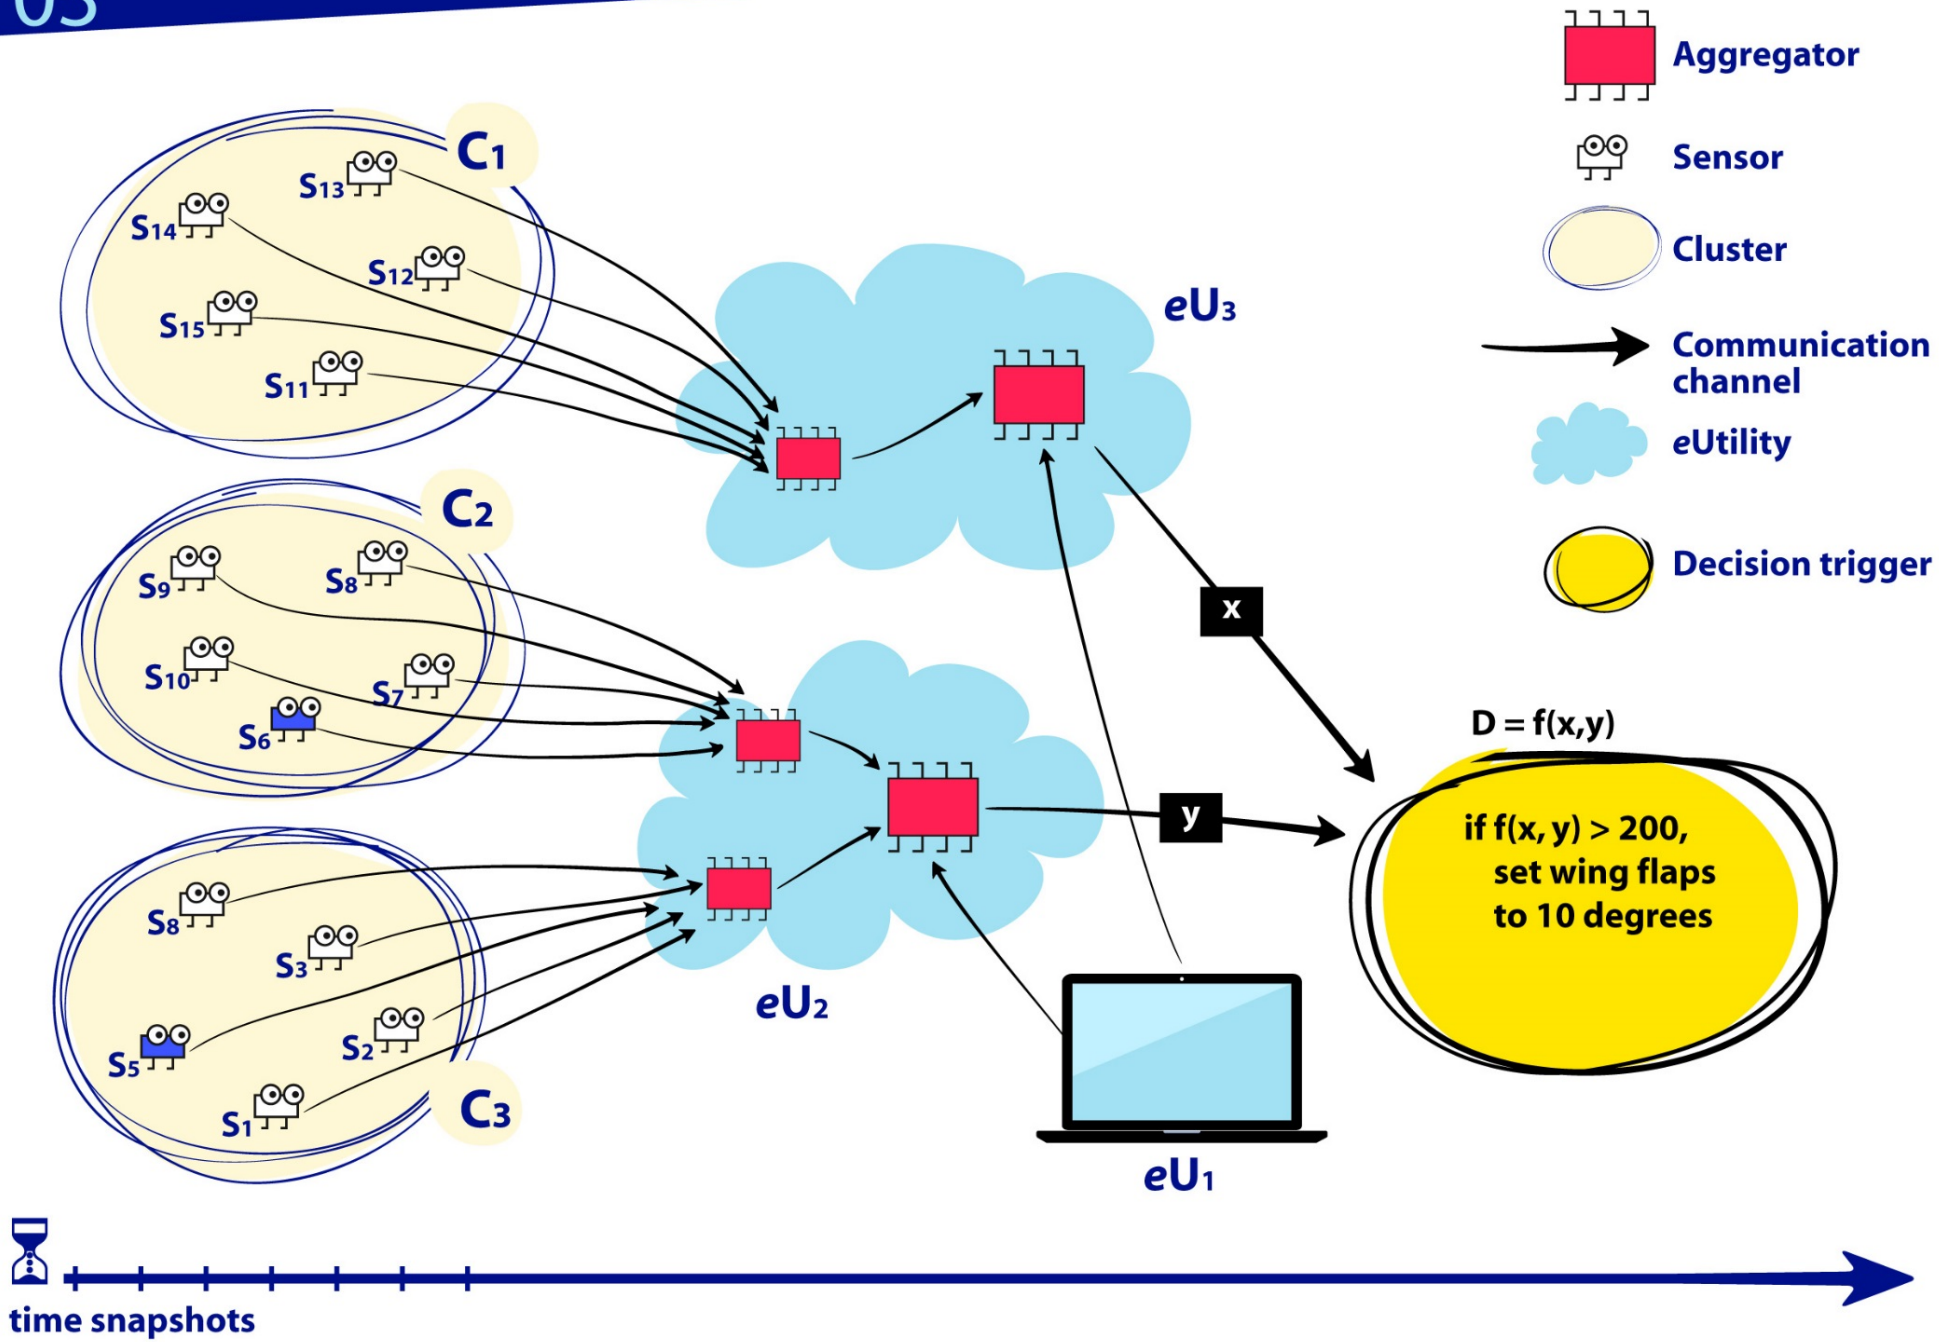
Aggregator
Sensor
Cluster
Communication channel
eUtility
Decision trigger
C1
S13
S14
S12
S15
S11
C2
S9
S8
S10
S7
S6
C3
S8
S3
S5
S2
S1
eU3
eU2
eU1
x
y
D = f(x,y)
if f(x, y) > 200,
set wing flaps
to 10 degrees
time snapshots

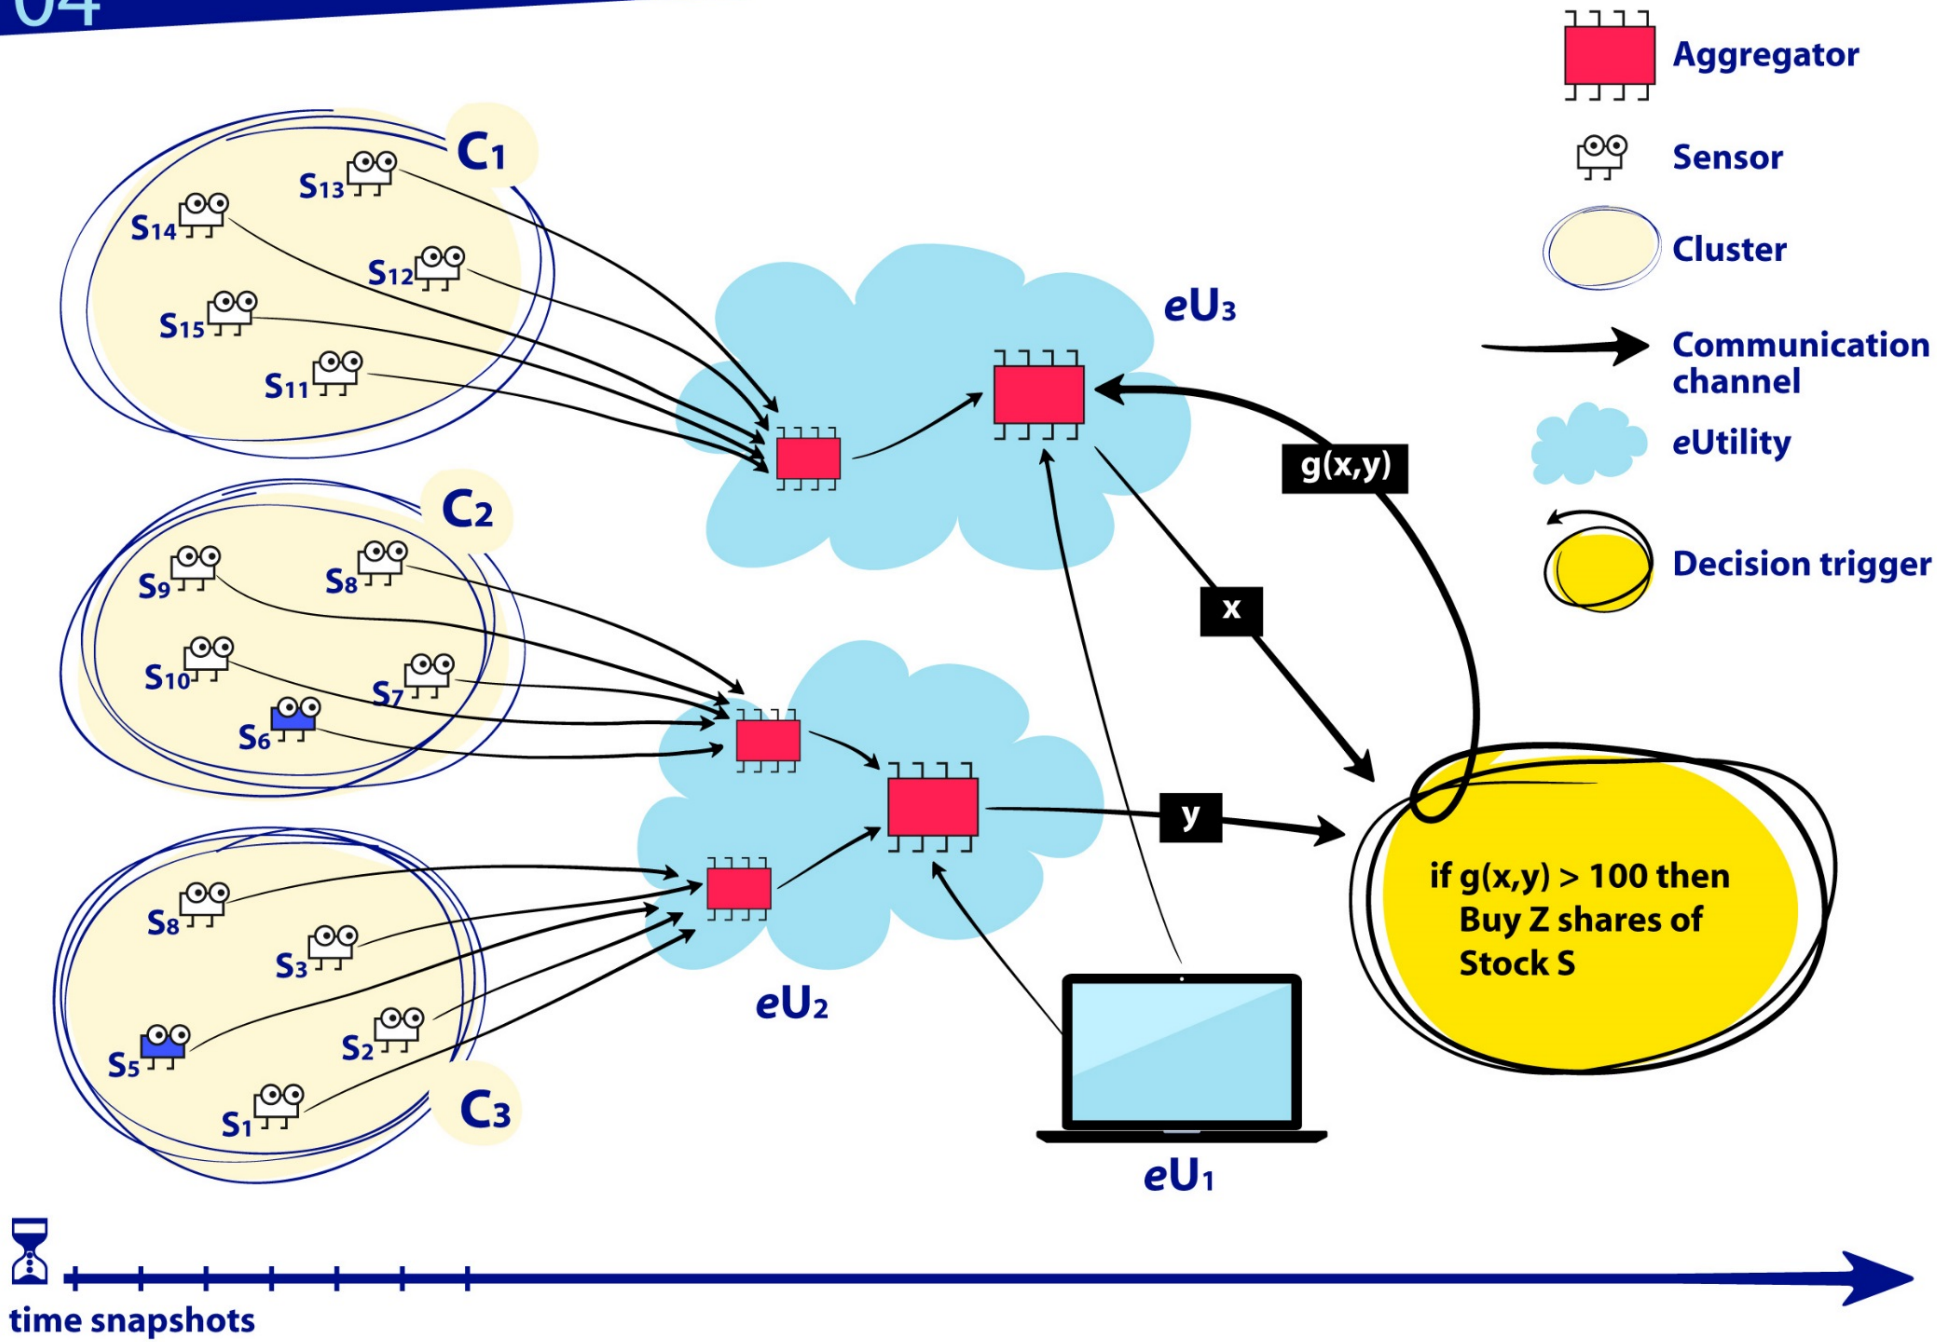
Aggregator
Sensor
Cluster
Communication channel
eUtility
Decision trigger
C1
C2
C3
S13
S14
S12
S15
S11
S9
S8
S10
S7
S6
S8
S3
S5
S2
S1
eU3
eU2
eU1
g(x,y)
x
y
if g(x,y) > 100 then
Buy Z shares of
Stock S
time snapshots

Security Examples Per Primitive

- ***Sensor:*** A smart building's temperature sensors are easily accessible and this particular system doesn't provide a means for validating the firmware's authenticity. An attacker substitutes the firmware with one that responds to remote commands. These sensors then become part of a botnet and can contribute to distributed denial-of-service (DDoS) attacks. This is an example of physical tampering and altering firmware.
- ***Aggregator:*** An attacker introduces a rogue sensor into a network that produces fake readings. These readings are passed as inputs to the aggregator function without any validation. The attacker launches a buffer overflow attack to gain root access to the entire middleware infrastructure (gateway). This is an example of an injection attack or buffer overflow.
- ***Communication channel:*** A wearable activity tracker is attached to a person's wrist and measures heart rate and blood pressure. It communicates via Bluetooth Low Energy (BLE) with the wearer's smartphone and forwards the data to a physician. Despite the fact that BLE takes specific actions to randomize the MAC address of the devices, the manufacturer neglected this feature. An attacker with a high-gain antenna can track the presence of the wearer in a crowd and create a movement profile. This is an example of eavesdropping on the communication channel.
- ***eUtility:*** A 'smart home' has a security camera installed at the front door that sends data to a corresponding cloud application that then forwards notifications and video footage to the homeowner's device after motion at the door is detected. An attacker conducts a DDoS attacks on the application provider's servers for two hours. They're able to break into the house without the user being notified. This is an example of a DDoS attack.
- ***Decision trigger:*** The decision trigger implementation accepts malicious inputs or potentially the outputs from the trigger are sniffed and released to competitors unbeknownst to the legitimate owner of the trigger. Either way, this is an example of data tampering and a loss of data integrity.

NoT Testability

- *Software testability* has a wealth of varying definitions. Today's definitions deal more with the ability to "test-out" the faults/defects.
- *Aside on testability*: the Domain Range Ratio (DRR) [Voas and Miller 1993] metric maps the cardinality of the input space to that of the output space for any "computed function": The greater the ratio, the easier it is to hide faults/failures during test.
- Consider that with respect to the previous two figures.

The "Many-to-two" Argument

- Assumption: NoT A contains faulty 'things' resulting in failures of the 'things'
  - Testing Scenario #1: 100 unique system tests and 100 unique outputs--- one-to-one mapping
  - Testing Scenario #2: 100 unique system tests and 2 unique outputs ['0', '1']--- many-to-two mapping, 50% distributed between '0' and '1'

- Because in Scenario #2 there is so little diversity in the output space, there is a likelihood that for any system test, A will produce a correct output during test, even though A is faulty (expect many faults/failures in a NoT). Therefore the "fault revealing" ability for system-wide tests is reduced for high DRRs (e.g., 100:2 versus 100:100), that is, A can almost "guess its way" to the correct output.

- And if the output distribution is not 50% distributed between '0' and '1', for example, 99% of the outputs are '1' and 1% are '0', it gets even worse for fault detection.

What Does This Suggest?

- A NoT, that for example, fires up an *actuator*, results in a *highly un-testable* NoT, due to so many influences (e.g., sensor data, *e*Utility inputs, aggregator computations), with simply a binary output from the decision trigger to the actuator.
- Therefore this hints that the use of internal probes (assertions) is the sole means to modify the ratio, and in doing so, boosts testability of a NoT.

Elements

1. **Environment** – The universe that all primitives in a specific NoT operate in; this is essentially the *operational profile* of a NoT. The environment is particularly important to the sensor and aggregator primitives since it offers context to them. An analogy is the various weather profiles that an aircraft operates in or a particular factory setting that a NoT operates in. This will likely be difficult to correctly define.
2. **Cost** – The expenses, in terms of time and money, that a specific NoT incurs in terms of the non-mitigated reliability and security risks; additionally, the costs associated with each of the primitive components needed to build and operate a NoT. Cost is an estimation or prediction that can be measured or approximated. Cost drives the design decisions in building a NoT.
3. **Geographic location –** Physical place where a sensor or *e*Utility operates in, e.g., using RFID to decide where a 'thing' actually resides. Note that the operating location may change over time. Note that a sensor's or *e*Utility's geographic location along with communication channel reliability and data security may affect the dataflow throughout a NoT's workflow in a timely manner. Geographic location determinations may sometimes not be possible. If not possible, the data should be suspect.

4. **Owner** - Person or Organization that owns a particular sensor, communication channel, aggregator, decision trigger, or *e*Utility. There can be multiple owners for any of these five. Note that owners may have nefarious intentions that affect overall trust. Note further that owners may remain anonymous. Note that there is also a role for an **operator**; for simplicity, we roll up that role into the owner element.
5. **Device_ID** – A unique identifier for a particular sensor, communication channel, aggregator, decision trigger, or *e*Utility. Further, a Device_ID may be the only sensor data transmitted. This will typically originate from the manufacturer of the entity, but it could be modified or forged. This can be accomplished using RFID for physical primitives.
6. **Snapshot** – an instant in time

Special Element: Snapshot

1. Because a NoT is a distributed system, different events, data transfers, and computations occur at different snapshots.
2. Snapshots may be aligned to a clock synchronized within their own network [NIST 2015]. A global clock may be too burdensome for sensor networks that operate in the wild. Others, however, argue in favor of a global clock [Li 2004]. This publication does not endorse either scheme at the time of this writing.
3. Data, without some "agreed upon" time stamping mechanism, is of limited or reduced value
4. NoTs may affect business performance – sensing, communicating, and computing can speed-up or slow-down a NoT's workflow and therefore affect the "perceived" performance of the environment it operates in or controls.
5. Snapshots maybe tampered with, making it unclear when events actually occurred, not by changing time (which is not possible), but by changing the recorded time at which an event in the workflow is generated, or computation is performed, e.g., sticking in a **delay()** function call.
6. Malicious latency to induce delays, are possible and will affect when decision triggers are able to execute.
7. Reliability and performance of a NoT may be highly based on (4) and (5).

Trustworthiness

Primitive or Element	Attribute	Pedigree Risk?	Reliability Risk?	Security Risk?
Sensor	Physical	Y	Y	Y
Aggregator	Virtual	Y	Y	Y
Communication channel	Virtual and/or Physical	Y	Y	Y
eUtility	Virtual or Physical	Y	Y	Y
Decision trigger	Virtual	Y	Y	Y
Geographic location	Physical (possibly unknown)	N/A	Y	Y
Owner	Physical (possibly unknown)	?	N/A	?
Environment	Virtual or Physical (possibly unknown)	N/A	Y	Y
Cost	Partially known	N/A	?	?
Device_ID	Virtual	Y	Y	Y
Snapshot	Natural phenomenon	N/A	Y	?

Summary

1. IoT is basically a standalone "brand" and catalogue of supporting technologies – IoT is not a singular technology
2. Discussing NoTs makes more "scientific" sense than discussing IoT. Can I compare your IoT to my IoT? No! Why? IoT is not measurable. NoTs can be defined, measured, and compared.
3. Primitives and elements offer "science."
4. The goal is to someday measure the Trust of a NoT:

   ***Trust* in some NoT *A*, at some snapshot *X*, is a function of NoT *A*'s assets ϵ {sensor(s), aggregator(s), communication channel(s), *e*Utility(s), decision trigger(s)} with respect to the members ϵ {geographic location, owner, environment, snapshot, cost, Device_IDs}, for each asset in the first set, when applicable.**

5. Public feedback agreed that there is *elegance* in this *simpl*e 5 + 6 part model to better answer: "What is IoT?"

Additional Takeaway Messages

1. 'Things' may be all software, hardware, a combination of both, and human.
2. A NoT may or may not employ 'things' connected to the Internet.
3. The number of 'things' in a NoT fuels functional complexity and diminishes *testability* unless *observability* is boosted by internal test instrumentation (assertions).
4. NoTs may bound scalability and complexity, and if true, may enhance an argument for trustworthiness since assurance techniques generally offer better efficacy to less complex systems.
5. Known threats from previous genres of complex software-centric systems apply to NoTs.
6. Security flaws and threats in NoTs may be exacerbated by the composition of 3rd party 'things.' This creates an *emergent* class of security 'unknowns.'
7. NoTs may have the ability to self-organize, self-modify, and self-repair when artificial intelligence (AI) technologies are introduced, e.g., neural networks, genetic algorithms, and machine learning. If true, NoTs could potentially rewire their security policy mechanisms and implementations, or disengage them altogether.
8. "After the fact" forensics for millions of composed, heterogeneous 'things', is almost certainly not possible in linear time.

9. ‘Things’ will be heterogeneous. Counterfeiting of ‘things’ may lead to seemingly non-deterministic behavior making testing’s results appear chaotic. Counterfeit ‘things’ lead to illegitimate NoTs. This is a *supply chain* concern.
10. Properly *authenticating* sensors may be a data integrity risk, e.g., the ‘who is who’ question. ‘Things’ may deliberately misidentify themselves.
11. ‘Things’ may be granted a nefarious and stealth connection capability, that is, coming and going in instantaneous time snapshots, leaving zero *traceability*. This is a “drop and run” mode for pushing external data into a NoT’s workflow. This may be mitigatable via authentication, cryptography, and possibly others. “Drop and run” affects trustworthiness.
12. *Actuators* are ‘things.’ If they are fed malicious data from other ‘things’, issues with life-threatening consequences are possible if the actuator operates in a safety-critical environment.
13. NoTs have workflows and dataflows that are highly *time*-sensitive – NoTs need communication and computation *synchronization*. Defective local/global clocks (timing failures) lead to deadlock, race conditions, and other classes of system-wide, NoT failures.

Challenges

1. 'Defining' IoT would be the hope for eventual standardization and measurement of IoT. But the mere fact that you cannot compare one IoT to another IoT demonstrates the non-actionable value of the term, IoT. IoT is nothing more than a collection or catalogue of IT technologies, components, and services.
2. Deciding what is a 'thing' in this new computing paradigm.
3. Deciding how to handle scalability issues.
4. Deciding how to handle heterogeneity issues.
5. Deciding how to handle pedigree and providence issues of 'things'.
6. Deciding how to address the economic trade-offs between the "ilities", such as security and reliability, when considered in the context of a network of 'things.'
7. Deciding how to address the technical trade-offs between the "ilities", such as security and privacy, when considered in the context of a network of 'things.'
8. Deciding on how a synchronization clock will work to make 'things' interoperate at the proper times.
9. Deciding about how 'things' will be certified, by who, and whether 'things' are even certifiable from a technical and economic standpoint.
10. Deciding on which past efforts and technologies in systems engineering, computer science, software engineering, and safety-critical systems (due to the role of actuation) are applicable will be very useful making the concept of networks of 'things' successful.
11. Deciding on whether 'things' induce any new security threats that were not anticipated.
12. Deciding on how networks of 'things' affect privacy, e.g., HIPPA?

Promising Approaches?

1. To address challenge 1, you either have to move away from the acronym IoT to the notion of a specifically-purposed network of 'things' or to a cyber-physical system. Comparing one IoT to another is 'jibberish' speak.
2. To address challenge 2, the definition of a 'thing' is vital – is it a human, computer, cloud, software, hardware, firmware, all of the above? This is a policy and standards decision, not a technical one.
3. To address challenge 3, scalability must be partially viewed as a function of the huge number of sensors in a network of 'things'. The number of sensors off-load 'big data', and that problem has few promising approaches yet. AI might offer a solution here.
4. To address challenge 4, heterogeneity is a composability problem. We've never been good with composability of large-scale, let alone large numbers of non-composable 'parts.' This requires basic research.
5. To address challenge 5, this is a policy issue. If you wish to bury your head in the sand and want no information concerning where your 'things' come from, then accept the consequences.
6. To address challenge 6, there is no solution. Basic research is needed.
7. To address challenge 7, there is no solution. Basic research is needed.
8. To address challenge 8, there is no solution. Basic research is needed based on distributed computing principles.
9. To address challenge 9, there is no solution. Certification is always a highly charged and political issue. Who certifies the certifier? Who creates the certification scheme? How do you handle mis-certification and re-certification? Basic research is needed.
10. To address challenge 10, there is no answer. Basic research is needed. But what we suspect is that this IoT computing paradigm is about 90% covered by previous technologies and in educational courses and scientific disciplines.
11. To address challenge 11, we are not clear whether 'things' induce new threats, however the combinatorics of heterogeneity, scalability, huge numbers of sensors, pedigree, and providence, have likely caused emergent classes of new threats. Basic research is needed here, as is new testing theory unique to networks of 'things.'
12. To address challenge 12, the answer is almost certainly. The more sensors, RFID tags, bar codes, surveillance you inject into an ecosystem, e.g., a hospital or factory, the more privacy diminishes. Basic research is needed, as well as privacy policies.

Thank You!

Jeff.voas@nist.gov

301-975-6622