


Deliverable reference: <b>D1.2</b>	Date: <b>09.11.2017</b>	Version: <b>V1.0</b>	Responsible partner: <b>SINTEF</b>
<h1>DiversIoT</h1>		Project co-funded by the Norwegian Research Council IKTPlus	
		<a href="http://its-wiki.no/wiki/DiversIoT:Home">http://its-wiki.no/wiki/DiversIoT:Home</a>	
Title:			
<h2>D1.2 Survey of IoT Diversity Enablers</h2>			
Editor(s): <b>Brice Morin</b>		Approved by: <b>Project Technical Manager: Franck Fleury</b>	
		Classification: <b>Unrestricted</b>	
Abstract / Executive summary:			
<p>Diversity by itself is not a challenge for the IoT. The IoT is indeed a very diversified ecosystem by nature.</p> <p>This survey gives an overview of diversity enablers that can be leveraged to build diversified IoT systems.</p>			
Document URL:		ISBN:	
		 <small>With funding from The Research Council of Norway</small>	

# Content

- Version History ..... 4**
- 1 Introduction ..... 5**
- 2 The very diverse nature of the IoT..... 7**
- 3 Implementation and Behaviour Diversity ..... 8**
  - 3.1 PROGRAMMING LANGUAGES ..... 8
  - 3.2 SEMANTICALLY-EQUIVALENT IMPLEMENTATIONS..... 10
  - 3.3 INFRASTRUCTURE & DISTRIBUTION..... 12
- 4 Communication Diversity ..... 13**
- 5 Managing Diversity in Space and in Time..... 14**
- 6 Conclusion: Challenges & Roadmap for further research ..... 15**

## Table of Figures

**NO TABLE OF FIGURES ENTRIES FOUND.**

### Version History

Version	Description	Date	Who
0.1	First complete version		Brice Morin

# 1 Introduction

The digitalized society, powered by the Internet of Things (IoT), will have enormous potential in terms of economic growth and improved quality of life. However, despite tremendous opportunities<sup>1</sup>, many companies are currently risk assessing their IoT adoption, because evidence show that it can have pernicious effect on security, privacy and safety for both individuals and the society as a whole<sup>2</sup>. Even more so, since organizations face an increasing pressure to comply<sup>3</sup> with a growing number of regulatory requirements related to management, control and monitoring of privileged access to sensitive data 11, which make organization liable as they have to “*protect it [personal data] from misuse and must respect certain rights of the data owners which are guaranteed by EU law*”. Gartner reports that more stringent regulations will lead to a rise of 40% in penalties imposed by regulatory bodies because of privacy and security issues<sup>4</sup>.

In nature, diversity has been a key mechanism for resilience for all forms of life, by enabling them to adapt to changes in the environmental conditions and evolve immunity to perturbations. More precisely in ecology, the *diversity–stability hypothesis* states that higher diversity within biological communities tends to increase resilience and reduces the risks of an ecological collapse<sup>5</sup>. Automatic software diversity is a promising solution, which is getting an increasing attention from the community<sup>6</sup>. However, because it goes against well-established software best practices, it is often impractical and current solutions are limited to low level mechanisms<sup>7</sup>. The scientific merit of this project is to challenge these well-established principles, introducing diversification allowing the development of more robust and trustworthy IoT systems. Recent advances in model compilation techniques offer new opportunities to apply diversification at a higher level of abstraction.

Recent multidisciplinary research on software engineering and ecology suggests that a promising way to approach resilience in software systems is to diversify software<sup>8,9</sup>, implying that each instance of a service has a different implementation and operates differently, still ensuring that its global behaviour is consistent and predictable. Thus, instead of exposing the very same code in millions of instances, each individual instance will be unique, making the overall system more resilient. One example nowadays used by most operating system is the randomization/obfuscation

---

<sup>1</sup> Gartner Says 6.4 Billion Connected "Things" Will Be in Use in 2016, Up 30 Percent From 2015, see <http://www.gartner.com/newsroom/id/3165317>

<sup>2</sup> Gartner Says Worldwide IoT Security Spending to Reach \$348 Million in 2016, see <http://www.gartner.com/newsroom/id/3291817>

<sup>3</sup> Protection of Personal Data, see <http://ec.europa.eu/justice/data-protection/>

<sup>4</sup> Gartner Research, Market Guide for Privileged Access Management, 2015

<sup>5</sup> Anthony R. Ives et al., «Stability and diversity of ecosystems,» *Science* 317, 58 – 62., 2007

<sup>6</sup> P. Larsen et al., «SoK: Automated Software Diversity,» i IEEE Symp. Security and Privacy, 2014.

<sup>7</sup> M. Stamp, «Risks of Monoculture,» *Communication of the ACM*, vol. 47, nr. 3, 2004.

<sup>8</sup> Laperdrix et al., «Beauty and the Beast: Diverting modern web browsers to build unique browser fingerprints,» in IEEE Symposium on Security and Privacy (S&P'16), 2016.

<sup>9</sup> F. Fleurey et al., «Multi-tier diversification in Web-based software applications,» *IEEE Software*, 32 (1), pp. 83-90, 2015.

of memory addresses<sup>10</sup>, to prevent a number of exploits, making it difficult for a malware program to jump at an offset in the memory and access critical data.

---

<sup>10</sup> Bhatkar, Sandeep, Daniel C. DuVarney, and Ron Sekar. "Address Obfuscation: An Efficient Approach to Combat a Broad Range of Memory Error Exploits." *USENIX Security Symposium*. Vol. 12. No. 2. 2003.

## 2 The very diverse nature of the IoT

IoT systems are composed by a highly heterogeneous interconnection of platforms and devices offering a wide diversity of capabilities<sup>11</sup>. On the one end of the continuum, cloud platforms provide virtually unlimited and "on-demand" resources in terms of computation power, storage and bandwidth. On the other end, the already vast and rapidly increasing number of smart objects, sensors, embedded systems and mobile devices connected to the Internet offers the connection to the users and to the physical world. These IoT devices:

- run on diverse platforms (ARM, MIPS, AVR, PIC)
- can be rather resource-constrained (as low as 2 KB RAM and 8 MHz CPU)
- or can be rather powerful (1GB RAM, 1GHz CPU)
- can run an OS (fully-fledged Linux for the most powerful, embedded Linux for the
- or can be too limited to even run an OS
- can communicate through different protocols: WiFi, Bluetooth, Bluetooth Low Energy, ZigBee, LoRa, Z-Wave, etc

While offering an immense potential for innovative IoT systems, the heterogeneity, diversity and vast distribution of the computing continuum represent daunting challenges<sup>12</sup>. Current software engineering approaches tend to provide support for managing and exploiting only parts of the continuum. For example, current cloud computing and service-oriented software engineering practices provide efficient abstractions for virtualizing the infrastructure in order for the software engineer to concentrate on the business logic of the applications. However these techniques merely support integration of mobile devices, sensors and actuators "as-a-service" and lack specific support for advanced exploitation of these small devices.

Furthermore, although sensor network nodes, gateways, smart-phones, and most smart-objects are provided with APIs and development environment which allow advanced exploitation of their capabilities. There are no common practices to expose the level of flexibility, languages, APIs and customization supported by different devices and platforms at the level of what is provided in cloud and service oriented approaches. Thus, most pervasive systems tend to be proprietary silos not easily exploitable by other systems as they offer only black box provisioning of higher level services.

---

<sup>11</sup> Morin, Brice, Nicolas Harrant, and Franck Fleurey. "Model-Based Software Engineering to Tame the IoT Jungle." *IEEE Software* 34.1 (2017): 30-36.

<sup>12</sup> Morin, Brice, Franck Fleurey, and Olivier Barais. "Taming heterogeneity and distribution in scps." *Proceedings of the First International Workshop on Software Engineering for Smart Cyber-Physical Systems*. IEEE Press, 2015.

## 3 Implementation and Behaviour Diversity

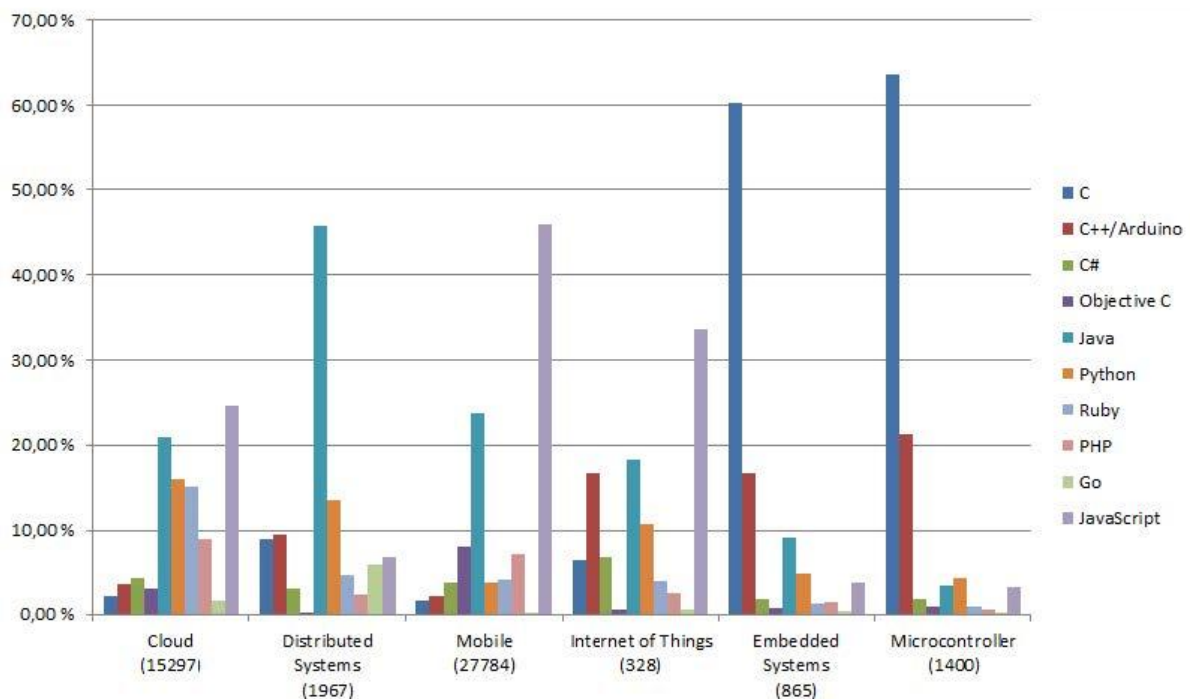
### 3.1 Programming Languages

A typical IoT infrastructure spans across a continuum of devices and platforms ranging from microcontroller-based devices up to Clouds. Software for the different classes of devices are typically built using different approaches and languages. In order to understand the skills and capabilities required to develop services on top of such an infrastructure, we queried a popular open-source repository (GitHub) to evaluate the heterogeneity of programming languages across the continuum. The following sets of keywords were used:

1. Cloud: server with virtually unlimited resources,
2. Microcontroller: resource constrained node (few KB RAM, few MHz)
3. Mobile: an intermediate node, typically a smartphone,
4. Internet of Things: Internet-enabled devices,
5. Distributed systems, as services exploiting CPS have to be distributed across the continuum,
6. Embedded systems, as a large and important part of the service implementations will run as close as possible to physical world, embedded into sensors, devices and gateways.

Figure 1 presents the results of those queries. The queried keywords are presented on the x axis together with the number of matches for that keyword. For each keyword, the y axis represents the popularity (in percent of the total number of matches) of each of the 10 most popular programming languages that we encountered. This simple study indicates that no programming language is popular across the whole continuum. A general trend indicates that Java and JavaScript (and to some extent, Python and Ruby) are popular in the higher-end of the continuum (cloud and mobile), whereas C (and to some extent, C++) is a clear choice for developers targeting embedded and microcontroller-based systems. Other languages do not score more 10% for any of the keywords. For all keywords except Cloud, the combined popularity of Java, JavaScript and C/C++ (i.e., the sum of the percentages) is above 70%. For Cloud, we observe a certain homogeneity with Python, Ruby also being very popular, so the combined popularity of Java, JS and C/C++ is only 50%. It is also worth noticing that the most popular language for a given keyword scores very poorly (less than 5%) for at least another keyword. While it might appear that a combination of C/C++,





**Figure 1: Programming languages used within different domains. The bars show the percentage of the primary programming language that used in GitHub repositories for each of the tags along the horizontal axis.**

JavaScript and Java should be able to cover the whole continuum of IoT, in practice it does not exclude the need for other programming languages. For example, the Fibaro Home Center 2 (a gateway for home automation based on the Z-Wave protocol) uses Lua as scripting language to define automation rules<sup>13</sup>. Another example is the BlueGiga Bluetooth Smart Module, which can be scripted using BGScript, a proprietary scripting language. This shows that each part of an infrastructure might require the use of a niche language, middleware or library to be exploited to its full potential. Other examples of such languages include micro-python<sup>14</sup>, micro-ruby<sup>15</sup>, Rust<sup>16</sup> or even resource-constrained JavaScript interpreters such as Espruino<sup>17</sup>.

A wide diversity of programming languages can be used and combined to implement IoT systems. Different languages will, for semantically equivalent programs, have a different impact on resources (CPU, memory, I/Os, etc), yielding diversified footprints.

<sup>13</sup> A brief introduction to Lua, Fibaro, see <http://www.fibarouk.co.uk/support/lua/brief-introduction-lua/>

<sup>14</sup> <https://github.com/micropython/micropython>

<sup>15</sup> <http://microruby.com/>

<sup>16</sup> <https://www.rust-lang.org/en-US/>

<sup>17</sup> <https://github.com/espruino/Espruino>

### 3.2 Semantically-equivalent implementations

Say that for a given language we can implement the same behaviour in many different ways. Give some examples of concrete strategies that could be applied.

For a given programming language, it exists infinite ways of coding the expected behaviour of a program. To take a very simple example, affecting the value 1 to a variable can be achieved by:

- `var i = 1`
- `var i = 1 + 1 - 1`
- `var i = f1() //call to a function returning 1`
- `var j = 1; var i = j`
- ...

This diversity can be achieved by producing different versions of the source code, or from a single source code, by letting the compilers to produce diversified machine code. A rather large community is providing techniques for compilers to diversify machine code automatically<sup>18, 19, 20</sup>.

Another interesting point of variability is the communication between distributed components. Strategies for distributions and communication are described in more details later in this document. In this section, we describe different communication patterns between two processes (A and B), independently of the actual way message are serialized and transported through the network. Figure 2 describes a few patterns that can be used and combined to introduce diversity when communicating, still preserving semantics.

A major choice is whether to “stream” data (A.) (i.e. sending multiple fragment of data, which we denote as a, b, and c) or to send data as a batch (B.) when component A and B need to exchange data. Those two patterns have different implications:

- streaming might be heavier on the network traffic, providing more opportunities to intercept messages. However, each message only carries a subset of the information.
- batch can reduce overall network traffic, reducing the probability that a message is intercepted. However, if intercepted, the message carries the whole information.

The timing (C.) and ordering (D.) of the messages that are sent can also be adjusted, so as to avoid exposing any patterns, which could be more easily interpreted.

---

<sup>18</sup> JAJODIA, Sushil, et al. (ed.). *Moving Target Defense II: Application of Game Theory and Adversarial Modeling*. Springer Science & Business Media, 2012.

<sup>19</sup> DAVI, Lucas, et al. Isomeron: Code Randomization Resilient to (Just-In-Time) Return-Oriented Programming. In: *NDSS*. 2015.

<sup>20</sup> Homescu, Andrei, et al. "Librando: transparent code randomization for just-in-time compilers." *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013.

The routing (E.) of the messages between A and B can also be adjusted, for example by introducing intermediate jumps so as to blur the origin of the message, though additional information should be added so as B can know that a message comes from A. This is a mechanism somehow similar to the Tor network.

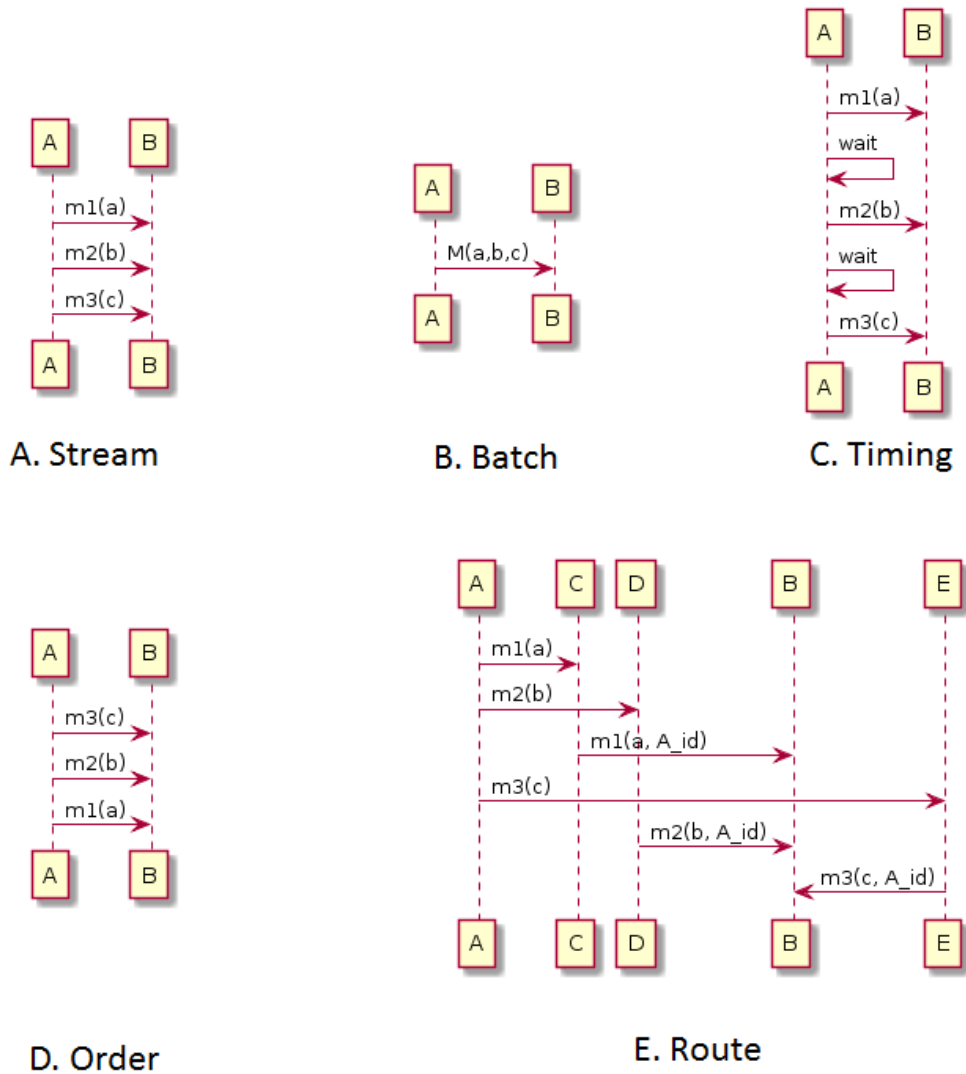
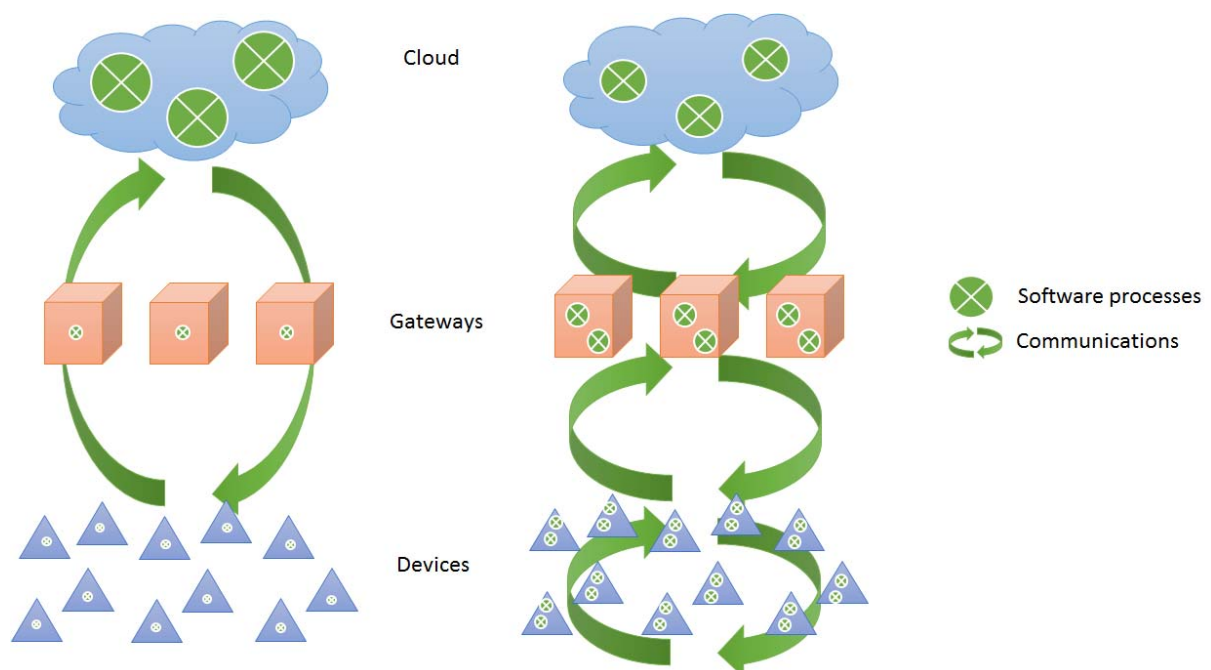


Figure 2: Communication patterns between components.

### 3.3 Infrastructure & Distribution

For a given piece of software, provided it is architected as a set of components able to communicate together asynchronously can be distributed in different ways on the available infrastructure<sup>21</sup>.

A typical, yet quite naive, distribution pattern for IoT systems is depicted on the left of Figure 3, where most of the logic is delegated to the cloud back-end. In this setup, devices send raw data to gateways, which simply forward to the cloud. Processes running in the cloud analyse those data and eventually device to enact some commands, which are sent to the gateways and forwarded to the proper devices.



**Figure 3: Distribution patterns for IoT software. Traditional centralized pattern on the left, and a more distributed approach on the right.**

A more robust way of distributing IoT systems shown on the right of the figure, as advocated by EU FP7 HEADS project<sup>22</sup>, is to distribute the logic in a more even way across the IoT infrastructure<sup>23</sup>. This way, decisions can be taken and enacted more locally and the overall system can be made more robust, as it can indeed work (potentially in a degraded mode) without a permanent connection to the cloud. This is for example the way the new Ikea Trådfri home-automation equipment work.

<sup>21</sup> VINOSKI, Steve. CORBA: integrating diverse applications within distributed heterogeneous environments. *IEEE Communications magazine*, 1997, 35.2: 46-55

<sup>22</sup> <http://heads-project.eu/>

<sup>23</sup> MORIN, Brice, et al. A generative middleware for heterogeneous and distributed services. In: *Component-Based Software Engineering (CBSE), 2016 19th International ACM SIGSOFT Symposium on*. IEEE, 2016. p. 107-116.

## 4 Communication Diversity

As we have described previously, a given IoT system can be deployed in many different ways and rely on different communication patterns. Communication itself can be implemented in very diverse ways, as we describe in this section.

Serialization is an important step in order to be able to send messages on the network for two components to communicate together. A large set of libraries and languages provide facilities to ease the serialization and deserialization of messages into strings or binary formats: Google's FlatBuffer<sup>24</sup>, Google's Protocol Buffer<sup>25</sup>, Apache Avro<sup>26</sup>, Apache (formerly Facebook's) Thrift<sup>27</sup>, Cap'n Proto<sup>28</sup>, SBE (Simple Binary Encoding)<sup>29</sup>, MessagePack<sup>30</sup>, eProxima Fast Buffers<sup>31</sup>, JSON<sup>32</sup>, XML<sup>33</sup>, etc.

In addition to these "standard" and very popular approaches most programming languages provide ways to manipulate strings and arrays or buffers of bytes, allowing for custom serialization formats.

Different ways of serializing a message will typically ultimately end up with different arrangements of bits on the network.

To actually transport a payload, different protocols can be used, such as HTTP/REST<sup>34</sup>, CoAP<sup>35</sup>, MQTT<sup>36</sup>, AMQP<sup>37</sup>, Bluetooth<sup>38</sup>, Bluetooth Low Energy<sup>39</sup>, Serial<sup>40</sup>, ZigBee<sup>41</sup>, LoRa<sup>42</sup>, SigFox<sup>43</sup>, LTE-M<sup>44</sup>, etc. Different protocols provide different tradeoff in terms of performances, energy consumption, etc.

---

<sup>24</sup> <https://google.github.io/flatbuffers/>

<sup>25</sup> <https://developers.google.com/protocol-buffers/>

<sup>26</sup> <https://avro.apache.org/>

<sup>27</sup> <http://thrift.apache.org/>

<sup>28</sup> <https://capnproto.org/>

<sup>29</sup> <https://github.com/real-logic/simple-binary-encoding>

<sup>30</sup> <http://msgpack.org/>

<sup>31</sup> <http://www.eprosima.com/index.php/products-all/eprosima-fast-buffers>

<sup>32</sup> <http://www.json.org/>

<sup>33</sup> <https://www.w3.org/XML/>

<sup>34</sup> <https://spring.io/understanding/REST>

<sup>35</sup> <http://coap.technology/>

<sup>36</sup> <http://mqtt.org/>

<sup>37</sup> <https://www.amqp.org/>

<sup>38</sup> <https://www.bluetooth.com/>

<sup>39</sup> <https://www.bluetooth.com/what-is-bluetooth-technology/how-it-works/le-p2p>

<sup>40</sup> <http://www.ni.com/white-paper/11390/en/>

<sup>41</sup> <http://www.zigbee.org/>

<sup>42</sup> <https://www.lora-alliance.org/>

<sup>43</sup> <https://www.sigfox.com/en>

<sup>44</sup> <https://www.link-labs.com/blog/what-is-lte-m>

## 5 Managing Diversity in Space and in Time

Diversity can not only be managed in space, i.e. by having diversified instances, but also in time, by having a given instance to change over time. In order to support diversity in time, the running system should be able to undergo changes at runtime, after it has initially be deployed. The challenge is to minimize downtime. Over the past 20 years, the software engineering community has actively been working on dynamically adaptive systems<sup>45</sup>, which are able to be reconfigured, or to reconfigure themselves<sup>46</sup>, at runtime. In the state of the practice, approaches like Docker containers<sup>47, 48</sup> (built on top of Linux containers and C-groups) now provide convenient way to package, deploy and update micro-services systems in a very agile and dynamic way.

---

<sup>45</sup> MORIN, Brice, et al. Taming dynamically adaptive systems using models and aspects. In: *Software Engineering, 2009. ICSE 2009. IEEE 31st International Conference on*. IEEE, 2009. p. 122-132.

<sup>46</sup> KEPHART, Jeffrey O.; CHESS, David M. The vision of autonomic computing. *Computer*, 2003, 36.1: 41-50.

<sup>47</sup> <https://www.docker.com/>

<sup>48</sup> STUBBS, Joe; MOREIRA, Walter; DOOLEY, Rion. Distributed systems of microservices using docker and serfnode. In: *Science Gateways (IWSG), 2015 7th International Workshop on*. IEEE, 2015. p. 34-39.

## 6 Conclusion: Challenges & Roadmap for further research

Diversity by itself is not a challenge for the IoT. The IoT is indeed a very diversified ecosystem by essence. This survey gave a brief overview of diversity enablers that can be leveraged to build diversified IoT systems.

The key research challenges, related to diversity, to be addressed in the rest of this project are:

- determine what kind of diversity is actually useful to implement robust, resilient and trustworthy IoT systems
- manage this “useful” diversity in a seamless and automated way so as to ensure that the different parts of an IoT system (typically devices, gateways and cloud) can still interoperate and collaborate together, in time (the same instance “looks” different at different points in time) and in space (different instances “look” different), though they might be “different”
- though automated management is needed to tame this diversity, it should:
  - not limit the potential of diversity, so as not to end up with a limited and predictable set of variants, which could easily be understood by attackers
  - not expose explicitly any key or footprint

DiversIoT will address those challenges by raising the level of abstraction of IoT systems. By adopting a Model-Driven generative approach to produce diversified concrete implementations, DiversIoT will be able to manage diversity at a more abstract level and in a more automated and systematic way. A prototype also delivered in Phase 1 shows the feasibility of such an approach.