

# A novel JXTA-based architecture for implementing heterogeneous Networks of Things

Filippo Battaglia<sup>a</sup>, Lucia Lo Bello<sup>a,\*</sup>

<sup>a</sup>*DIEEI - Department of Electrical, Electronic and Computer Engineering, Viale Andrea Doria, 6, 95125 Catania*

---

## Abstract

This paper presents EmbJXTAChord, a novel peer-to-peer (P2P) architecture that integrates the good features of different sources, such as JXTA, EXI, CoAP, combining and augmenting them to provide a framework that is specifically devised for developing IoT applications over heterogeneous networks.

EmbJXTAChord provides for several interesting properties, such as, distributed and fault-tolerant resource discovery, transparent routing over subnetworks, application protocol independence from the transport protocol in narrowband Wireless Sensor Networks, thus eliminating the need for using dedicated software or configuring custom gateways to achieve these functionalities.

Moreover, EmbJXTAChord offers native support not only for TCP/HTTP, but also for Bluetooth RFCOMM and 6LoWPAN, thus opening to a broad range of IoT devices in supernetworks composed of networks using different interconnection technologies, not necessarily IP-based. In addition, EmbJXTAChord offers security over heterogeneous networks providing support for secure peergroups (even nested) and for group encryption, thus allowing for unicast and multicast communication between groups of objects sharing the same resources. The users of the proposed architecture will benefit from an integrated solution and the applications developed on the proposed framework will be able to reconfigure themselves, adapting automatically to the network topology of the execution environment.

Finally, EmbJXTAChord provides jxCOAP-E, a new CoAP implementation that leverages on the transport mechanisms for heterogeneous networks offered by EmbJXTAChord. jxCOAP-E enables to realize a RESTful service architecture for peer-to-peer narrowband or broadband networks composed of devices connected via Ethernet, Wi-Fi, Bluetooth, BLE or IEEE 802.15.4. Differently from CoAP, jxCOAP-E provides a distributed and fault-tolerant service discovery mechanism and support for secure multicast communications. The paper presents EmbJXTAChord, discusses all the relevant design challenges and presents a comparative experimental performance assessment with state-of-the-art solutions on commercial-off-the-shelf devices.

---

**Note:** *This is a preprint version of an Elsevier accepted paper. The formal version of the work is available at <https://doi.org/10.1016/j.comcom.2017.11.002>*

©2017. This manuscript version is made available under the CC-BY-NC-ND 4.0 license <http://creativecommons.org/licenses/by-nc-nd/4.0/>

---

\*Corresponding author

Email addresses: [fbattaglia@unict.it](mailto:fbattaglia@unict.it) (Filippo Battaglia), [lobello@unict.it](mailto:lobello@unict.it) (Lucia Lo Bello)

3GPP	3rd Generation Partnership Project	NAT	Network Address Translation
6LoWPAN	IPv6 over Low power Wireless	PAN	P2P Peer-to-Peer
AES	Advanced Encryption Standard	PAN	Personal Area Network
ATT	Attribute Protocol (BLE)	RaspPI	Raspberry PI
BLE	Bluetooth Low Energy	REST	Representational State Transfer
BNEP	Bluetooth Network Encapsulation Protocol	RPV	Rendezvous PeerView
CM	EmbJXTAChord Compressor Manager	RTT	Round Trip Time
CoAP	Constrained Application Protocol	SICS	6LoWPAN
CSR	Cambridge Silicon Radio	SOAP	Simple Object Access Protocol
DHT	Distributed Hash Table	SRDI	Shared Resource Distributed Index
DTLS	Datagram Transport Layer Security	TLS	Transport Layer Security
ETH	Ethernet	URI	Uniform Resource Identifier
EXI	Efficient XML Interchange	UDDI	Universal Description Discovery and Integration
GATT	Generic Attribute Profile (BLE)	UTF	Unicode Transformation Format
IETF	Internet Engineering Task Force	WAVE	Wireless Access in Vehicular Environment
IMX	Intermessage compression	WG	Working Group
LRW	Limited Range Walker	WS	Web Service
LTE	Long Term Evolution	WSN	Wireless Sensor Network
MIME	Multipurpose Internet Mail Extensions	XML	eXtensible Markup Language
MQTT	Message Queue Telemetry Transport	adv	Advertisement
MSA	Module Specification Advertisement	rdv	Rendezvous
MTB	Message Transport Binding	uPNP	Universal Plug and Play

Table 1: List of the acronyms and abbreviations used in the paper.

## 1. Introduction and motivation

A typical domestic scenario of today includes several computers, smart phones and appliances connected through wired Ethernet, power lines or IEEE 802.11 (Wi-Fi) wireless links, which can, in turn, access sensors or actuators through multiple wireless protocols, such as Wi-Fi, Bluetooth or IEEE 802.15.4. For instance, a palmtop can run an application that sets the start time for the washing machine that, in turn, can defer its work based on the constraints of the power meter and on the energy output currently provided by the solar panels.

This scenario is only a small-scale example of the Internet of Things (IoT), in which a number of interconnected cooperative smart objects provide multi-agent services [1]. However, none of the frameworks proposed in the IoT field so far is able to fully guarantee two fundamental properties, i.e., the application source code independence from the underlying network topology and the suitability for any kind of subnetwork. Currently the IoT applications must be tailored to the middleware and the network configuration, as the available solutions are mainly incompatible [2]. In fact, each middleware exploits its own API and data formats for requests and responses.

In the past, in an attempt of standardization, the use of web services for domestic or industrial applications based on SOAP <sup>1</sup> (Simple Object Access Protocol) was proposed [3][4][5]. Other works proposed REST (Representational state transfer), a paradigm where all resources are accessible by a common HTTP client (as a browser) [6]. Unfortunately, both SOAP and REST leverage on the transmission of large HTML or XML documents, which can be too bandwidth-consuming for low-bandwidth wireless sensor networks (WSN) protocols, such as the IEEE 802.15.4 [7][8][9].

As a consequence, new lightweight and efficient protocols for IoT communication between smart devices, such as CoAP (Constrained Application Protocol)

<sup>1</sup>For convenience of the reader, the acronyms most frequently used in this paper are reported in Tab. 1 and 2.




Figure 1: A typical scenario for a heterogeneous network.

[10][11], MQTT (Message Queue Telemetry Transport) [12] or SMQ (Simple Message Queries) [13][14], were recently developed [15].

Under CoAP, a new protocol recently proposed by the IETF CoRE Working Group, all the functionalities are seen as a set of resources, identified by a URI (Uniform Resource Identifier) and all the operations are performed through four methods (GET, PUT, POST and DELETE). Under MQTT and SMQ, instead, sensor readings are seen as *topics*, whose change can be notified by a *central broker* to a set of subscribers.

The protocols cited above were designed minimizing the functionalities provided to the developer, thus obtaining bandwidth-efficient and energy-efficient communications between resource-constrained devices [16][12].

The solutions based on CoAP, MQTT, SOAP-WS or REST [17][15] need further components to provide the functionalities of distributed service discovery [18] [19][20], fault tolerance [21] and resource protection [22]. Moreover, these IoT protocols use the underlying IP network layer, assuming the availability of such a layer for each link-layer protocol.

The envisaged vision of all devices interconnected via IP determined the proliferation of middleware solutions for IoT that are designed to work over IPv4/IPv6. However, in a heterogeneous network made up of devices connected through different kinds of protocols at the physical, network and transport levels, located at different places and separated by firewalls, gateways and NATs, the creation of an IPv6 network can be difficult. For example, let us consider Fig. 1. The laptop is connected to a subnetwork consisting of a printer and a mobile phone that is connected to the laptop via Bluetooth or Bluetooth Low Energy (BLE), as this protocol is less power-consuming than Wi-Fi. A firewall operates between the first and the second subnetwork and a 6LoWPAN smart gateway between the second and the third subnetwork. The first two subnetworks are based on IPv4, while the third one on IPv6/6LoWPAN. Some issues can prevent the delivery of the CoAP messages between the laptop or the mobile phone and the WSN in the third subnet, for instance

- **Bad firewall configuration.** The firewall can prevent the delivery of the packets to the CoAP standard port;
- **Different network protocols.** The three networks use different network protocols, therefore there are different addressing schemes (32b for IPv4, 128b for IPv6);
- **Difficult use of IP in Bluetooth networks for some devices.** The mobile phone firmware might not support the Ethernet emulation for Blue-

tooth (BNEP, the Bluetooth Network Encapsulation Protocol), but only connections via RFCOMM (serial port emulation). In such a case, a custom gateway RFCOMM/IP should be implemented. Some issues may occur even if the mobile phone can use some of the new BLE features, such as the IPSP (Internet Protocol Support Profile) [23] that allows to create a 6LoWPAN network over BLE (6LoBLE). For instance, the BLE nodes that work as IPSP routers cannot be connected to a IPv4 subnetwork without an IPv6/IPv4 gateway. Moreover, some operating systems might also have limitations that hinder the use of BNEP or IPSP [24].

- **Complex security implementation.** IoT communications can be secured using IPSec [25], TLS (Transport Layer Security) [26] or DTLS (Datagram Transport Layer Security) [27]. TLS and DTLS are application-level protocols that work over TCP and UDP, respectively. IPSec is a network-level protocol that may require a difficult configuration [28], which makes the protocol prone to interoperability issues, and specific customizations (for instance to work with 6LoWPAN [29]). Unfortunately, all these protocols are unsuitable for non-IP networks (such as the Bluetooth RFCOMM piconets).
- **Difficult implementation of multicast group communications.** The standard request-response interaction model used by CoAP supports only end-to-end message exchange. For this reason, the IETF CoRE WG released the RFC 7390 [30] that defines how CoAP should support *group communications*. As this feature leverages on IP multicast packet delivery, it cannot be implemented when the physical layer supports only unicast communications (such as in the BT and BLE network cases [23]) or when the bridges between subnetworks prevent the propagation of these messages. Moreover, the RFC 7390 does not state how to implement *secure* group communications among nodes [31][32]. This is an important limitation, as the TLS and DTLS implementations currently available do not support multicast communications, despite some proposals were published [33] [34].

In such a situation, the only solution would be to reconfigure all the subnetworks in order to use the same addressing schema (IPv6) and to manually configure the firewall and the routing tables of the border routers, so as to manage packet delivery between a device of the first and of the third subnet.

These issues can be overcome using a peer-to-peer (P2P) protocol, working at the application level, able to support different network protocols through dedicated modules, thus abstracting from the link technology actually adopted and providing for automatic bridging. This way, the mobile phone can communicate to the 6LoWPAN device (and vice versa) using a chain of bridges (the laptop, the firewall between the first and the second subnet, the smart gateway between the second and the third subnet) in a seamless way. This solution would have been considered unsuitable for IoT in the past, because the use of a P2P protocol for message exchanging determines a consistent overhead compared to the use of the bare IPv4/IPv6. Nowadays, P2P communication became an interesting option for IoT, thanks to the availability on the market of small-sized embedded boards, such as the Raspberry PI [35], or the Raspberry PI-3 [36] that provide

enough computational power to support P2P protocols even on low-cost COTS hardware. The P2P protocol should provide the following functionalities:

- Host resolution by name. In a heterogeneous network, a device should be addressed by name and not by the IP address. In fact, IP or the Dynamic Host Configuration Protocol (DHCP) might be unavailable in the device subnetwork or the DHCP server could dynamically change the assigned IP address <sup>2</sup>;
- Network protocol translation. As using IP over each link is not necessary, any protocol can be adopted. In the described scenario, the laptop transfers the IoT messages from the IP section of the network to the Bluetooth RFCOMM section transparently to the application level;
- HTTP tunnelling. Message delivery is possible through the firewall without manual reconfiguration;
- Routing over subnetworks. The P2P protocol must be able to deliver the messages between any pair of devices, even when they belong to different subnetworks (for instance, a BT mobile phone and an IEEE 802.15.4 sensor node). The P2P protocol should be able to automatically determine the information needed for routing, and routing should be performed by the application level without reconfiguring the IP forwarding tables in the router or in the operating system.
- Message propagation over subnetworks. If a peer needs to send the same message to all other members of the group, it can use the multicast transmission provided by the link-layer. If multicast transmission is not available (as in the case of BT or BLE networks), or if routers prevent the transmission of multicast packets between the subnetworks, the P2P framework can anyway propagate the message, transmitting a copy of the message to each peer through unicast connections in a single transaction. This is possible as the coordinator nodes maintain a list of all the peers currently present in the peer group.

This paper provides two main contributions. The first one is EmbJXTA-Chord, a novel peer-to-peer (P2P) architecture specifically devised for developing IoT applications over heterogeneous networks. EmbJXTAChord integrates the good features of different sources, such as JXTA, EXI, CoAP, combining and augmenting them to provide a framework for the development of IoT applications that is able to offer several nice properties, such as:

- Service architecture based on a RESTful-API;
- Distributed and fault-tolerant service discovery;
- Support for secure *nested* peer groups;
- Routing over subnetworks transparent for the application layer, that allows to hide the presence of gateways;

---

<sup>2</sup>For an IP-only homogeneous network the host resolution by name can be provided by using a protocol such as ZeroConf. However, ZeroConf is unsuitable for a heterogeneous network due to its multicast-based design [37].

- Support for HTTP tunnelling and NAT traversal;
- Availability of a distributed content management system integrated in the system (based on advertisements);
- Agnosticism about the network protocol that is actually used.

To the best of our knowledge, there is no other open-source solution that supports all these features together. The implementation of the proposed architecture "mimics" the good properties of JXTA (without being JXTA) and integrates several technologies that were recently developed in the IoT field. For instance, EXI and a new compression protocol allow to reduce the bandwidth required for transmission, CoAP allows the realization of a RESTful architecture, the new MTB enables to support new transport protocols. In EmbJXTAChord such technologies work transparently for the developer, thus realizing a fully-integrated framework that simplifies the development of IoT applications.

In particular, EmbJXTAChord uses a compression algorithm that allows to extend some interesting features borrowed from JXTA (such as distributed and fault-tolerant resource discovery, transparent routing over subnetworks, application protocol independence from the transport protocol) in narrowband Wireless Sensor Networks (WSN), thus eliminating the need for dedicated software or custom gateways.

Moreover, EmbJXTAChord not only supports TCP/HTTP, but adds native support for Bluetooth RFCOMM and 6LoWPAN. This paves the way to its adoption in a broad range of IoT devices in supernetworks composed of subnets not necessarily based on IP. EmbJXTAChord solves the interoperability issues of the physical and transport protocols, as such protocols are transparently managed by the underlying P2P overlay layer.

EmbJXTAChord also fosters security over heterogeneous networks, adding support for secure peergroups (even nested) and for group encryption, thus allowing for unicast and multicast communication between groups of objects sharing the same resources.

The second main contribution of the paper is jxCOAP-E, a module of EmbJXTAChord that realizes a new version of CoAP that allows to deploy *heterogeneous networks of things* consisting of several subnetworks that are seen by the application level as a whole RESTful system, regardless of the actual connection topology or the available bandwidth. JxCOAP-E provides a distributed, fault-tolerant resource discovery paradigm and improves the CoAP multicast security model supporting *group encryption*, that allows to create *groups of objects* sharing the same resources, unlike UDP CoAP, which uses a centralized server for resource discovery and supports only unicast secure connections via DTLS.

The paper is organized as follows. Sect. 2 recaps JXTA benefits and limitations for implementing IoT applications over heterogeneous networks, while Sect. 3 addresses related work. Sect. 4 describes the EmbJXTAChord features, while Sect. 5 presents the experimental results and an assessment of the performance achievable by EmbJXTAChord using the Raspberry PI or the Raspberry PI-3. Sect. 6 deals with a scenario suitable for the proposed solution.

Sect. 7 presents two small examples of EmbJXTAChord programming, thus showing the effort required to the developer for using the middleware. Sect. 8

	JXTA protocol	Description
ERP	Endpoint Routing Protocol	It routes JXTA messages over different subnetworks, even if separated by gateways and NATs, working transparently to the upper layers of the middleware.
PRP	Peer Resolver Protocol	It is the connectionless communication protocol that allows the exchange of <i>messages</i> between nodes ( <i>peers</i> ) belonging to different subnetworks. Both unicast and multicast communications are supported.
RP	Rendezvous Protocol	It is executed by all rendezvous peers. The protocol manages a Distributed Hash Table (DHT) that allows to find all advertisements of the peergroup.
PDP	Peer Discovery Protocol	It works in synergy with the Rendezvous protocol. It provides an API that allows to find the advertisements describing the peers, peergroups and services available within the peergroup.
PBP	Pipe Binding Protocol	It allows a reliable connection-oriented communication between a pair of nodes of the peergroup, regardless of the features of the underlying transport protocol.
PIP	Peer Information Protocol	It provides statistics about network traffic among peers.

Table 2: Description of EmbJXTAChord protocols. These modules are reimplementations of the original JXTA 2.7 protocols.

provides experimental assessments of the power consumption for a single node in some typical use cases.

Finally, Sect. 9 gives conclusions and hints for further work.

## 2. EmbJXTAChord vs JXTA

### 2.1. Overview of JXTA

The first version of JXTA was released by Sun Microsystem (today Oracle) in 2001, in order to allow communications over heterogeneous networks consisting of several subnets based on different transport protocols and addressing schemes. JXTA consists of 6 protocols, whose functionalities are summarized in Tab. 2.

JXTA creates an overlay abstraction layer that manages all peers (i.e., nodes) through a uniform 128b address (named PeerID).

The peers are organized in *peergroups*. As represented in Fig. 2, peergroups can be deployed (i.e. created) regardless of the communication technology used between nodes. Moreover, peergroups can be nested. The support for peergroups allows to protect resources from unauthorized accesses. For instance, in Fig. 2 the peer named *WSN\_e2* can use the resources available in the PeerGroup1 (named NetPeerGroup), while it cannot access to the resources available in the PeerGroup2. Conversely, the peers *WSN\_e0* and *WSN\_e1* can access to the resources available both in the PeerGroup1 and in the PeerGroup2.

Each JXTA peer can work in one of the following modes: *adhoc*, *edge*, *rendezvous* (rdv), and *relay* [38].

The *edge* mode is commonly used by all nodes that work as clients, as it loads all standard JXTA protocols, but requires an active connection to a rendezvous peer (rdvpeer) before starting. The *rendezvous* nodes also provide the functionalities needed for maintaining operative the JXTA infrastructure, as they store part of a register, named SRDI (Shared Resource Distributed Index), that is used to implement distributed resource discovery within the whole group [39]. The *adhoc* nodes are resource constrained devices providing only a minimum set of functionalities. The *relay* nodes are similar to the edge nodes, but they can propagate multicast messages among two subnetworks regardless of the support for multicast datagrams at network level.

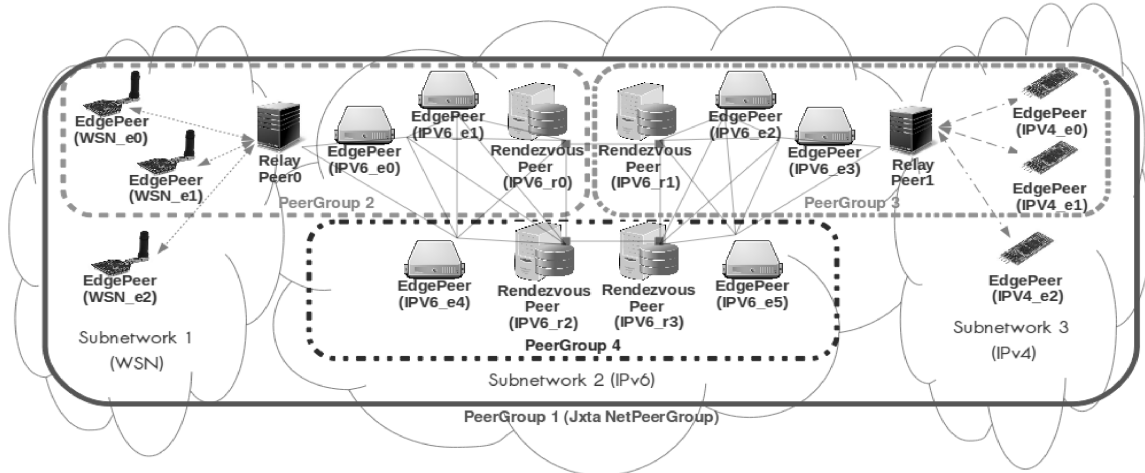


Figure 2: Representation of a heterogeneous network using EmbJXTAChord.

Two peers can communicate using virtual channels, named pipes for asynchronous unicast connections or jxta-sockets for synchronous unicast connections. Pipes can be reliable or unreliable, whereas sockets are always reliable. In addition, JXTA provides also propagate pipes, which ensure multicast communications within a group. A content management service based on XML documents (xmldocs), named advertisements (adv), is integrated in the protocol. When an advertisement is published by a peer, it can be found by all the other peers of the group through the Peer discovery protocol (PDP), which provides the discovery of new resources (services, nodes, contents etc.) within a group of peers.

## 2.2. What EmbJXTAChord borrows from JXTA

EmbJXTAChord inherits the following features of JXTA:

- *Independence from the transport protocol.* EmbJXTAChord can potentially use any transport protocol, using dedicated modules, named Message Transport Binding (MTB) modules. Besides the MTB modules for TCP-IP, HTTP tunnelling (that allows the traversal of firewalls and NATs) and multicast UDP, that were supported also in the latest version of JXTA (v2.7) [40], EmbJXTAChord integrates MTB modules for Bluetooth RFCOMM and 6LoWPAN communication protocols (see Sect. 4.1), thus supporting a wide range of devices that are typically used in IoT applications;
- *Support for heterogeneous networks.* EmbJXTAChord provides routing over subnetworks thus allowing communications between peers belonging to subnets using different addressing schemes and network protocols. Each peer can transparently promote itself as gateway (or sink). Operations as address translation or hop-by-hop delivery are automatically performed by the EmbJXTAChord overlay level;
- *Group security.* Peers can be grouped in peergroups, sharing a set of privileged access resources (pipes, jxCOAP-E services). The groups can



be nested, thus creating a hierarchy of privileged peers. The unicast and multicast communications within a group can be encrypted;

- *Group propagation.* A message can be propagated within a group even if the network level does not support multicast datagram propagation;
- *Distributed and fault-tolerant service discovery.* Unlike SOAP, CoAP and REST, which usually use a centralized server for service discovery (UDDI), EmbJXTAChord redundantly distributes all the group information in a set of *rendezvous peers*. This strategy, already adopted by JXTA, ensures high fault tolerance and allows to maintain the network consistency even under a high churn-rate (i.e. the rate of peers joining and leaving the group [39]).

The main engine of EmbJXTAChord is divided into 6 modules that reimplement the functionalities of the six protocols of JXTA 2.7. Thus, these modules maintain the same names of the corresponding ones in JXTA 2.7 (see Tab. 2 for the functionalities). For instance, the protocol that manages the resource discovery in EmbJXTAChord is named PDP (Peer Discovery Protocol) as in JXTA 2.7. For sake of clearness, we will refer to these components as *JXTA protocols* both in EmbJXTAChord and in JXTA 2.7 cases. Moreover, EmbJXTAChord maintains the message structure of JXTA 2.7 (divided in MessageElements, see Sect. 4.3). For this reason, in the rest of the paper we will refer to *JXTA messages* both in EmbJXTAChord and in JXTA 2.7 cases.

### 2.3. Differences between EmbJXTAChord and JXTA

There are at least three differences between EmbJXTAChord and JXTA:

- *A new API, named SJXTA* (simplified JXTA) (see Sect. 4). One of the main limitation of JXTA is its very complex Application Programming Interface (API). The SJXTA API manages several operations (peer, group or service discovery, pipe opening, creating or joining a new group), that are performed by the EmbJXTAChord engine providing few informations, thus simplifying the developer's job;
- *The new Hypercompression algorithm* (see Sect. 4.3). A second important drawback of JXTA is its high bandwidth consumption. JXTA messages are encoded as bandwidth-consuming xmldocs, JXTA peer, peergroup, and pipe IDs are encoded as strings. Most items are redundant, i.e., they are repeated several times in the same message or in multiple consecutive messages. The algorithm used to manage rendezvous peers (named *loosely-consistent Distributed Hash Table limited range walker*)(lcDHT-LRW) [39] and the source routing adopted for delivering messages over different subnets are bandwidth-expensive [41]. EmbJXTAChord uses the Hypercompression algorithm, that reduces the bandwidth required for message transmission, working together with a Rendezvous protocol (see Sect. 4.2), based on the more bandwidth-efficient Chord [42] algorithm;
- *The new jxCOAP-E component.* JxCOAP-E allows the communication between server and client peers through a standardized (and well-known) RESTful interaction model. Moreover jxCOAP-E allows to port, in a very simple way, the applications originally written for TCP-IP homogeneous

networks into the EmbJXTAChord hybrid environment. JxCOAP-E overcomes another important limitation of JXTA, i.e. the lack of a service interaction model. In fact, JXTA specifications include an underdeveloped concept of *jxta-service*, but nothing is said about *service interaction*. This raises the need for writing application-specific source code [40].

### 3. Related work

In the last years, researchers proposed several solutions aimed to create a service architecture over a heterogeneous network. The most interesting approaches were the use of P2P protocols, modular middleware solutions and middleware that implements the Web of Things (WoT) paradigm.

Moreover, other recent advancements in Bluetooth Smart, Vehicular and LTE technologies can be used in the proposed solutions, in order to integrate also this kind of subnetworks.

About the use of P2P protocols for IoT applications, in [43] XMPP (Extensible Messaging and Presence Protocol) was proposed as a solution for implementing a heterogeneous network without the need for a gateway. XMPP supports multiple transfer protocols via extension modules, but it is unsuitable for narrowband networks, as it exploits a protocol for the transmission of binary arrays that is based on the bandwidth-greedy Base64 encoding scheme. Conversely, EmbJXTAChord supports a binary transfer mode that allows to optimize the bandwidth occupation.

Despite a theoretical IoT model based on JXTA was presented in [44], there are only few works that exploit JXTA for IoT applications. *JxSensor* is a project aimed to integrate JXTA in a WSN [45]. A set of sensors is linked to a sink, managed as a virtual peer, which can be addressed using the common JXTA methods. JxSensor is a translation gateway which runs on the sink that is placed between the JXTA network and the WSN. The JXTA requests (responses) are delivered by the computer to (from) the gateway, where a component, named MoteAdapter, translates them into a set of commands that can be interpreted by the WSN sensors and actuators. Only the commands for the actuators and the data gathered by the sensors are transmitted through the narrowband link, not the standard JXTA messages. Therefore, the sensors and actuators can exchange data with the rest of the JXTA network through the translation gateway, but they cannot use all the JXTA functionalities.

An early version of jxCOAP-E based on JXTA 2.7, named *jxCOAP*, was presented in [46], as a component of the Javascript runtime container *jxActinium*. However, jxActinium is not devised for narrowband WSN. There are multiple reasons for this, such as, a wasteful use of bandwidth. no support for Bluetooth and 6LoWPAN and the use of the old loosely-consistent DHT limited range walker (lcDHT-LRW) algorithm for the Rendezvous Protocol. In Sect. 5.3 a comparison between jxCOAP-E (the version in EmbJXTAChord) and jxCOAP (the jxActinium version in [46]) is provided for the sake of completeness.

In a modular middleware sensors and actuators are driven and controlled by the main server through specific drivers loaded from a cloud. The main server also converts, before transmission, the directives of the middleware into a set of short commands in a format understandable by the sensors or actuators. For

example, SenseWrap exploits modules named Virtual Sensors in order to communicate to sensors and actuators [47]. The Virtual Sensors hide the hardware and the network protocol that are effectively used for communication, performing protocol conversion and exposing to the clients a uniform interface based on TCP/UDP-IP. The configuration is realized through the Zeroconf protocol. Unfortunately, the work in [47] does not provide information about the security policies for the middleware. Conversely, EmbJXTAChord can directly work over other network protocols than IP, without any message encapsulation or conversion into IP packets, and does not need Zeroconf for autoconfiguration.

The solution proposed in [48], named Smart Home Gateway, exploits a central server running a modular system based on OSGi (Open Service Gateway initiative), in order to manage the whole set of connected devices. When a new device is connected to the home network, the smart gateway automatically loads from a server in the cloud the correct driver (such as a new OSGi component) for the related Controller Device. A Controller Device is a modem able to communicate to a subset of devices exploiting a specific communication protocol (e.g., X10, Insteon, ZigBee). The work assumes that each Controller Device can be connected directly to the smart gateway, as no support is provided for hop-by-hop delivery between subnetworks. As a consequence, a failure in the smart gateway immediately affects the whole system. Furthermore, the smart gateway manages the whole security protocol. Conversely, EmbJXTAChord supports hop-by-hop delivery between subnetworks based on different network protocols and also allows to create smart environments in which a rendezvous peer failure does not determine the failure of the whole system.

Moreover, unlike IP-based solutions such as Smart Home Gateway or SenseWrap, EmbJXTAChord provides support for both unicast and multicast secure communications, also implementing several access levels through peer groups.

Web-of-Things is a paradigm that uses the architecture, the protocols and the services used in the Web in order to discover, manage and integrate smart objects into the global Internet [49]. In the WoT vision, sensors and actuators should be managed using a simple Web browser. In the last years, many frameworks were presented in this field, aimed to efficiently interconnect devices belonging to different subnetworks. In general, each of these frameworks includes a web server, that runs one or more web services representing the state of sensors and actuators by means of a RESTful API, a supervisor that runs an execution engine, thus elaborating the gathered data, and one or more *smart gateways* that interface sensors and actuators using a lightweight protocol.

The most used protocols for WoT communication are CoAP [11], MQTT[12][15], SMQ[13] [14] and ActiveMQ [50] (see Tab. 3). In the development of EmbJXTAChord, we have chosen to create the jxCoAP-E component (i.e. CoAP over JXTA) because CoAP offers several advantages over other IoT protocols. For instance, comparing with the Message Queue Telemetry Transport (MQTT) [15][12], Simple Message Queries (SMQ) [13] [14] and ActiveMQ [50] protocols, CoAP provides a request/response communication model that does not require a central broker.

However, CoAP, SMQ and MQTT are based on IP, so they inherit all the limitations about connectivity in the heterogeneous networks described in Sect. 1. As a result, differently from jxCOAP-E, they cannot leverage on features

	jxCOAP-E	CoAP	MQTT	SMQ	ActiveMQ
Supported transport protocols	TCP,UDP,HTTP, RFCOMM,SICS	UDP IP	TCP IP	TCP IP	TCP, UDP, HTTP, XMPP Websockets
Secure transport protocols for unicast communications	TLS over JXTA MTBs (TCP,HTTP, RFCOMM,SICS)	DTLS	TLS	TLS	HTTPS, TLS Secure WebSockets
Support for secure nested peergroups	AES group encryption	-	-	-	-
Communication model	Request-Response Publish-Subscribe	Request-Response Publish-Subscribe	Publish-Subscribe	Request-Response Publish-Subscribe	Publish-Subscribe
One-to-one communications	Supported	Supported	-	Supported	-
Multicast communications	Supported	Only if multicast UDP available [30]	-	-	Supported
Secure multicast communications	Supported (AES group encryption)	-	-	-	Supported
Central broker required	No	No	Yes	Yes	Yes
Distributed service discovery	Supported	-	-	-	Yes
Routing over subnetworks	Supported	-	-	-	-
In-band compression	Supported	-	-	Only topic names	-
NAT traversal	Supported (via JXTA PRP)	-	-	Supported (via WebSockets)	Supported (via WebSockets)

Table 3: Comparison between jxCOAP-E and other IoT communication protocols.

such as routing over subnetworks, peergroups or secure multicast connections. Moreover, the central broker used by MQTT is a potential point of failure and it can become overloaded when too nodes are active at the same time. Conversely, jxCOAP-E does not require a centralized broker and allows to use simultaneously more servers on different peers within a peergroup.

Node-RED [51][52] and WoTKit processor [53] are WoT platforms that were developed to efficiently interconnect devices belonging to different subnetworks (see Tab. 4). They provide an execution engine that runs programs where software modules and devices are represented as entities connected by wires. WoTKit provides the Dashboard, a browser-based interface aimed to manage and run the programs. FRED (a Frontend for Node-RED) [54] is a commercial version of Node-RED (produced by SenseTecnic) that supports the simultaneous execution of multiple flows in a multiuser environment.

EVERYTHING [55] is a commercial platform (provided by the Evrythng company) aimed to integrate whatever physical object in the cloud. Evrythng associates each physical object to a Web Object that allows to remotely control the item using a RESTful interface. Evrythng provides also the THNGHUB gateway [56], that allows to integrate into the cloud non-IP based devices (such as Bluetooth or ZigBee sensors).

The WoTKit, Node-RED, FRED and THNGHUB gateways support several transport protocols for the communication with end-devices (see Tab. 4). However, gateways send the gathered data to the execution engine using IP-based protocols (such as MQTT or WebSockets). As a consequence, these frameworks are affected by all the limitations already described in Sect. 1 when they are used in a multi-hop heterogeneous network or when the support for nested secure peergroups is necessary. Moreover EVERYTHING and THNGHUB are fully

	EmbJXTA Chord	WoTKit [53]	Node-RED [51]	FRED [54]	Evrythng [55]	SenseWrap [47]	Smart Home Gateway[48]	HPS [57]
Supervisor Entity	Required only to deploy peergroups	ActiveMQ broker	MQTT broker	MQTT broker	Cloud managed by the company	SenseWrap server	Central OSGi server	HTTP Proxy Server
Embedded devices that can be used as gateway nodes	RaspPI	RaspPI Arduino	RaspPI Arduino	RaspPI Arduino	Every dev. able to run ThngHub	Every dev. able to run Virt. Sensors	-	Every dev. able to run BLE stack
Protocols/interfaces for end-devices supported	BT, ZigBee, Wi-Fi, RaspPI, GPIO	BT, ZigBee, Wi-Fi,	BT, ZigBee, Wi-Fi, ZWave, RaspPI, GPIO	BT, ZigBee, Wi-Fi, ZWave, RaspPI, GPIO	BT, ZigBee, Wi-Fi, ETH	BT, ZigBee, Wi-Fi, Sun SPOT	X10, ZigBee, Insteon, uPNP	BLE, GATT/ATT
Communication protocols (gateway nodes-to-supervisor)	JXTA (TCP/TLS, BT, SICS, HTTP tunn.)	ActiveMQ	MQTT (TCP/TLS) WebSockets	MQTT (TCP/TLS) WebSockets	MQTT, HTTP, TCP/TLS, WebSockets	Proprietary (over TCP)	-(OSGi server is directly connected to sensors)	GATT to HTTP translation
Communication protocols (supervisor-to-user)	JXTA (TCP/TLS, BT, SICS, HTTP tunn.)	HTTP	MQTT (TCP/TLS) WebSocket HTTP	MQTT (TCP/TLS) WebSocket HTTP	HTTP, HTTPS, REST	Proprietary (over TCP)	HTTPS, REST	HTTP, HTTPS, REST
Publish/Subscribe communication model	Supported (jxCOAP-E)	Supported (ActiveMQ)	Supported (MQTT)	Supported (MQTT)	Supported (MQTT, CoAP, WebSocket, HTTP)	Supported (proprietary protocol)	Supported (via OSGi Event Manager)	Supported (GATT notify followed by HTTP POST)
Multiple users on a single execution server	Supported	Supported	Not supported	Supported	Supported	No Execution Engine provided	Engine in OSGi server supports all users	No Execution Engine provided
Secure nested peergroups	Supported	-	-	-	-	-	-	-
Every peer can become a gateway	Automatic	-	-	-	-	-	-	-
Unicast comm. between end-devices	Supported (via JXTA PBP)	-	Supported	Supported	-	-	-	-
Multicast communications	Supported (via JXTA RP)	-	-	-	-	-	-	-
Secure multicast communications	Supported (via Group Encryption)	-	-	-	-	-	-	-
Service discovery	Distributed (jxCOAP-E servers discovered via JXTA PDP)	Centralized	Centralized	Centralized	Hubs discovered locally via mcast. Connect. to cloud req. for auth. and sync	Virtual Sensors discovered using ZeroConf	OSGi driver registers in the Central Server the new devices	Centralized (Discovery performed by the GATT central node)
Routing over subnetworks	Supported	-	-	-	-	-	-	-
In-band compression	Supported	-	-	-	-	-	-	-
NAT traversal in unicast communication.	Supported	-	Only via WebSockets	Only via WebSockets	-	-	-	-
Platform accessible via web browser	-	Dashboard	Dashboard	Dashboard	Dashboard	-	WebServices RESTful API	WebServices RESTful API

Table 4: Comparison between EmbJXTAChord and some alternative IoT middleware solutions.

dependent on a connection to the servers maintained by the Evrythng company. Conversely, EmbJXTAChord supports non-IP transport protocols but, differently from THNGHUB, it is able to create independent cloud architectures.

About the Bluetooth hybrid networks, in 2010 the Bluetooth Special Interest Group (SIG) released the version 4.0 of the BT specifications, thus defining a new standard named Bluetooth Low Energy (BLE, also known as Bluetooth Smart). From v4.2, Bluetooth Smart introduced new features such as HPS (HTTP Proxy Service) [57] and IPSP (Internet Protocol Support Profile) [58][23]. The HPS gateways provide to HTTP-to-GATT translation, thus allowing to read/write some features of the BLE sensors (named GATT *Characteristics*) using a RESTful interface. IPSP [58][23] allows to implement a 6LoWPAN network over BLE (6LoBLE). A BLE node (named *Router*) pro-

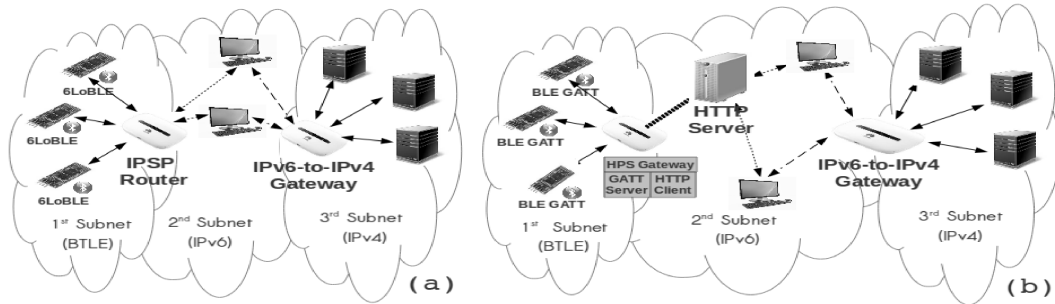


Figure 3: Representation of a heterogeneous network that includes BLE nodes (a) using the IPSP Router and 6LoWPAN over BLE (b) using the HPS gateway.

vides to reroute the 6LoWPAN packets coming from the other BLE nodes to the external (non-Bluetooth) subnetwork.

Unfortunately, IPSP supports only IPv6. Moreover, no support for routing over subnetworks is provided. As a consequence, both IPSP and HPS are unable to overcome the issues described in Sect. 1. For instance, in Fig. 3 a peer in the third subnetwork (IPv4-based) could not communicate to any BLE device in the first subnetwork. The only way to overcome this issue would be to install an IPv4/IPv6 gateway between the second and the third subnetwork and to reconfigure the routing tables in order to deliver the packets to/from the IPSP or HPS gateway. EmbJXTACHord, instead, is able to replace the functionalities of an IPv4/IPv6 gateway thus routing the messages between every node of the three subnetworks without these reconfiguration steps.

It is important to point out that EmbJXTACHord could work also over 6LoBLE, using the Message Transport Binding for IPv6. However, this solution would not be ubiquitous, because IPSP can be used only if it is supported by the operating system (for instance, IPSP is not supported by the Microsoft operating systems older than Windows 10 [24] and by the versions of BlueZ (Linux) older than v5.0). In order to measure the performance of EmbJXTACHord using BLE, we used the RaspberryPI-3 that integrates a chipset that is compliant with the v4.1 of the Bluetooth standard [59]. The results are shown in Sect. 5.

Other interesting enhancements for hybrid networks are related to the V2X (Vehicle-to-Everything) paradigm. The communication between vehicles and roadside units (RSUs) can be realized using DSRC (Dedicated Short Range Communication, a wireless technology based on the WAVE standard developed by the IEEE 1609 WG) [60], or LTE cellular technology [61]. 3GPP Release 12 introduced new specifications for low-power devices with low bandwidth requirements, such as *LTE-Direct* (also known as *ProSe*) [62] [63], aimed to promote LTE as a general solution for M2M communications. *LTE-Direct* defines procedures for device discovery and communication, thus allowing to group several nearby mobile phones into a cluster, controlled by a *clusterhead* (i.e. a mobile phone acting as a master). This way, LTE communications can work even when the base station is not available [62]. 3GPP Release 14 allows to use *LTE-Direct* as alternative for DSRC, or to connect the mobile phone with the vehicle (C-V2X, cellular V2X paradigm) [64].

EmbJXTACHord is not a competitor for LTE and DSRC, as it works at the application layer, while the previously cited technologies mainly work at the

PHY and MAC layers. Rather, EmbJXTAChord is a framework that can be advantageously integrated with LTE and V2X technologies.

As LTE-Direct is interoperable with Wi-Fi Direct, but not with BLE or IEEE 802.15.4, EmbJXTAChord can be used to integrate this kind of devices. Moreover LTE clusters cannot be nested, while EmbJXTAChord supports secure nested peer groups, thus allowing the transmission of a message only to a subset of the nodes in the LTE network.

EmbJXTAChord needs only the name of a peer to establish a connection, regardless of the change of the assigned IP address, that is a frequently occurring event in DSRC vehicular networks.

Finally, LTE and DSRC do not define any standard solution for content discovery and caching, despite it was observed that these functionalities can be advantageously used in vehicular networks [61]. Conversely, EmbJXTAChord makes available for the applications its own technology based on the Peer Discovery and Rendezvous protocols.

#### 4. EmbJXTAChord features

Fig. 4 shows the EmbJXTAChord architecture. EmbJXTAChord provides a new API, named SJXTA (Simplified JXTA), aimed to simplify the application development. The idea is that the developers should provide only the minimum set of information needed for a given operation. The following new concepts are introduced:

- *The network objects (netobj).* These are Java objects containing references to all the service instances and to all the configuration parameters of a group. When starting a new JXTA session, the developer initializes a standard netobj related to the JXTA NetPeerGroup [40]. The only parameters required are the *PeerName*, an optional *PeerProperties* string, and a 64b mapped *Options* parameter. Each group is associated to a netobj, therefore SJXTA automatically provides a new child netobj each time a new child group is created or is joined by a peer. Using the netobj methods, the developer can read and configure all the pieces of information about the current JXTA instance (such as *peername*, *PeerID* etc.), and can access to all the resources available in the group;
- *Advertisement generation.* Whenever a peer creates a new resource (peer-group, peer, pipe, socket or service), a new advertisement is generated and published. In JXTA 2.7, these complex operations are delegated to the developer. Conversely, in EmbJXTAChord the advertisements related to the new built resources are automatically generated and published by the SJXTA level;
- *Connecting and accepting pipes, connecting and accepting sockets.* In SJXTA sockets and pipes are addressable through their names, therefore there is no need for looking up the relevant advertisement within the group. When it is necessary to create a virtual channel allowing the connection by other peers, the developer instantiates a new *AcceptingSocket* (that provides an `accept()` method). The only parameters required for this operation are the netobj of the related peer group, the *SocketName*

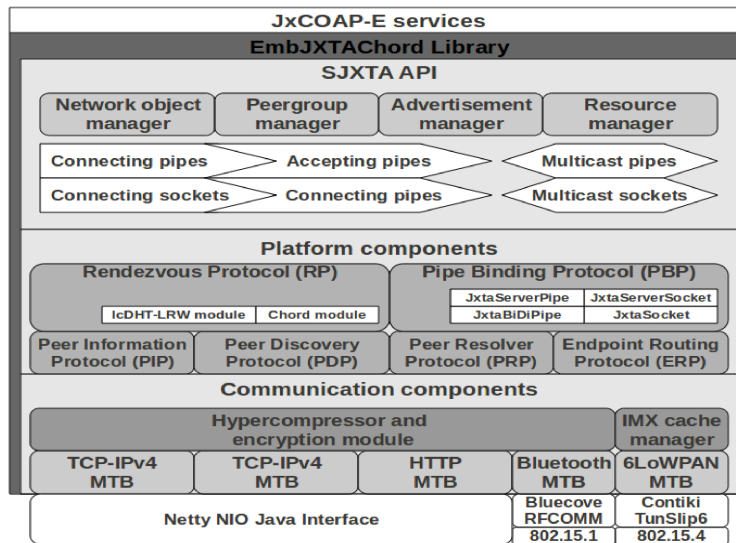


Figure 4: The EmbJXTAChord architecture.

and (optionally) a *SocketProperties* string. At the remote side, when it is necessary to access an accepting socket, whose name is known in the group, the developer creates a *ConnectingSocket* object and calls its `connect(SocketName)` method. SJXTA looks for an advertisement with the given *SocketName* and then establishes the connection. SJXTA sockets support only synchronous data transmission. Conversely, accepting and connecting pipes also support asynchronous data transfer (i.e. through a *listener*);

- *Multicast pipes and multicast sockets.* These objects allows to propagate data towards all peers of the group;
- *Resource discovery.* Any EmbJXTAChord resource is addressed by name, as the SJXTA level transparently provides to look up the related advertisement. Moreover SJXTA provides methods for *resource discovery* that return a name list of the resources that are compliant with a set of requirements. The developer can run the look up operations using a first filter on the resource type (*peer*, *group* or *pipe*), a second filter on a regular expression on its name, and a third filter on its properties.

#### 4.1. The new Message Transport Binding modules

While JXTA 2.7 only supports TCP-IP, UDP-IP or HTTP-tunnelling connections, EmbJXTAChord introduces two new Message Transport Binding (MTB) modules, aimed to support Bluetooth and IEEE 802.15.4 WSN networks.

The Bluetooth Message Transport Binding exploits Bluecove [65], a Java development library implementing JSR-82 specifications [66][67] that supports unicast connections to BT devices through a wide range of protocol stacks: Microsoft, Widcomm or Bluesoleil (under Windows) and BlueZ (under Linux) [68]. The mandatory protocol RFCOMM (serial port emulation) is used, thus



supporting up to 60 simultaneous connections and ensuring the widest compatibility<sup>3</sup>. Each device is addressable through its BT name or 96b address (`btsp://00228372FFC0` is an example of JXTA EndpointAddress for a BT device).

The SICSLOWPAN Message Transport Binding was designed for supporting WSNs based on the IEEE 802.15.4 standard. It is used, for instance, by the ZigBee protocol, but since many available ZigBee stacks are commercial, API-incompatible or with unmaintained sources (such as FreakZ or Open-ZB [70]), here a workaround was adopted. If the target EndpointAddress is of the `sicslowpan://[ipv6addr]` type, EmbJXTAChord establishes a connection creating an IPv6 tunnel to the serial port of the local mote. This is a device, equipped with a microcontroller and an IEEE 802.15.4-compliant transmitter, that runs a bridge application based on Contiki (an operating system for tiny devices that supports RPL routing and 6LoWPAN transport protocols) [71] [72] [73]. This solution ensures a wide hardware support, as any Contiki-compatible mote can be used. SICSLOWPAN MTB configures the mote at startup, requesting its link-local IPv6 address, which is registered in the peer advertisement in order to allow communication exchanges with other nodes of the WSN. Moreover, the MTB runs an algorithm for segmentation and reassembly that allows for multicast propagation of JXTA messages over the WSN, thus overcoming the limitation on the maximum size of the IP packets in 6LoWPAN (i.e., not larger than 1280B [74]).

#### 4.2. The rendezvous protocol

EmbJXTAChord replaces the bandwidth-greedy lcDHT-LRW rendezvous protocol with an alternative implementation based on Chord [42], an algorithm based on the theory of consistent hashing. This solution, which was successfully tested in JXTACh [75], is able to effectively improve the look up time of the advertisements and the bandwidth usage, as it entails less network traffic. Unlike lcDHT-LRW, Chord does not require that the local Rendezvous PeerView (RPV) (that in Chord is named *fingertable*) be exchanged between rendezvous nodes when no peer joins or leaves the group<sup>4</sup> (only small messages for *fingertable stabilization* and *predecessor checking* are periodically exchanged [42]).

Fig.5 shows the advertisement discovery process using the Chord Rendezvous Protocol. In EmbJXTAChord, each advertisement is characterized by a 128b hash value (key), but here in Fig.5 and in the case described in this Section, for a simpler representation, we assume that 32b hash values are used.

Each peer (rendezvous or edge) manages a local cache containing the advertisements recently discovered/published and whose lifetime is not yet expired. In Fig. 5 each rdvpeer is characterized by:

- a 32b *Chord identifier* (128b IDs are used in the real implementation) that

<sup>3</sup> Despite Bluetooth might be supported also using BNEP (network emulation), this is not a general solution, as some operating systems, such as Windows XP, do not support the piconet Group Ad-hoc Network role [69] or multiple simultaneous connections.

<sup>4</sup> Under JXTA 2.7 each rdvpeer selects, every  $T_{LRW}$  seconds, a *random sized* subset of the other rendezvous nodes (randomly chosen) where to send its own copy of the current RPV (RPV-exchange). This process is aimed to obtain the convergence of the RPV tables of all the rdvpeers in the group (see [39] for further details).

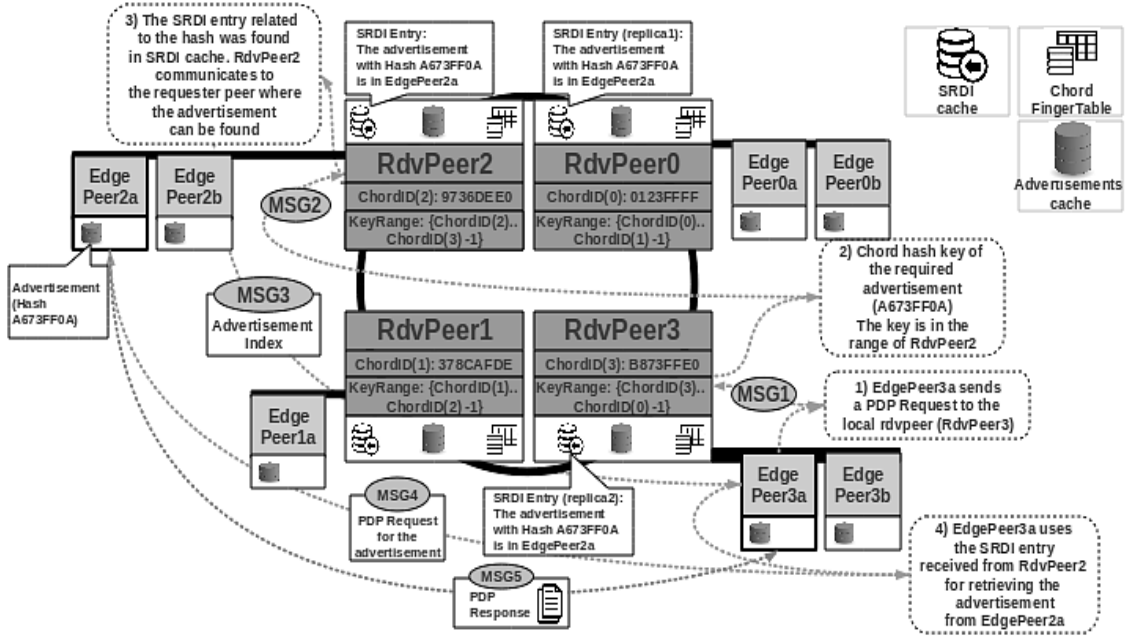


Figure 5: A representation of advertisement discovery using Chord RP protocol.

is computed when the peer joins the group by hashing the peername and the endpoint address;

- a *key range*, that defines the range of 32b keys (128b in the real implementation) for which the rdvpeer is responsible in Chord circular domain.

Each rdvpeer is connected to one or more edgepeers (for instance, in Fig.5 EdgePeer2a and EdgePeer2b are linked to RdvPeer2).

When an edgepeer (*publishing edgepeer*) needs to publish a new advertisement, the following steps are performed:

- The publishing edgepeer stores the advertisement in its own local cache and sends a PDP query to the local rdvpeer (*publishing rdvpeer*);
- The publishing rdvpeer computes the hash key of the new advertisement and performs a Chord look-up operation [42] aimed to determine the *target rdvpeer*, i.e. the peer responsible for the hash key. For instance, in the scenario shown in Fig.5, if the 32b hash of the advertisement is A673FF0A, the target rdvpeer is RdvPeer2 because it is responsible for the key range {9736DEE0..B873FFDF};
- The publishing rdvpeer contacts the target rdvpeer, that stores a new SRDI entry (*advertisement hash, JXTA PeerID of the storing node*) in its own local SRDI cache (this new entry is named *advertisement index* in JXTA terminology [39]);
- The publishing rdvpeer contacts also the two rdvpeers closest to the target rdvpeer in the Chord circular domain, thus registering two replicas of the

SRDI entry, in order to implement JXTA fault-tolerant strategy [39]. For instance, in Fig.5 the two replicas of the SRDI entry are stored in the cache of RdvPeer1 and RdvPeer3.

It is important to observe that advertisements are never transferred through the Chord domain. Only SRDI entries need to be transferred when a new rdvpeer joins or leaves the group. Fig. 5 shows the steps performed when a peer (in the example *EdgePeer3a*) starts the discovery process of an advertisement:

- EdgePeer3a contacts (MSG1) the local rdvpeer (*discovering rdvpeer*, RdvPeer3 in the example);
- The discovering rdvpeer computes the hash key of the required advertisement (A673FF0A in the example) and performs a Chord lookup operation, thus determining the *target rdvpeer* responsible for such key (RdvPeer2). It was demonstrated [42] that the lookup operation requires, with high probability,  $O(\log_2 N)$  hops within the Chord domain, if all the fingertables of the  $N$  rdvpeers are already stabilized;
- The discovering rdvpeer contacts (MSG2) the target rdvpeer that looks for the SRDI entry related to the hash in its own SRDI cache (in the example in Fig. 5 the entry is immediately found);
- The target rdvpeer returns (MSG3) the found SRDI entry (i.e. the advertisement index) to the discovering edgepeer (EdgePeer3a);
- The index returned to EdgePeer3a contains the ID of the peer *where* the advertisement is stored. Finally, EdgePeer3a contacts EdgePeer2a (MSG4), thus retrieving the document (MSG5).

If the Chord fingertable of some rdvpeers is not stabilized yet, or if some rdvpeer is unreachable, the discovering rdvpeer might not find the necessary SRDI entry in the cache of the target rdvpeer. In such a case, a strategy based on a limited-range walker [39] is performed, thus looking for the replicas of the SRDI entry in the near rdvpeers of the Chord domain.

The cost for maintaining DHT consistency is evaluated in Sect. 5.4.

#### 4.3. Hypercompression algorithm

EmbJXTAChord uses a new binary message format. Each message consists of several namespaces (ns), each of which contains several *MessageElements* (msgelem), featured by a name, a MIME type <sup>5</sup>, a content and (optionally) a signature. EmbJXTAChord inherits such message structure from JXTA 2.7 (version 1.0 of the JXTA specifications [40]). However, EmbJXTAChord and JXTA encode messages using two different binary formats, therefore a peer that runs EmbJXTAChord cannot communicate with a peer that runs JXTA 2.7.

The binary format used by JXTA 2.7 suffers from several drawbacks:

---

<sup>5</sup>EmbJXTAChord and JXTA 2.7 indicate the content of a MessageElement (*text/plain*, *application/octet-stream* etc..) using the media types defined by the MIME (Multipurpose Internet Mail Extensions) standard [76].




Figure 6: The architecture of the EmbJXTAChord compressor manager

- *XML and string contents are uncompressed.* In JXTA 2.7, the bandwidth is mainly used for names, MIME types and xmldocs, which are transmitted in text form;
- *JXTA PeerIDs are transmitted as long strings.* Each 128b ID requires bytes for the prefix (for example `urn:jxta:cbid-`), followed by 32 characters;
- *Endpoint addresses are uncompressed.* Addresses as `tcp://x.y.z.w:port` are transmitted as strings. In some cases, this can be expensive. For instance, an IPv6 address (128b, i.e. 16B in binary format) may require more than 40B if transmitted as an UTF-8 string;
- *The elements belonging to the same message are often redundant.* IDs and addresses for peers and peergroups can be repeated several times in the same message.

EmbJXTAChord exploits a Compressor Manager (CM) (see Fig. 6) that works cutting out both *intramessage* and *intermessage* redundancies. It intercepts all outgoing (or incoming) messages, transparently compressing (uncompressing) them.

In the first compression step (S1) the outgoing JXTA message is passed to a *preprocessing module*, that deals with some large MessageElements (Chord-Walker (CWL), EndpointRouterMessage (ERP), EndpointDestinationAddress (EDA) and JxtaWireHeader (JWH)), splitting them into a sequence of single partial “subelements” (pipe, group or peer IDs, network addresses, 128b keys) in the form `prefix://id`, that are ready to be compressed by the next stages.

In the second step (S2), the system scans the content of all XML and string elements, looking for any sequence in the form `prefix://seq`, thus creating a *dictionary of custom well-known sequences* (DCWS) that depends on the specific message. All sequences found in the content of each msgelem are replaced with an alias consisting of a control char  $CH_1$  + the 8b ID identifying in DCWS the replaced custom sequence. Next, the operation is repeated using a *dictionary of standard well-known sequences* (DSWS) whose items are standardized and known a priori by all peers. The sequences found in the content of each msgelem are replaced with an alias  $CH_2$  + the 8b ID identifying in DSWS the replaced item. For each message, the DCWS is increasingly reordered and transmitted using a differential format, while the DSWS is not transmitted.

In the third step (S3), the CM performs the *intramessage compression* through the following operations:

- S3a: The pair (*name*, *MIME type*) is compressed into a 4B stream. The first byte is a control char  $CH_3$  declaring some compression options for the msgelem, the second byte identifies the MIME type, the third byte is an alias ID referred to a dictionary of *well-known msgelem names* and the fourth one is a 8b suffix code;
- S3b: Hexadecimal numbers (such as IDs, custom and well-known sequences, etc.) are packed in binary form;
- S3c: IPv4, IPv6 and BT addresses are binary-packed;
- S3d: Strings are compressed using the deflate algorithm (gzip). EmbJX-TACHord exploits a gzip dictionary containing several sequences that appear frequently in JXTA data stream (such as `urn:jxta:cbid-`, `jxta://cbid-..`);
- S3e: The ByteArray elements are compressed using the deflate algorithm (gzip);
- S3f: The xmldocs are compressed through OpenEXI, an open-source and free implementation for Java of the EXI encoder [77][78]. An optimized schema, named `jxta.xsd`, is used to improve the efficiency of the EXI+GZIP combination.

The compression of hexadecimal numbers and addresses is performed in two steps (s3b, s3c) because of the different binary formats used. For the numbers, the format contains some control bytes and the binary representation of the hexadecimal value. Special codes may be used to indicate that a well-known subsequence was contained in the main sequence before compression and the original position (the CM implements this function using a *dictionary of well-known binary subsequences* (DBS) that is similar to DSWS). For the addresses, a more complex, protocol-specific, binary format is used, that is able to indicate also the network prefix (`tcp://`, `btsp://` etc.) and the position of dots and colons.

In the fourth step (S4), the CM performs the *intermessage compression* (IMX)<sup>6</sup>. The receiver and the sender maintain a IMX cache containing the  $c_X$  xmldocs and the  $c_S$  strings recently sent (received) more often to (from) the other side. The cache is updated whenever a new message is sent (received) and increasingly reordered on the basis of the 64b content hash. When the transmitter finds in the local IMX cache the msgelem to send, the latter is replaced with a short 8b code  $K_C$  in the transmitted packet. The receiver gets the missing msgelem from the  $K_C$ -th position in its own local IMX cache. Special codes are provided for registering/purging elements, thus maintaining the caches of both sides coherent.

The final step (S5) consists in the compression of the whole CM chain output through a 2nd-level gzip deflater. The result is encrypted (if it is required, as explained in Sect. 4.5) and finally transmitted on the channel. An assessment

---

<sup>6</sup>As this step is computationally expensive, it can be advantageously used only for the slowest channels (it is used for SICS MTB, but not for TCP, HTTP and BT MTBs).

of the performance measured for the Compressor Manager is provided in Sect. 5.1.5.

#### 4.4. *jxCOAP-E services*

As said in Sect.1, EmbJXTAChord provides jxCOAP-E, a modified version of CoAP specifically devised to work over the JXTA communication components. JxCOAP-E allows to merge the features of CoAP (compactness and RESTful interaction) and JXTA (support for hybrid networks).

EmbJXTAChord leverages on the service discovery architecture defined by JXTA specifications [40], thus allowing a server peer to provide *custom JXTA services* to the other client peers of the group. The communication to a JXTA service is provided by the Peer Resolver Protocol (PRP). When the server starts a new service, a new resolver listener is registered in the PRP module. A client can access to a service provided by a server (which must belongs to the same group), through the PRP method `ResolverQuery()`, providing the server *PeerID*, the *ServiceName*, and the *ServiceParam* string. The server performs the operation required and answers using the method `ResolverResponse()`. The interaction model is stateless and connectionless, as message delivery through PRP is unreliable[40].

In order to ensure that all peers in the group are informed about the availability of a new service, the server peer publishes a *module specification advertisement* (MSA), containing the name and description of the service.

Differently from JXTA 2.7, that do not define rules on the interaction model used for services, thus leaving this complex task to the developer, EmbJXTAChord exploits the RESTful interaction model provided by CoAP.

When a service is deployed, a new instance of a *jxcoap*<sup>7</sup> *virtual server* is started and then bound to a new *JXTA low level service* (jxllservice), whose module specification advertisement is published within the group. When a node needs to access a service, it runs a new instance of a *jxcoap virtual client* (vclient), which encapsulates each CoAP request into a PRP message that is sent to the jxllservice of the server peer. The resources provided by the jxcoap server are addressed through strings named *jxURI* [46]. The client interacts with the server through an API consisting of 4 operations (jxGET, jxPOST, jxPUT and jxDELETE) that require a jxURI specifying the target resource.

The access to a jxcoap service consists of 3 steps. First, the client looks for the MSA advertisement using the SJXTA resource discovery API. If the peer is an authorized member of the group the server belongs to, the advertisement is found, thus retrieving the PeerID of the server and initializing a new virtual client instance (*binding*). EmbJXTAChord replicates on multiple rendezvous peers the reference to the node that stores the service advertisement (see Sect. 4.2), thus ensuring it can be found even if some rendezvous peers are unreachable. Next, the client retrieves from the jxcoap server the list of the available resources through a jxGET operation to the standard jxURI

---

<sup>7</sup>We hereby refer to *jxcoap* for properties and concepts that are common to both the implementations, i.e., to the jxCOAP proposed in [46] and to the jxCOAP-E proposed in this paper.

“./well-known/core”. Finally, the client accesses all the resources provided by the bound server. An assessment of jxCOAP-E performance is in Sect. 5.3.

#### 4.5. Secure peergroups

EmbJXTAChord supports a new feature named *AES group encryption*, that allows the creation of secure peergroups. This feature overcomes two limitations of JXTA 2.7:

- In JXTA 2.7 multicast communications and PRP message exchange cannot be protected [79];
- JXTA 2.7 supports TLS-based unicast connections (*secure pipes*) between any pair of peers, even if the underlying transport protocol is not TCP. However, as the JXTA specifications do not define a Certification Authority Service, it is up to the developers to manage the transmission of the X.509 certificates (public keys) [38].

EmbJXTAChord protects the PRP and multicast communications within a group in the following way. When a *rdvpeer* (named *group owner*) creates a new peergroup, it can optionally state a *Group Traffic Encryption Key* (GTEK), which is used by an AES-128 encrypter integrated in the 2nd level gzip deflater (step S5 in Sect. 4.3). Hence, only the peers that know such a password can join the group, find the resource advertisements and decode the PRP messages. GTEK may be a Preshared key, or it can be requested by the edge to the group owner via TLS before joining the child group (Certificate-based scheme).

EmbJXTAChord introduces also a new version of peergroup advertisement that is used by TLS for retrieving the public keys needed to establish a connection. When a new edgepeer is joining a new group, it looks for the peergroup advertisement that includes the PeerID and the X.509 certificate of the group owner. Next, the edge contacts the owner requesting the group fingertable, thus providing its peer advertisement and X.509 certificate. In this way, the group owner maintains a *central keystore* with a list of all the public keys of the connected peers (*group members*), while each member maintains a *peripheral keystore* with at least the public key of the group owner. Each pair of peers in the group can use the keystore in the group owner to acquire the peer advertisement (and thus the public key) of the recipient, thus being able to establish a secure unicast TLS connection.

#### 4.6. EmbJXTAChord in an ISO/OSI network stack

Fig. 7 shows a comparison between a novel architecture using jxCOAP-E over EmbJXTAChord and a classic RESTful architecture based on CoAP over UDP. The comparison is represented using two ISO/OSI network stacks. The stack assumes that the architectures work over an IPv4-based Ethernet network, therefore the physical, data-link and network layers are identical in both cases.

The transport layer (OSI layer 4) is different because CoAP works over an unreliable protocol such as UDP whereas jxCOAP-E can use several reliable protocols (TCP, HTTP, 6LoWPAN, Bluetooth RFCOMM).

The session layer (OSI layer 5) does not contain any component in both cases, because in Fig.7 we consider only the single-application case.




Figure 7: A comparison between jxCOAP-E and classic UDP CoAP ISO/OSI network stacks.

The presentation layer (OSI layer 6) for jxCOAP-E contains the EmbJXTAChord *communication components* (i.e. the Hypercompressor and the Message Transport Binding modules). They provide to encrypt/decrypt and to compress/uncompress JXTA messages. These components also provide some *overlay functionalities* that allow to use a uniform addressing scheme based on 128b JXTA PeerIDs, thus making transparent to the upper layer the transport protocol that is actually used. In the UDP CoAP case, the presentation layer provides only to encrypt messages through the DTLS protocol.

The application layer (OSI layer 7) for jxCOAP-E contains the EmbJXTAChord *platform components*. They are the six modules of the framework (see Tab.2) that provide the applications (and to the jxCOAP-E services) with the typical functionalities of the JXTA framework (resource discovery, unicast and multicast communications, traffic information).

In the case of UDP CoAP, the OSI layer 7 is much simpler. It consists of a RESTful service framework (such as Californium [80]) that implements the request/response and publish/subscribe CoAP interaction models. As CoAP works over an unreliable and connectionless transport protocol (UDP), the RESTful service framework includes a module for *block-wise transfer* [81] of large payloads, which provides functionalities such as segmentation, reassembly and retransmission. JxCOAP-E does not need such functionalities, as it works over reliable transport protocols.

#### 4.7. Support for heterogenous networks

Fig. 8 shows how EmbJXTAChord works over different transport protocols in a heterogeneous network. The *PeerBridge* node acts as a gateway between *Peer1* (connected through an Ethernet link using IPv4) and *Peer2* (connected through an IEEE 802.15.4 link using 6LoWPAN). The Endpoint Routing Protocol (ERP) determines (transparently to the application layer) that the PeerBridge node can be used as a gateway to reach Peer2 (this information is made



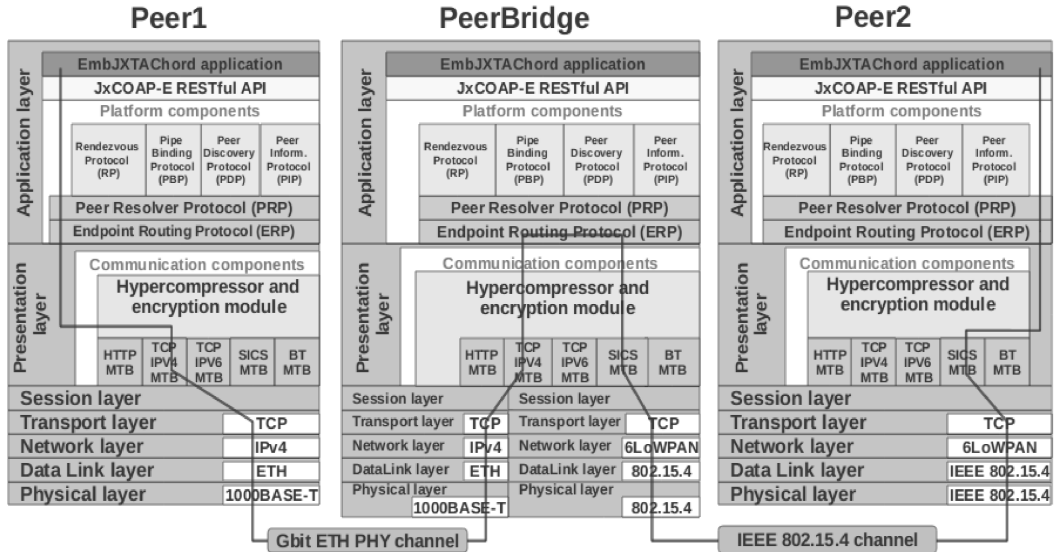


Figure 8: A representation of the communication between two peers belonging to different subnetworks, using OSI-inspired network stacks. *PeerBridge* acts as a gateway between the Ethernet network (IPv4) and the IEEE 802.15.4 WSN network (6LoWPAN).

available to all peers in the group using a route advertisement).

Once *PeerBridge* promoted itself to the gateway role, ERP automatically provides to uncompress/decode the messages coming from the Ethernet source node. Next it encodes, compresses and reroutes them to the IEEE 802.15.4 destination node.

The main advantage of using EmbJXTAChord in this scenario is that routing over subnetworks is automatically performed at the application level, without requiring the reconfiguration of the routing tables in the switches<sup>8</sup>, the installation of a *border-router* between the IPv4 and 6LoWPAN subnetworks [83] or the configuration of proprietary software for IPv4-to-IPv6 tunnelling [84].

## 5. Experimental results

The performance of EmbJXTAChord was tested under three respects. Sect. 5.1 compares the performance of EmbJXTAChord and JXTA 2.7. Sect. 5.2 measures the performance of an unicast connection (*connecting pipe*) between two peers, thus demonstrating that the proposed framework can work on hybrid narrowband networks (Bluetooth, 6LoWPAN), making transparent the presence of a gateway through routing over subnetworks. Sect. 5.3 deals with the performance of jxCOAP-E over homogeneous and heterogeneous networks, thus demonstrating that EmbJXTAChord can work on low-cost COTS hardware

<sup>8</sup>The configuration can be manually performed by the sysadmin, or using tools such as the *Router Advertisement Daemon* (radvd) that leverages on the IPv6 Neighbour Discovery Protocol (NDP) [82]. However, NDP is available only for the IPv6 subnetworks. Furthermore, radvd does not support multiple hops over three or more subnetworks. Finally, radvd is available only for Linux and its installation and configuration are not user-friendly tasks.

such as the RaspberryPI and RaspberryPI-3 boards with latency times that are acceptable for a wide range of applications.

### 5.1. *EmbJXTAChord vs JXTA 2.7*

Some tests aimed to compare the performance of EmbJXTAChord and JXTA 2.7 were performed. Sect. 5.1.1 measures the overhead determined by EmbJXTAChord encryption and compression respect to JXTA 2.7, in order to verify if the proposed framework is unsuitable for using on RaspPI or RaspPI-3 boards.

Sect. 5.1.2 measures the network overhead determined by Chord periodic message exchange, thus paving the way for the tests in Sect. 5.1.3 and Sect.5.1.4 that measure the message overhead in unicast and multicast connections. All these trials demonstrate the effectiveness of the Hypercompression algorithm in reducing the bandwidth waste respect to JXTA 2.7.

Finally, Sect. 5.1.5 measures the performance of the Hypercompression algorithm in function of the enabled compression schemes.

#### 5.1.1. *Execution times of common operations*

The latency values for some common EmbJXTAChord operations in an Ethernet network were measured using a modified version of the JxtaBench project [85].

Six testbed configurations (TB1-TB6) were considered. In the first one (TB1), the peer0 (server) was configured as *rendezvous* and the peer1 (client) was configured as *edge*. In the second one (TB2), both peers were configured as *rendezvous*. The peers were personal computers (PC) with an AMD Phenom II 3.0 Ghz CPU, running KUbuntu OS v15.10 and Oracle JRE1.7. All machines worked in single-core mode. In TB3-TB4 and in TB5-TB6 the peer roles were the same of TB1-TB2, respectively, but with a different hardware. The TB3-TB4 tests were performed using two Raspberry PI model B+, based on a single-core SoC with a 850 MHz ARMv11 CPU and 512 MB RAM [35]. The TB5-TB6 tests were performed using two Raspberry PI-3 model B based on the BCM2837 SoC, integrating four 1.2 Ghz ARMv8 CPUs and 1 GB RAM [36]. All RaspPI tests were performed using Raspbian OS and Oracle JRE1.8 for the 32b ARM processors.

Finally, all the tests were repeated with JXTA 2.7 (using *lcDHT-LRW*, without compression) in order to evaluate the overhead determined by EmbJXTAChord. During the test, the remote node created an accepting pipe and a new custom peergroup. The local node (*edge* or *rendezvous*) first joined the *NetPeerGroup* and the custom peergroup, then looked up for the other peers advertisements and pipes and, finally, connected to the remote pipe. Observing the execution times for the most common operations, that are shown in Tab. 5, some conclusions can be reached:

- Some operations on the server side (op1-3) and on the client side (op4-6) require very long times (especially on RaspPI, TB3-4). They might be still acceptable for some cases, as these operations are executed only once at startup. The reason for these long times is that each JXTA module can serve only the requests of a single group, therefore deploying or joining a new peergroup requires the JVM loads a new instance of all the active JXTA modules;

Testbed configuration	TB1	TB2	TB3	TB4	TB5	TB6	
server node	rdvpeer	rdvpeer	rdvpeer	rdvpeer	rdvpeer	rdvpeer	
client node	edgepeer	rdvpeer	edgepeer	rdvpeer	edgepeer	rdvpeer	
CPU	AMD Phenom2 (PC)	AMD Phenom2 (PC)	ARM11 (RaspPI mod. B+)	ARM11 (RaspPI mod. B+)	ARM11 (RaspPI-3 mod. B)	ARM11 (RaspPI-3 mod. B)	
Cores used for testing	1(3Ghz)	1(3Ghz)	1(850Mhz)	1(850Mhz)	4(1.2Ghz)	4(1.2Ghz)	
<b>Startup operations</b>							
EmbJXTAChord (compression and group encryption enabled, Chord)							
op1	start(server node)(*)	1312	1312	32501	32501	10040	10040
op2	deploy CustomPeerGrp(*)	525	525	1377	1377	708	708
op3	create a new pipe(*)	16	16	358	358	88	88
op4	start(client node)	1654	1671	26313	27029	6155	6691
op5a	init NetPeerGroup (rdv)	-	2009	-	27869	-	5597
op5b	join NetPeerGroup (edge)	2318	-	37452	-	17798	-
op6a	join CustomPeerGroup (rdv)	-	2766	-	22866	-	6604
op6b	join CustomPeerGroup (edge)	16743	-	24795	-	17243	-
JXTA 2.7 (compression and group encryption disabled, lcdHT-LRW)							
op1	start(server node)(*)	1179	1179	29678	29678	11812	11812
op2	deploy CustomPeerGrp(*)	39	39	860	860	224	224
op3	create a new pipe(*)	2	2	136	136	14	14
op4	start(client node)	1548	1320	31549	27534	6056	5394
op5a	init NetPeerGroup (rdv)	-	752	-	3928	-	3062
op5b	join NetPeerGroup (edge)	1647	-	28799	-	16620	-
op6a	join CustomPeerGroup (rdv)	-	1119	-	5695	-	3494
op6b	join CustomPeerGroup (edge)	16145	-	17019	-	16440	-
<b>Common operations</b>							
EmbJXTAChord (compression and group encryption enabled, Chord)							
op7	discovery other peers	13	32	1220	1292	134	255
op8	discovery remote pipe adv	18	27	712	861	141	230
op9	bind remote pipe	26	42	1347	1459	189	271
JXTA 2.7 (compression and group encryption disabled, lcdHT-LRW)							
op7	discovery other peers	10	11	680	290	51	95
op8	discovery remote pipe adv	6	10	255	283	57	60
op9	bind remote pipe	16	13	714	385	69	82

Table 5: Latency times for startup and common operations (ms)

The values were measured on the client node, except for the ones marked with (\*) that were measured on the server node.

- Some startup operations are longer for the edgepeers than for the rdvpeers. The edgepeers run NetPeerGroup joining (op5b) and CustomPeerGroup joining (op6b), that are performed contacting a local rdvpeer in order to register the new edgepeer, publish its own advertisement and retrieve the advertisement related to the custom peergroup using RP and PDP protocols (see Sect. 4.2). The rdvpeers run NetPeerGroup initialization (op5a) and CustomPeerGroup joining (op6a), that only require to initialize and stabilize their own RPV. They do not need to find the advertisement index of the custom peergroup using the RP protocol, as some SRDI entries are automatically transferred in their own cache when they join the Chord group (Chord key transfer, see [42]);
- The execution times of some common operations (op7-9) are higher for EmbJXTAChord than for JXTA 2.7. This is reasonable because of the overhead due to EmbJXTAChord compression and group encryption. The overhead for the op7-9 was lower than 20ms for PC (TB1-TB2), 1.1s for RaspPI (TB3-TB4) and 200ms for RaspPI-3 (TB5-TB6);
- For the rdvpeers, the execution times of NetPeerGroup initialization (op5a) and CustomPeerGroup joining (op5b) are longer in EmbJXTAChord than

	EmbJXTA Chord	EmbJXTA Chord	EmbJXTA Chord	EmbJXTA Chord	JXTA 2.7
	Compr. Config C0	Compr. Config C1	Compr. Config C2	Compr. Config C3	Compr. Config C4
Rdv protocol	Chord	Chord	Chord	Chord	LRW
Large MsgEl comp.(S1)	yes	yes	no	no	no
Well-known seq.(S2)	yes	yes	yes	no	no
String comp.(S3d)	yes	yes	yes	no	no
ByteArray comp.(S3e)	yes	yes	yes	no	no
XML EXI comp.(S3f)	yes	yes	no	no	no
Intermsg comp.(S4)	yes	no	no	no	no
GZIP deflater(S5)	yes	yes	no	no	no

Table 6: Compressor configurations

in JXTA 2.7, especially on RaspPI (TB3-TB4). This is mainly due to the initialization and stabilization of the 128b Chord Fingertable that is computationally expensive (see Sect. 4.2).

In conclusion, EmbJXTAChord can be used for the common operations also on low-cost COTS hardware such as RaspPI and RaspPI-3. The higher overhead comparing with JXTA 2.7 is the price to pay for the enhancements in compression and security. However, the use on the RaspPI is suitable only if the application does not require low startup times. Otherwise, the multicore RaspPI-3 (TB5-6) is preferable.

#### 5.1.2. Average DHT management overhead

In the second test, a number  $r$  of computers, with  $r = \{2, 4, 6, 8, 10\}$ , equipped with the same CPU used for the measurements described in Sect. 5.1.1, were connected to the same LAN, with the aim of measuring the *DHT management overhead*, i.e., the average bandwidth occupation due to the periodic message exchanges of Chord [42] or lcDHT-LRW [39].

All the computers were configured as *rendezvous* peers. Five configurations for the compressor manager (C0-C4, see Tab. 6) were used. The first one (C0) exploits all EmbJXTAChord features, whereas the latter (C4) is equivalent to JXTA 2.7 exploiting the lcDHT-LRW protocol.

The  $k$ -th rendezvous peer ( $k = \{0..(r-1)\}$ ) was started at the time instant  $t^{start}(k) = k \cdot 60s$ . The rendezvous peers sent two types of Chord service messages: *Chord stabilize* with periods  $T_{stab} = 8s$  and *Chord predecessor-checking* with periods  $T_{pred.chk} = 4s$ , respectively. Moreover, they performed a *fix-fingers* operation every  $T_{fix.fing} = 4s$  (these messages are required for checking that the rendezvous nodes are still alive or for the updating of Chord fingertable in the nodes, see [42] for more details). For JXTA 2.7, lcDHT-LRW was configured so that RPV exchanges occurred every  $T_{LRW} = 8s$ .

The Tab. 7 shows the average bandwidth overhead  $\phi_{dht-mgm}(r) = \Phi_{n.trx.bytes}/\Delta T$  where  $\Phi_{n.trx.bytes}$  is the number of bytes sent by rdv0 to the other rendezvous peers during the observation time  $\Delta T = 60s$  between two following joining events.

When EmbJXTAChord is used, increasing the number of rdvpeers determines higher values of management overhead. It is interesting to observe that this may be different for JXTA 2.7 (see for instance the case  $r=10$ ) because of the randomness determined by the RPV-exchange operation periodically performed by lcDHT-LRW algorithm [39] (see Sect.4.2).

	Nr rdv peers (r)	EmbJXTA Chord Compr. Config C0	EmbJXTA Chord Compr. Config C1	EmbJXTA Chord Compr. Config C2	EmbJXTA Chord Compr. Config C3	JXTA 2.7 Compr. Config C4
0 → all	2	0.201	0.245	0.858	1.468	1.119
0 → all	4	0.201	0.391	1.112	1.496	2.431
0 → all	6	0.215	0.533	1.841	2.766	3.721
0 → all	8	0.203	0.289	1.014	1.353	10.794
0 → all	10	0.182	0.358	1.180	1.566	1.606

Table 7: Average DHT management overhead (KB/s)

	Nr rdv peers (r)	EmbJXTA Chord Compr. Config C0	EmbJXTA Chord Compr. Config C1	EmbJXTA Chord Compr. Config C2	EmbJXTA Chord Compr. Config C3	JXTA 2.7 Compr. Config C4
0 → 1	4	0.097	0.176	0.529	0.532	1.930
0 → 2	4	0.000	0.000	0.000	0.000	0.286
0 → 3	4	0.104	0.215	0.582	0.964	0.215

Table 8: Average DHT management overhead between two rendezvous nodes (KB/s)

In the best case the average bandwidth overhead was reduced from 10.794 KB/s (JXTA 2.7) to 0.203 KB/s (EmbJXTAChord). It is interesting that the value for  $\phi_{dht-mgm}(r)$  measured for JXTA 2.7 sometimes can be lower than the one measured for EmbJXTAChord when all the compression schemes are disabled, as the lcDHT-LRW algorithm is more efficient than Chord when few rendezvous peers are used [75]. However, when all the compression schemes are enabled, EmbJXTAChord always outperforms JXTA 2.7, as it requires less bandwidth.

### 5.1.3. Average pipe message overhead

As JXTA adds some information to the data transmitted, a third test was performed in order to measure the *additional message overhead*  $\Delta m$  added to the payload when connecting or accepting pipes are used. A group of 4 rendezvous peers (rdvpeer 0,1,2,3) was deployed letting them join the group, in sequence, with a 60s offset from each other. When all 4 rdvpeers joined the group, the Chord predecessor of rdvpeer0 was rdvpeer1 and the successor was rdvpeer3. Since then, for a duration of  $\Delta T = 240s$ , the average *DHT management overhead*  $\phi_{dht-mgm(0 \rightarrow n)} = \Phi_{n.trx.bytes} / \Delta T$  for the messages sent from rdvpeer0 to rdvpeer-n ( $n = 1, 2, 3$ ) was measured (Tab. 8)<sup>9</sup>. As a further trial, the rdvpeer0 was connected to an *accepting pipe* created by rdvpeer3, after all the 4 rdvpeers had joined. The rdvpeer0 sent  $N$  messages containing  $\Phi_{payload.size}$  bytes (randomly generated each time) to rdvpeer3, and then closed the pipe. The parameter *transmitted-message-vs-payload-size ratio*  $\gamma$  is defined as:

$$\gamma = \frac{\Phi_{n.trx.bytes} - \phi_{dht-mgm(0 \rightarrow 3)} \Delta T}{N \Phi_{payload.size}} = \frac{\Phi_{n.trx.bytes}^*}{N \Phi_{payload.size}} \quad (1)$$

where  $\Phi_{n.trx.bytes}$  is the number of bytes sent by rdvpeer0 (measured by Wireshark) during  $\Delta T = 240s$  since the pipe connection and  $\phi_{dht-mgm(0 \rightarrow 3)}$  is the overhead due to DHT management measured in the previous trial.

<sup>9</sup>Using Chord protocol (C0,C1,C2,C3 configuration), there was no traffic towards rdvpeer2 as it was neither a predecessor nor a successor of rdvpeer0.

		payload size	EmbJXTA Chord Compr. Config C0	EmbJXTA Chord Compr. Config C1	EmbJXTA Chord Compr. Config C2	EmbJXTA Chord Compr. Config C3	JXTA 2.7 Compr. Config C4
$\gamma$	0 → 3	256B	2.081	2.363	3.353	4.790	4.859
	0 → 3	1024B	1.326	1.408	1.589	1.949	1.964
	0 → 3	10240B	1.035	1.041	1.063	1.074	1.096
$\Phi_{msg.size}$	0 → 3	256B	533B	605B	858B	1226B	1244B
	0 → 3	1024B	1358B	1442B	1627B	1996B	2012B
	0 → 3	10240B	10599B	10661B	10892B	11002B	11228B

Table 9: Transmitted-message-vs-payload-size ratio  $\gamma$  and average message size  $\Phi_{msg.size}$  (B) for accepting/connecting pipes (4 rendezvous peers)

		payload size	EmbJXTA Chord Compr. Config C0	EmbJXTA Chord Compr. Config C1	EmbJXTA Chord Compr. Config C2	EmbJXTA Chord Compr. Config C3	JXTA 2.7 Compr. Config C4
$\gamma$	0 → 3	256B	2.747	3.388	7.013	10.373	9.636
	0 → 3	1024B	1.478	1.622	2.492	3.338	3.159
	0 → 3	10240B	1.052	1.069	1.153	1.223	1.215
$\Phi_{msg.size}$	0 → 3	256B	703B	867B	1795B	2655B	2467B
	0 → 3	1024B	1514B	1661B	2552B	3418B	3235B
	0 → 3	10240B	10776B	10946B	11815B	12533B	12451B

Table 10: Transmitted-message-vs-payload-size ratio  $\gamma$  and average message size  $\Phi_{msg.size}$  (B) for multicast pipes (4 rendezvous peers)

$\Phi_{n.trx.bytes}^*$  is the number of bytes due to message transmission, therefore the average size of the transmitted messages  $\Phi_{msg.size}$  can be defined as  $\Phi_{n.trx.bytes}^*/N$ .

Replacing in eq. 1, the following condition holds:

$$\gamma = \frac{\Phi_{msg.size}}{\Phi_{payload.size}} \quad (2)$$

where  $\gamma$  is the *transmitted-message-vs-payload-size ratio* (if  $\gamma < 1$ ,  $\gamma^{-1} > 1$  measures the *compression efficiency* of the algorithm).

In the performed tests, the `rdvpeer0` transmitted  $N = 25000$  messages with a payload of  $\Phi_{payload.size} = 256, 1024, 10240B$ , using the 5 configurations reported in Tab. 6. The measured values for  $\gamma$  and  $\Phi_{msg.size}$  are reported in Tab. 9<sup>10</sup>. Using 256B payloads, the  $\gamma$  values measured for JXTA 2.7 and for EmbJXTAChord are respectively 4.86 and 2.08 (for resulting average message sizes respectively of 1244B and 533B). Hence, JXTA 2.7 appears to be unsuitable for the smallest messages as the overhead is too high with respect to the payload. Using 1024B payloads, the  $\gamma$  values measured for JXTA 2.7 and for EmbJXTAChord are respectively 1.96 and 1.32 (for resulting average message sizes respectively of 2012B and 1358B). Using 10240B payloads, the measured  $\gamma$  values are respectively 1.096 and 1.035 (for average message sizes of 11228B and 10599B, respectively). In conclusion, EmbJXTAChord improves the bandwidth occupation in all tested cases.

#### 5.1.4. Average multicast pipe message overhead

The fourth test measured the overhead related to *multicast pipes* (mpipe). EmbJXTAChord supports the multicast delivery of messages using the *Chord*

<sup>10</sup>As TCP is reliable, the *isReliable* option for JXTA pipes was always disabled before tests. For JXTA 2.7 data (C4 compr. mode), when `lcDHT-LRW` algorithm is used, the `rdvpeer3` can receive RP messages not only from `rdvpeer0` but from every rendezvous peer of the group. As a consequence, a different formula is used ( $\Phi_{msg.size} = (\Phi_{n.trx.bytes} - \phi_{dht-mgm(all \rightarrow 3)})/N$ ).  $\Phi_{n.trx.bytes}$  and  $\phi_{dht-mgm(all \rightarrow 3)}$  were measured using Wireshark at `rdvpeer3`.




Figure 9: Compression ratio (%) with different compression schemes.

*walker propagation*, whereas JXTA 2.7 exploit the LRW walker [75][39]. When a *rdvpeer* needs to propagate a content through its group, it sends the message to its successor which, in turn, retransmits *rdvpeer*-by-*rdvpeer* through the whole Chord circular domain. Next, each rendezvous peer retransmits the message to its connected edgepeers, thus realizing the propagation of the content through the whole group.

The trial was performed under the same conditions of the third test. The *rdvpeer0* created a *mpipe*, the *rdvpeers* 1-3 joined the group in sequence with a 60s offset from each other, and once Chord fingertable stabilization occurred (so that the successor of *rdvpeer0* became *rdvpeer3*), the *rdvpeer3* started to receive the data sent in multicast by *rdvpeer0*. The *rdvpeer0* propagated  $N = 25000$  messages with a payload of  $\Phi_{payload.size} = 256, 1024, 10240B$ , using the 5 configurations in Tab. 6. About the traffic between *rdvpeers* (0,3), the *average size of the transmitted messages*  $\Phi_{msg.size}$  and the *transmitted-message-vs-payload-size ratio*  $\gamma$  can be found through eq. (2) and (1).

A comparison between the values for  $\gamma$  reported in Tab.9 and Tab.10 shows that the overhead is larger for multicast pipes than for unicast ones. Moreover, the values for  $\gamma$  in Tab.10 show that EmbJXTAChord improves the bandwidth occupation also when multicast pipes are used.

### 5.1.5. Performance of the compression algorithm

The fifth test was aimed to measure the performance of the Compressor Manager (CM) using different compression schemes. A single 1024B message was created, consisting of 128 sequences, each made up of a random byte  $x$  repeated 8 times. The message was sent through a unicast, a multicast and a secure pipe between two rendezvous peers. The sizes of the uncompressed message were 2006B, 3443B and 2473B, respectively.

Fig. 9 shows the message sizes, measured enabling or disabling the different compression schemes (compression configurations CP1-CP8). The larger size reduction is provided by the EXI compressor (S3f) and by the compression of strings, addresses and numbers (s3b-s3d). With all the compression features enabled, the message can be compressed down to 31% (unicast pipe), 25% (mul-

link	TLS	sender	recv	1KB	2KB	4KB	10KB	1KB	2KB	4KB	10KB
				Raspberry edgepeers				PC edgepeers			
BT →BT	-	edgepeer0	edgepeer1	4.53	8.63	17.08	37.50	54.45	65.17	128.51	135.02
BT →BT	yes	edgepeer0	edgepeer1	2.34	4.02	8.21	13.82	35.38	44.44	94.95	81.70
BT →BT→ETH	-	edgepeer0	rdvpeer0	2.57	4.17	10.94	33.32	59.56	65.33	101.81	124.30
BT →BT→ETH	yes	edgepeer0	rdvpeer0	1.87	3.47	6.57	11.23	33.25	56.79	79.76	77.56
SICS→SICS	-	edgepeer2	edgepeer3	0.60	1.67	2.77	3.98	0.67	2.10	3.31	4.42
SICS→SICS	yes	edgepeer2	edgepeer3	0.12	0.27	0.51	0.74	0.48	1.07	2.02	2.33
SICS→SICS→ETH	-	edgepeer2	rdvpeer0	0.52	1.08	2.05	1.55	0.54	1.62	3.16	3.54
SICS→SICS→ETH	yes	edgepeer2	rdvpeer0	0.24	0.53	0.99	1.03	0.45	0.82	1.83	1.11
SICS→BT	-	edgepeer2	edgepeer0	0.63	1.87	2.24	3.16	0.66	1.99	3.26	3.40
SICS→BT	yes	edgepeer2	edgepeer0	0.25	0.92	0.80	1.95	0.35	1.06	1.11	2.27

Table 11: Transfer rate measured on heterogeneous network (KB/s)

link	TLS	sender	recv	1KB	2KB	4KB	10KB
				Raspberry-3 edgepeers			
BLE →BLE	-	edgepeer0	edgepeer1	49.1	91.1	126.3	120.9
BLE →BLE	yes	edgepeer0	edgepeer1	40.8	79.2	94.8	103.8
BLE →BLE→ETH	-	edgepeer0	rdvpeer0	43.5	86.1	112.2	115.3
BLE →BLE→ETH	yes	edgepeer0	rdvpeer0	34.3	76.2	81.3	93.2

Table 12: Transfer rate measured for RaspberryPI-3 using Bluetooth Low Energy (KB/s)

ticast pipe) and 41% (secure pipe) of the original size. On secure pipes the CM shows the worst performance, as compression is affected by the AES-128 encryption and by the cipher-block chaining (CBC) padding [86]. Moreover, the messages produced by EmbJXTAChord are shorter than the ones that could be obtained applying a gzip deflater to JXTA messages thus demonstrating the effectiveness of the Hypercompression algorithm.

### 5.2. Performance on Bluetooth-IEEE 802.15.4 heterogeneous network

In order to show that EmbJXTAChord can work on a narrowband architecture consisting of embedded devices, a test was performed on a heterogeneous network consisting of one PC (rdvpeer0), that communicates using the IPv4 protocol with another PC (rdvpeer1) acting as a sink and four RaspPI devices (edgepeers 0,1,2,3), connected to rdvpeer1 using RFCOMM over Bluetooth (BT v2.1) and 6LoWPAN, respectively (see Fig. 10). The rdvpeers (0,1) were linked through a Gigabit Ethernet wire. The edgepeers (0,1) were connected to rdvpeer1 through two USB BT adapters based on a Silicon Cambridge transmitter. The edgepeers (2,3) used two boards STM Dizum MB950 [87] equipped with a M3 Cortex CPU, 256 KB RAM and an IEEE 802.15.4 transmitter (the operating system was Contiki 2.7 with NullMAC layer<sup>11</sup>).

In the first trial, the edgepeer1 created an accepting socket that was used by the edgepeer0 to send  $N = 1000$  messages using a direct BT connection. Each payload contained  $\Phi_{payload.size} = 1024, 2048, 4096, 10240B$ , randomly chosen. The same procedure was repeated establishing a direct SICS connection between the edgepeer3 and the edgepeer2. In the second trial, the edgepeer0 (edgepeer2) sent data via BT (via SICS) to rdvpeer1 which acted as a *sink* for rdvpeer0. In the third trial, the data exchange was between the edgepeers (2,0) (the rdvpeer1 acted as a *bridge* between BT and SICS subnets, thus allowing communication between nodes not directly linked at physical level). The three

<sup>11</sup>ContikiMAC is not available for MB950, therefore NullMAC is the default option for this microcontroller unit.






Figure 10: The topology of the heterogeneous network used for testing.

trials were repeated using a secure socket based on the TLS protocol. Finally, the whole experiment was repeated replacing the RaspPI boards with computers, which were configured as in Sect. 5.1.1. Tab. 11 shows that the transfer rates  $\beta = N\Phi_{payload.size}/T_{trx}$ , where  $T_{trx}$  is the time needed to send all the messages, heavily depend on the nodes computational power, due to the compression and encryption overhead. A fair interpretation of the performance on 6LoWPAN must take into account that the throughput of the Contiki devices at the application level is consistently lower than the raw theoretical data rate of the IEEE 802.15.4 radio interface (256Kb/s). For example, in a work that measured the real end-to-end performance of some motes connected to a PC through a serial port, a real throughput between 32 and 70 Kb/s (i.e. 4 and 8.75 KB/s) for various models of transceivers is reported [88].

The best throughput is achieved using larger payloads. This suggests to pack multiple short messages into a single large one, to optimize the performance.

Finally, some trials were repeated using some RaspberryPI-3 equipped with a Bluetooth Smart transmitter (v4.1). Tab. 12 shows that RaspPI-3 with BLE largely outperforms the old RaspPI with the old Bluetooth v2.1.

### 5.3. *jxCOAP-E performance*

The last experiment, made up of several trials, was aimed to measure the performance of a server based on jxCOAP-E under several aspects.

#### 5.3.1. *Scalability test*

In the first trials, the scalability of the jxcoap server was measured. First a configuration consisting of an rendezvous PC *server* and  $s = [1..7]$  PCs (edgepeers) was deployed. Each edgepeer ran  $v = 10$  jxcoap virtual client instances (see Sect. 4.4) for  $T = 180s$ , therefore the rendezvous server responds to  $n = s*v \in [10..70]$  virtual clients in total during the trial. All the peers were equipped with the same AMD CPU used in Sect. 5.1.1 and connected via Ethernet. Each vclient sent a jxGET request containing a parameter  $\eta$  to a *coap* service<sup>12</sup> running on the server, then waited for a response before repeating the operation. All the trials were performed using the CoAPBench benchmark, modified to work over jxCOAP-E [80].

Four services were used: *HelloWorld*, *Fibonacci*, *Newton* and *SortSquare-Root*. The first service responded to a client request with a “Hello world” string.

<sup>12</sup> In this experimental section we refer to UDP CoAP for the original version proposed by CoRE WG and to *coap* for concepts that are common to UDP CoAP, jxCOAP and jxCOAP-E.

The second one returned a string containing the  $\eta$ -th number of the Fibonacci sequence. The third service returned an array containing the square roots of the first  $\eta$  integer numbers, determined by the server using the Newton method iterated  $\xi = 25000$  times and encoded as 8B double precision float numbers. The fourth service received a xmldoc from the client containing  $\eta$  float numbers and replied with another xmldoc containing the related square roots, determined by the server using its CPU floating unit and encoded as strings, increasingly sorted through the bubble-sort algorithm. When the same value of  $\eta$  is used, the size of the uncompressed payload included in the *coap* response is larger for the *SortSquareRoot* service than for the *Newton* one.

As these trials were aimed to measure the scalability of the *coap* server, only the number  $s = [1..7]$  of connected edgepeers varied, while all the requests were made with  $\eta = 25$ . All trials were repeated using UDP CoAP, jxCOAP-E within the *NetPeerGroup* and jxCOAP-E within a *peerGroup* secured by AES-128 (see Sect. 4.5). Fig. 11(a) shows that using UDP CoAP the *SortSquareRoot* service scaled better than the *Newton* service, as the second one requires a higher elaboration time. Conversely, Fig. 11(b-c) show that using jxCOAP-E the *SortSquareRoot* service was outperformed by the *Newton* service, as the first one entails a larger overhead due to EXI compression. The effect of AES-128 encryption on the jxCOAP-E response rate was minor.

Fig. 11 also shows that the response rate provided by the UDP CoAP server was much higher than the one provided by the jxCOAP-E server. This is a reasonable result, because *EmbJXTAChord* manages the complex structure of the JXTA messages [40], made up of several namespaces and *MessageElements* (see Sect.4.3), together with the operations for compression and encryption. In order to measure the overhead introduced by compression, a further experiment was performed. In an Ethernet network, the jxCOAP-E server was connected to  $s = 4$  edgepeers, each running  $v = 10$  *jxcoap vclient* instances. Each *vclient* connects to the *HelloWorld*, *Fibonacci*, *Newton* and *SortSquareRoot* services for  $T = 180s$ , using  $\eta = 25$  for all requests. Each trial was performed without and with AES-128 encryption. The experiment was repeated multiple times, each time changing the enabled compressor schemes (CP1..CP8) accordingly to the table in Fig.12(a).

Fig.12(b-c) show that the measured response rates (without and with AES encryption) increased more and more when multiple compression schemes were disabled, because of the lower overhead determined by the compression. The largest improvements were measured for the *HelloWorld* and *Fibonacci* services, that are characterized by the shortest elaboration times. Obviously, disabling multiple compression schemes determines a higher bandwidth consumption, but in the Ethernet case this does not significantly affect the performance.

Conversely, when a narrowband network is used, the compression schemes implemented by *EmbJXTAChord* can improve both the response rate and the bandwidth utilization. In order to prove this assertion, the previous experiment was repeated using a SICS network made up of a single jxCOAP-E server and  $s = 2$  edgepeers, each running  $v = 1$  virtual clients. Fig. 12(d-e) show the measured response rate. It is interesting to observe that in some trials the default CP8 configuration (all compression schemes enabled) is outperformed by the CP5 configuration (S1,S2,S5 compression schemes disabled, see Sect.4.3). This happens because the CP5 configuration in some cases (see Fig.9) can provide low size in transmission (but not the *lowest* possible size) together with low




Figure 11: Scalability measures for the jxcoop server. The graph shows the response rate measured by each edgepeer during the test. (a) Using UDP without compression. (b) Using jxCOAP-E (all compression schemes enabled). (c) Using jxCOAP-E within a group with compression and AES-128 encryption enabled.

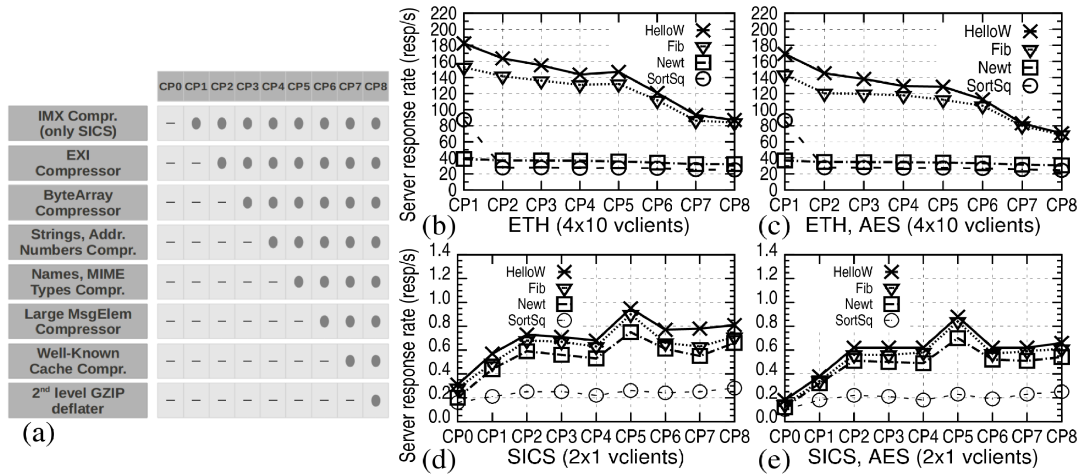


Figure 12: Scalability measures for the jxcoop server. The graph shows the average response rate measured by each edgepeer, using several compressor configurations. (a) The compressor configurations used for the trials. (b-c) Response rate for ETH trials without and with AES-128 encryption. (d-e) Response rate for SICS trials without and with AES-128 encryption.

compression overhead. The CP5 configuration can be used as an alternative to the default CP8 configuration when the application needs to optimize the response rate and the number of nodes connected to the sink peer is low ( $s = 2$  in the described test). Otherwise, if the minimization of the bandwidth required for transmission is mandatory as many nodes communicate at the same time to the sink, the default CP8 configuration is preferable.

In conclusion, EmbJXTAChord is suitable for all the cases where the advantages described in Sect.2 make acceptable for the developer the larger overhead caused by compression and encryption. In order to allow the optimization of the transmission strategy used for each link, EmbJXTAChord adds some *feature bits* to the message that signal to the recipient node the compression schemes that are currently in use.

### 5.3.2. jxCOAP-E latency over a homogeneous network (no encryption)

The following tests were aimed to measure the average round-trip latency ( $\mathcal{L}(\eta)$ ) of the jxcoop requests sent to the Fibonacci, Newton or SortSquareRoot service over a homogeneous network. Such a metric, that refers to all clients, can be determined in function of  $\eta$ , averaging all the measured RTT (round-trip




Figure 13: (a-f) Latency values measured for jxCOAP-E on a homogeneous network using 1 rendezvous server and 4 PC edgepeers without and with AES encryption; (g-l) Performance comparison between jxCOAP-E and jxCOAP under SICS and BT networks.

times) of the jxcoap requests received by the server (the average is done overall, i.e. regardless of which of the  $n$  vclients sent the request).

A PC rendezvous *server* was connected to  $s = 4$  edgepeer PCs, each running  $v = 1$  vclient (PC-only configuration). Hence, the server responded to  $n = s * v = 4$  vclients in all. Each test was performed for  $T = 180s$ , using the values in  $S = \{5, 10, 15, 20, 25, 30\}$  for the  $\eta$  parameter. The tests were performed on a network made up of one server and of four vclients linked through Ethernet (ETH), Bluetooth (BT v2.1) or IEEE 802.15.4 (SICS) links. The trials were performed using jxCOAP-E within the standard NetPeerGroup.

The whole experiment was performed using the PC-only configuration and, later, replacing the 4 PCs with 4 RaspPI model B+ boards (RaspPI configuration) and with 4 RaspPI-3 model B boards (RaspPI-3 configuration).<sup>13</sup> The CPU and OS were the same described in Sect. 5.1.1. The trials related to the RaspPI-3 configuration were performed using Bluetooth transmitters compatible with BLE v4.1.

Fig. 13(a-f), Fig. 14(a-f) and Fig. 14(g-l) show the latency values measured respectively for PC-only, RaspPI and RaspPI-3 configurations. The slope changes in the  $\mathcal{L}$  curves are mainly determined by the time required for compression (which varies on  $\eta$ ) and by the compression ratio that, not being constant, affects the payload size and thus the transmission time. The SICS values are

<sup>13</sup>The *coop block-wise transfer mode* [81] was disabled during tests, as RFCOMM, 6LoWPAN and TCP are reliable and connection-oriented protocols. For SICS tests only, the intermessage compression (see Sect. 4.3) was used with the parameters  $c_S = 30$  and  $c_X = 10$ . In this experiment and in the following ones, the Hypercompressor was configured to use all compression schemes (CP8 default configuration).




Figure 14: Latency values measured for jxCOAP-E on a homogeneous network (a-f) using 1 rendezvous server and 4 RaspPI edgepeers without and with AES encryption; (g-l) using 1 rendezvous server and 4 RaspPI-3 edgepeers without and with AES encryption.




Figure 15: (a-c) Performance comparison between jxCOAP-E and UDP CoAP over an Ethernet-based homogeneous network (1 rendezvous server, 4 edgepeers) (d) The latency values determined for UDP CoAP over an Ethernet network.

affected also by the overhead due to the IMX cache manager (see Sect. 4.3), as the CM must calculate the hash of each xmldoc or string msgelem before transmission.

The results show that using ETH, BT or BLE  $\mathcal{L}(\eta) < 1.2s$  for all configurations. Moreover,  $\mathcal{L}(\eta) < 0.5s$  for BLE and RaspPI-3, thus ensuring a responsive behaviour for a wide set of applications (home management and automation, ambient assisted living, health monitoring). The latencies measured for SICS, instead, are significantly higher. This makes the use of EmbJXTACHord over SICS networks suitable only for applications that do not require a reactive behaviour (i.e. smart metering, data gathering, enviromental monitoring). Performance under SICS can be improved tailoring the IMX cache parameters and the compression configuration on the basis of the traffic generated by the specific application.

	Ethernet ( $\beta^{ETH}$ )			Bluetooth ( $\beta^{BT}$ )			SICS ( $\beta^{SICS}$ )			Hybrid ( $\beta^{BT SICS}$ )		
	Fib	Newt	SortSq	Fib	Newt	SortSq	Fib	Newt	SortSq	Fib	Newt	SortSq
PC	19.0%	7.5%	17.6%	1.8%	1.0%	2.7%	23.3%	11.6%	19.1%	17.6%	7.57%	14.7%
RaspPI	22.1%	17.8%	41.3%	9.8%	8.4%	11.5%	24.3%	18.6%	21.9%	18.4%	10.2%	38.5%
RaspPI-3(*)	22.5%	15.0%	17.5%	1.6%	2.2%	2.8%	24.3%	18.1%	20.4%	12.3%	9.4%	32.6%

Table 13: Average  $\mathcal{L}$  overhead values (%) determined by the AES-128 encryption on homogeneous and heterogeneous networks.

### 5.3.3. Comparison between jxCOAP-E and UDP CoAP over an Ethernet-based homogeneous network

These trials compared the latency values for the three jxcoap services to the ones measured for a UDP CoAP server, using Ethernet and PC-only configuration. Fig. 13(a) and Fig. 15(d) show the  $\mathcal{L}_{jxcoapE}$  and  $\mathcal{L}_{udp.coap}$  times for each value of  $\eta$ , while Fig. 15(a-c) show the ratio  $\rho(\eta) = \mathcal{L}_{jxcoapE}(\eta)/\mathcal{L}_{udp.coap}(\eta)$ . Fig. 15(d), which refers to UDP CoAP, indicates that the Newton service requires the highest elaboration time, followed by the SortSquareRoot and the Fibonacci ones. Despite this, the  $\mathcal{L}_{jxcoapE}$  times detected for the Newton service are always lower than the ones detected for the SortSquareRoot service, thus indicating that the jxCOAP-E latency depends more on the compression and delivery times (which are affected in turn by the payload size and by the channel bandwidth) than on the elaboration time. The SortSquareRoot service requires more time for compression than the Newton service, as the payload returned in the response by the first one is a xmldoc that is EXI+GZIP compressed, while the payload returned by the second one is a ByteArray that is only GZIP compressed (see Sect. 4.3).

### 5.3.4. Encryption overhead

In order to measure the overhead determined by AES encryption, the trials described in Sect. 5.3.2 were repeated within an AES-128 secure peer group. Fig. 13(d-f), Fig. 14(d-f) and Fig. 14(j-l) show the  $\mathcal{L}_{jxcoapE.AES}(\eta)$  times measured when AES-128 encryption is enabled. Assuming that  $\mathcal{L}_{jxcoapE.AES}(\eta) = \mathcal{L}_{jxcoapE}(\eta) + \Delta\mathcal{L}_{el}(\eta) + \Delta\mathcal{L}_{trx}(\eta)$ , where  $\Delta\mathcal{L}_{el}(\eta)$  is the elaboration time needed for encryption and  $\Delta\mathcal{L}_{trx}(\eta)$  is the additional time needed for the transmission of the larger payloads due to the use of secure peer groups, the overhead caused by AES on a single trial was defined as

$$\alpha(\eta) = \frac{\mathcal{L}_{jxcoapE.AES}(\eta)}{\mathcal{L}_{jxcoapE}(\eta)} - 1 = \frac{\Delta\mathcal{L}_{el}(\eta) + \Delta\mathcal{L}_{trx}(\eta)}{\mathcal{L}_{jxcoapE}(\eta)} \quad (3)$$

The  $\alpha(\eta)$  values, measured using for  $\eta$  the values in  $S = \{5, 10, 15, 20, 25, 30\}$ , were averaged thus obtaining the  $\beta = avg_{(\eta \in S)} \alpha(\eta)$  values reported in Tab.13.

Tab.13 shows that the overhead values are mostly below 25%. For instance, using RaspPI-3 configuration, the latency values measured for the SortSquareRoot service on the BLE network were  $\mathcal{L}(\eta = 20) = 450.37\text{ms}$  without encryption and  $\mathcal{L}(\eta = 20) = 460.72\text{ms}$  with encryption (overhead of 10.35ms, 2.30%). The low overhead measured for BLE makes this very suitable for applications requiring responsiveness.

The trials on SICS experienced the highest overhead. For instance, using PC configuration, the latency values measured for Newton service on the SICS network was  $\mathcal{L}(\eta = 10) = 1053.24\text{ms}$  without encryption and  $\mathcal{L}(\eta = 10) = 1161.54\text{ms}$  with encryption (overhead of 108.30ms, 10.28%).

Using RaspPI configuration, the latency values measured for the SortSquare-Root service on the SICS network were  $\mathcal{L}(\eta = 20) = 2830.63\text{ms}$  without encryption and  $\mathcal{L}(\eta = 20) = 3503.99\text{ms}$  with encryption (overhead of 673.36ms, 23.79%). The values measured for the Fibonacci service on SICS were  $\mathcal{L}(\eta = 20) = 1468.47\text{ms}$  without encryption and  $\mathcal{L}(\eta = 20) = 1841.96\text{ms}$  with encryption (overhead of 373.49ms, 25.43%).

These overhead values are determined by the additional elaboration time  $\Delta\mathcal{L}_{el}(\eta)$  for AES encryption and by the additional transmission time  $\Delta\mathcal{L}_{trx}(\eta)$ . In fact, the use of a secure peergroup increases the message size as additional data (such as the PeerGroupID) must be transmitted in some MessageElements. Moreover,  $\Delta\mathcal{L}_{trx}(\eta)$  is affected also by CBC padding that can determine a large overhead in a WSN because of the small size of the IEEE 802.15.4 frames [86]. The measured latency values indicate that the use of EmbJXTAChord in a WSN is suitable only with fast processors and for applications that do not require the transmission of large payloads.

### 5.3.5. Comparison between jxCOAP-E and jxCOAP

These trials compared the performance of the jxCOAP-E version integrated in EmbJXTAChord and of the early jxCOAP version based on JXTA 2.7 integrated in jxActinium [46]. Using the PC-only configuration, in order to measure the performance of jxCOAP, all compression schemes were disabled and all latencies were measured again for BT and SICS links. Next, the ratio  $\sigma(\eta) = \mathcal{L}_{jxcoop.jxact}(\eta)/\mathcal{L}_{jxcoopE}(\eta)$  was computed for each value of  $\eta$ . Fig. 13(j-l) show that on the slow SICS links jxCOAP-E largely outperforms jxCOAP with latency improvements between 144% and 280% ( $\sigma^{SICS}(\eta)$  is bounded between 2.44 and 3.80). Moreover Fig. 13(g-i) show that jxCOAP-E outperforms jxCOAP also on the BT links, with latency improvements between 1% and 6% ( $\sigma^{BT}(\eta)$  is bounded between 1.00 and 1.06). In all the cases, the reduction of the transmission time compensates for the overhead introduced by the compression.

### 5.3.6. Latency over a heterogeneous network made up of Bluetooth and 6LoWPAN subnetworks

A PC rendezvous *server* (rdvpeer0) was connected through a BT link to a second PC rendezvous (rdvpeer1) that was in turn connected to  $s = 4$  edgepeer PCs, each running  $v = 1$  vclients, through SICS wireless links (Fig. 16(c)). In this way, rdvpeer1 worked as a *bridge* between the BT and SICS subnetworks. The latencies of the jxcoop requests were measured for the Fibonacci, Newton and SortSquareRoot services, using jxCOAP-E within the standard NetPeerGroup and within an AES-128 secure peergroup. The experiment was performed using the PC, RaspPI and RaspPI-3 configurations. The results show that EmbJXTAChord is able to ensure a seamless communication between the server working in the BT subnetwork and the vclients working in the SICS subnetwork. Fig. 16 shows that the  $\mathcal{L}_{jxcoopE}$  times are in most cases higher than the ones measured for the homogeneous network, because of the overhead due to the network translation performed by the bridge. The  $\beta^{BTSICS}$  overhead values measured using AES-128 encryption are reported in Tab. 13.

### 5.4. The cost of Chord consistent hashing

As described in Sect. 4.2, each rdvpeer of the peergroup implements the Chord DHT algorithm, maintaining in memory its own fingertable.




Figure 16: Latency values measured for a BT-SICS hybrid network (a-b) using 1 rdv server and 4 PC edgepeers without and with AES encryption; (d-e) using 1 rdv server and 4 RaspPI edgepeers without and with AES encryption; (f-g) using 1 rdv server and 4 RaspPI-3 edgepeers without and with AES encryption; (c) the hybrid network architecture used for testing.

Number of nodes ( $N$ )	Messages per rdvpeer ( $\log_2 N$ )	Single msg size (B) (compressed)	Overall msg size (B) (compressed)	Single msg size (B) (uncompressed)	Overall msg size (B) (uncompressed)
8	3	479	1437	1732	5196
64	6	479	2874	1732	10392
128	7	479	3353	1732	12124
2048	11	479	5269	1732	19052
4096	12	479	5748	1732	20784
16384	14	479	6706	1732	24248

Table 14: Overall size (B) of the messages required to each of the  $N$  rdvpeers for updating Chord fingertable when a new rdvpeer joins the group.

In [42] it was demonstrated that joining or leaving a new rdvpeer requires to update the informations of other  $O(\log_2 N)$  rdvpeers, each of which needs in turn, with high probability,  $O(\log_2 N)$  Chord messages to update its own fingertable. Using EmbJXTAChord, each of the messages needed for updating the fingertable (named *find\_successor* requests in Chord terminology) contains a 128b *start field* together with other informations about the sender node.

In our experiments, it was measured that the transmitted size of these messages was 479B (when all compression schemes were enabled) and 1732B (without compression). Tab.14 shows that the overall size of the messages transmitted by each rdvpeer of the group when a new joining/leaving event occurs, assuming that they are in number of  $O(\log_2 N)$ . The table shows that compression allows to maintain very low the transmission cost even with a high number of rdvpeers.

When a new rdvpeer joins the group, it is necessary also to transfer part of the SRDI entries into the cache of the new member. Each SRDI entry requires further 32B (256b) (128b for the hash value of the advertisement and 128b for the PeerID of the storing node). Unfortunately, the overall size of the SRDI update message is not predictable a priori as it depends on the number of SRDI entries (i.e. on the number of stored advertisements).



## 6. A typical scenario for a smart home

In order to show the advantages of using JXTA in the field of IoT, let us consider a scenario consisting in a small residence made up of some three-storey houses. Some services are shared by all the tenants of the residence. Moreover, each house contains some smart objects, sensors and actuators. The smartphone of each member of the family can be used for person authentication at the family's house entrance. While a webcam, connected to a small PC acting as the house Authentication Peer (AP), acquires the face image of a family member, in their pocket the smartphone runs a small application, based on EmbJXTAChord, that transmits to the AP the face biometric template through a BT connection. A robust face authentication algorithm that works using small-sized templates can be found, for instance, in [89]. Once the AP has recognized the identity of the family member, it sends some commands (using jxCOAP-E) to all the smart devices in the house, thus opening the main door and customizing several parameters of the environment (temperature and lightning level in the rooms, kind of music to play etc.).

Sensors and actuators can be connected to one or multiple smart peers, each one made up of a Raspberry PI-3 [36] or of a Raspberry PI-Zero [90] connected via USB to a 6LoWPAN transmitter board. In this way, all the sensor boards can access all the functionalities of the whole peergroup (including the services provided by laptops, by PCs connected via Ethernet or by BT devices such as mobile phones) leveraging on the hop-by-hop delivery and on the routing over subnetworks. No proprietary hardware for bridging ETH/BT/SICS is required, because the single rendezvous peer of the floor acts as a gateway between the Bluetooth or IEEE 802.15.4 subnetworks and the rest of system. Another advantage is that the GPIO port of the Raspberry can be directly connected to digital sensors and actuators [91].

All the smart devices, sensors, actuators and domestic appliances in a house belong to the same HomePeerGroup (HPG), protected using AES-128 encryption. Each smart device in the group can use the REST-ful interface for coordinating its own operations with the activities currently performed by the other devices. For instance, each appliance can read the power available at the moment, using this information to defer the starting time of a job.

If several HPGs exist, one for each house, a device in a house cannot access or see the services or devices installed in another house. However, all devices installed in the residence can be grouped into the Residence Peergroup (RPG), in order to share some of their own functionalities. For example, in each of the houses of the residence there is a smart monitor able to connect to the entryphone in order to see who is at the residence entrance. In order to preserve confidentiality about the visitors of each house, the entryphone and the smart monitor create a secure socket via TLS before starting the video connection. Moreover, any two members of the RPG use their smart monitors to communicate through a secure connection.

In the proposed architecture, a new peer can be added to the system without any user intervention. Once a peer has published an advertisement within the group, all its functionalities are exploitable by all the peers of the same group, even if they are not in the same subnetwork.

## 7. Server and client sample programs

About the scenario described in Sect. 6, a server-side sample program is shown in Listing 1. The rendezvous server (whose JXTA name is *PeerServer*) deploys a custom peergroup named *ChildPeerGroup0*. Next, in the custom peergroup a new jxCOAP-E service named *SJXTA\_CoapService0* is created. The communication is protected using the preshared-key *GROUP\_PASSW* (0000).

```
package sjxta.main;

import sjxta.*;
import sjxta.classes.*;

public class FaceAuthentication extends JxCoapResourceBase
{
    public FaceAuthentication(String name)
    {
        super(name);
    }

    public void handleGET(JxCoapExchange exchange)
    {
        // Here is the code for face authentication and response
        // to the client peer
    }
}

public class Main
{
    private static final String GROUP_PASSW = "0000";

    public static void main(String[] args) throws java.io.IOException
    {
        SJXTA_NetObj MyNetObj = new SJXTA_NetObj();
        SJXTA_NetObj nObj_MyChildPeerGroup = new SJXTA_NetObj();
        SJXTA_AuthInitializer_PSE MyAuthInit = null;

        // Server configured as rendezvous peer

        int ErrCode = SJXTA_NetObj.Init (MyNetObj,
                                         "PeerServer",
                                         SJXTA_CONST.RENDEZVOUS);

        if (ErrCode==0) // No errors. NetPeerGroup is now associated to
        {
            // the network object named MyNetObj

            ErrCode = SJXTA_NetObj.StartTheConnection(MyNetObj, 0);

            if (ErrCode==0) // No errors.
            {
                // Deploy the child peergroup (protected by AES-128 using GROUP_PASSWD)
                MyAuthInit = new SJXTA_AuthInitializer_PSE (GROUP_PASSWD);

                ErrCode = SJXTA_PeerGroup.CreateNewPeerGroup (MyNetObj,
                                                             nObj_MyChildPeerGroup,
                                                             "ChildPeerGroup0",
                                                             MyAuthInit, 0);

                if (ErrCode==0) // New group created successfully. The new group
                {
                    // is now associated to the network object
                    // nObj_MyChildPeerGroup

                    // Deploy a new jxCOAP-E virtual server within the group
                    // (it provides the service named SJXTA_CoapService0)
                    SJXTA_Service MyNewService = new SJXTA_Service();
                    ErrCode = SJXTA_ServiceManager.DeployNewService (nObj_MyChildPeerGroup,
                                                                    MyNewService,
                                                                    "SJXTA_CoapService0",
                                                                    null, 0);

                    if (ErrCode==0)
                    {
                        // Start the new jxCOAP-E virtual server
                    }
                }
            }
        }
    }
}
```

```

        ErrCode = SJXTA.ServiceManager.StartTheService(MyNewService,
                                                    null,
                                                    0);

        if (ErrCode==0)
        {
            // Add a new resource to the jxCOAP-E server
            MyNewService.AddResource (new FaceAuthentication("faceauth"));

            // Wait for the connection of a jxCOAP-E virtual client
            System.out.println ("Hit any key to continue...");
            System.in.read();

            ErrCode = SJXTA.ServiceManager.RevokeService(MyNewService,
                                                    10000,
                                                    0);
        }

        SJXTA.NetObj.StopTheConnection(MyNetObj, 0);
        System.exit(0);
    }
}
}
}
}
}
}
}

```

Listing 1: The source code of the server-side program.

The client-side program is shown in Listing 2. The edge client needs to contact a rdvpeer to start operations. The address of the local rdvpeer can be found through multicast (when ETH or SICS links are used) or providing the address explicitly.

```

package sjxta_main;

import sjxta.*;
import sjxta_classes.*;

public class Main
{
    private static final String GROUP_PASSW = "0000";

    public static void main(String[] args) throws java.io.IOException
    {
        SJXTA.NetObj MyNetObj = new SJXTA.NetObj();
        SJXTA.NetObj nObj.MyChildPeerGroup = new SJXTA.NetObj();
        SJXTA.AuthToken_PSE.String MyAuthToken = null;

        // Client configured as edge peer

        int ErrCode = SJXTA.NetObj.Init (MyNetObj, "PeerClient", SJXTA.CONST.EDGE);

        if (ErrCode==0) // No errors. NetPeerGroup is now associated to the
        { // network object named MyNetObj

            // Set the BT address of the local rdvpeer
            SJXTA.NetConfig.AddSeedRendezvous(MyNetObj, "btspp://#RAPTOR:1");

            // Try to join the NetPeerGroup
            ErrCode = SJXTA.NetObj.StartTheConnection (MyNetObj,
                                                    SJXTA.CONST.NO_RDV_AUTOSTART);

            if (ErrCode==0) // No errors.
            {
                // Create a new authentication token
                MyAuthToken = new SJXTA.AuthToken_PSE.String (GROUP_PASSW);

                // Join to the new secure peergroup using the new token for decryption
                ErrCode = SJXTA.PeerGroup.JoinToPeerGroup (MyNetObj,
                                                            nObj.MyChildPeerGroup,

```



Device	CPU	Link to server	Transm. chipset	jxCOAP-E test	Power cons. (mW) (before test) $\mathcal{P}_b$	Power cons. (mW) (during test) $\mathcal{P}$	Average latency (ms) $\mathcal{L}_{jxcoopE}$	Batt. Voltage (V) $V_{batt}$	Batt. Max Charge (mAh) $C_{batt}$	Estim. life time (Min) $T_{life}^{min}$	Estim. life time (Max) $T_{life}^{max}$
RaspPI	BCM 2835	ETH	built-in	SortSqRoot	1004.0	1161.5	355.96	5.0	10000	43h02m	49h48m
			SendDataBlock	1004.0	1159.2	307.62	5.0	10000	43h07m	49h48m	
		BT v2.1	CSR	SortSqRoot	1064.7	1250.0	1153.2	5.0	10000	40h00h	46h57m
			BT v2.1	SendDataBlock	1064.7	1202.4	984.1	5.0	10000	41h35h	46h57m
SICS	STM MB950	SortSqRoot	1294.8	1482.0	2428.5	5.0	10000	33h44m	38h36m		
		SendDataBlock	1294.8	1432.6	2956.1	5.0	10000	34h54m	38h36m		
RaspPI-3	BCM 2837	ETH	built-in	SortSqRoot	1156.9	1584.0	45.09	5.0	10000	31h33m	43h13m
			SendDataBlock	1156.9	1536.0	38.68	5.0	10000	32h33m	43h13m	
		BT v4.1	built-in	SortSqRoot	1297.4	1831.6	520.4	5.0	10000	27h17m	38h32m
			SendDataBlock	1297.4	1783.4	495.0	5.0	10000	28h02m	38h32m	
		SICS	STM MB950	SortSqRoot	1482.0	2101.5	2414.3	5.0	10000	23h47m	33h44m
				SendDataBlock	1482.0	2152.8	2844.2	5.0	10000	23h13m	33h44m
WiFi	built-in	SortSqRoot	1200.0	1968.0	54.78	5.0	10000	25h24m	41h40m		
		SendDataBlock	1200.0	1920.0	53.93	5.0	10000	26h02m	41h40m		
Oukitel U7 Plus	Mediatek 6737	BT v4.0	built-in	SortSqRoot	773.5	855.0	632.1	4.35	2500	12h43m	14h03m
			SendDataBlock	773.5	810.0	587.9	4.35	2500	13h25m	14h03m	
		WiFi	built-in	SortSqRoot	824.4	914.0	282.1	4.35	2500	11h53m	13h11m
			SendDataBlock	824.4	868.3	333.1	4.35	2500	12h31m	13h11m	

Table 15: Measures about energy consumption and battery lifetime estimation.

fact, EmbJXTACHord can automatically determine the provider node looking for the Module Specification Advertisement of the service (see Sect. 4.4) within the custom peergroup. Once connection is established, EmbJXTACHord transparently provides to compress, encrypt and route over multiple subnetworks (if it is needed) the jxCOAP-E messages.

## 8. Power measures

The last experiment measured the power consumption and the lifetime of a client node (*mobile device*) connected to a jxCOAP-E server, when only a battery pack is available as power source. The test was repeated multiple times, using mobile devices such as a Raspberry PI, a Raspberry PI-3 and one Android smartphone (an Oukitel U7 Plus based on a Mediatek MTK6737 CPU with four 1.3 Ghz cores).

For each trial, the PC and the mobile device were configured as the rendezvous server and the edgepeer client, respectively. Both server and client worked within a custom peergroup, thus performing AES encryption. For each trial described in this section, the mobile phone (edgepeer) ran only a single vclient ( $v = 1$ ). For each of the mobile devices under test, two jxCOAP-E trials were run: *SendDataBlock* and *SortSquareRoot*. The *SendDataBlock* trial was devised to simulate operations that are frequently performed in a smart home. For instance, in the scenario described in Sect.6, a smartphone that transmits to the AP server the biometric template of a user, exploiting the face authentication algorithm described in [89], needs to send requests whose payload is 4KB long, next waiting for a response. During the *SendDataBlock* trial the client sent to the server consecutive requests containing a payload of  $\eta = 4800$  bytes (randomly generated). For each request, the server sent a response whose payload is a 256-chars long string. The *SortSquareRoot* service was already described in Sect. 5.3.1. For this service, the parameter  $\eta = 30$  was always used in all the trials described in this section. Each trial was run for  $T = 180s$ .

A digital wattmeter (connected to the USB port of the mobile device) was used to measure the power consumption during the jxCOAP-E transmission ( $\mathcal{P}$ ). The whole procedure was repeated using different links (ETH, BT, SICS, Wi-Fi) between the server and the mobile device. Before starting each trial, the power consumption  $\mathcal{P}_b$  of the device was measured when the transmitter is powered on, but no elaboration or transmission is performed. A comparison between  $\mathcal{P}$  and  $\mathcal{P}_b$  allows to estimate the power consumption due to message transmission and cpu elaboration.

The measures for Wi-Fi and BT referring to RaspPI-3 were performed using the transmitter modules integrated in the BCM43438 chipset [59]. The measures related to the Oukitel smartphone were realized using the USB port as the only power source (i.e. after disconnecting the phone battery). As the smartphone ran Android 6.0, a dedicated version of EmbJXTAChord based on OpenJDK 9 [92] was developed for the experiment.

The lifetime of the mobile device was calculated assuming the RaspPI and RaspPI-3 connected to a ROMOSS battery [93] (voltage  $V_{batt} = 5$  V, maximum charge  $C_{batt} = 10000$  mAh), and the smartphone connected to the battery provided by the producer (voltage  $V_{batt} = 3.8$  V, maximum charge  $C_{batt} = 2500$  mAh). An estimation of the battery lifetime can be obtained calculating the minimum and maximum value through the formulas  $T_{life}^{min} = (3600 \cdot V_{batt} C_{batt}) / \mathcal{P}$  and  $T_{life}^{max} = (3600 \cdot V_{batt} C_{batt}) / \mathcal{P}_b$ .

Tab. 15 shows that RaspPI-3 ensures lower latencies than RaspPI in all trials, but at the cost of a higher power consumption. Moreover, the use of SICS determines a high power consumption, as the Raspberry boards need to be connected to an external MB950 daughter board. Tab. 15 also shows that all the tested mobile devices, in the conditions previously described, can remain fully operative for many hours without the need for recharging the battery.

## 9. Conclusions

EmbJXTAChord enables to support IoT applications allowing any sensor or actuator device to be connected in secure *peergroups of objects*, regardless of both the transport protocol (TCP, HTTP, Bluetooth RFCOMM, 6LoWPAN) and the features of the link that is actually used for communication. The P2P protocol provides functionalities such as node and resource discovery, secure peergroups, routing over subnetworks, reliable and unreliable connections between nodes. In order to ensure good performance even on narrowband networks, EmbJXTAChord exploits a compression protocol that reduces the size of a JXTA message to 25% of the original one. The protocol is light enough to run on RaspberryPI and RaspberryPI-3 boards, thus allowing the implementation of low-cost heterogeneous IoT networks based on a RESTful architecture.

jxCOAP-E allows the creation of a RESTful architecture over the heterogeneous network. The support for AES-128 protected peergroups allows to create group of peers that shares services and resources, regardless of the underlying network topology. jxCOAP-E leverages on the underlying P2P architecture in order to provide a distributed and fault-tolerant service discovery. In all the trials performed over Ethernet, Bluetooth and Bluetooth Smart networks, jxCOAP-E provided latency values that are low enough to allow the use for applications that require a responsive behaviour (such as home management and automation, assisted living or health monitoring). Conversely, the latency

values measured over a 6LoWPAN wireless sensor network were higher, but they are still acceptable for tasks that do not require a high level of responsiveness (such as smart metering and environmental control).

Future work will deal with the implementation of new Message Transport Binding modules for other protocols, such as G3-PRIME for smart grid [94], and a plugin for the browser to allow monitoring the JXTA peer group resource through a REST-ful API.

## References

### References

- [1] E. Borgia, The internet of things vision: Key features, applications and open issues, *Elsevier Computer Communications* 54 (2014) 1–31. doi:<http://dx.doi.org/10.1016/j.comcom.2014.09.008>.
- [2] J. Mineraud, O. Mazhelis, X. Su, S. Tarkoma, A gap analysis of internet-of-things platforms, *Elsevier Computer Communications* 89–90 (2016) 5–16, *Internet of Things Research challenges and Solutions*. doi:<http://dx.doi.org/10.1016/j.comcom.2016.03.015>.
- [3] A. N. Lee, J. L. M. Lastra, Data aggregation at field device level for industrial ambient monitoring using Web Services, in: *9th IEEE Int. Conf. on Ind. Informatics (INDIN)*, 2011, 2011, pp. 491–496. doi:[10.1109/INDIN.2011.6034929](https://doi.org/10.1109/INDIN.2011.6034929).
- [4] L. Lo Bello, O. Mirabella, N. Torrisi, Modelling and evaluating traceability systems in food manufacturing chains, *Proc. of Int. Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises (2004)* 173–179doi:[10.1109/ENABL.2004.44](https://doi.org/10.1109/ENABL.2004.44).
- [5] L. A. Amaral, R. T. Tiburski, et al., Cooperative Middleware Platform As a Service for Internet of Things Applications, in: *Proc. of the 30th Annual ACM Symposium on Applied Computing, SAC '15*, ACM, New York, NY, USA, 2015, pp. 488–493. doi:[10.1145/2695664.2695799](https://doi.org/10.1145/2695664.2695799).
- [6] X. Feng, J. Shen, Y. Fan, REST: An alternative to RPC for Web services architecture, *Proceedings of First International Conference on Future Information Networks (2009)* 7–10doi:[10.1109/ICFIN.2009.5339611](https://doi.org/10.1109/ICFIN.2009.5339611).
- [7] IEEE 802.15.4 Documents, available online: <https://mentor.ieee.org/802.15/documents>.
- [8] E. Toscano, L. Lo Bello, A topology management protocol with bounded delay for Wireless Sensor Networks, in: *2008 IEEE Int. Conf. on Emerging Tech. and Factory Automation*, 2008, pp. 942–951. doi:[10.1109/ETFA.2008.4638508](https://doi.org/10.1109/ETFA.2008.4638508).
- [9] E. Toscano, L. Lo Bello, Comparative assessments of IEEE 802.15.4/ZigBee and 6LoWPAN for low-power industrial WSNs in realistic scenarios, in: *2012 9th IEEE Int. Workshop on Factory Communication Systems*, 2012, pp. 115–124. doi:[10.1109/WFCS.2012.6242553](https://doi.org/10.1109/WFCS.2012.6242553).

- [10] C. Bormann, A. P. Castellani, Z. Shelby, CoAP: An application protocol for billions of tiny internet nodes, *IEEE Internet Computing* 16 (2) (2012) 62–67. doi:10.1109/MIC.2012.29.
- [11] Z. Shelby, C. Bormann, et al., The Constrained Application Protocol (CoAP), RFC 7252 (June 2014).
- [12] D. Thangavel, X. Ma, A. Valera, H. X. Tan, C. K. Y. Tan, Performance evaluation of MQTT and CoAP via a common middleware, in: 2014 IEEE Ninth International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), 2014, pp. 1–6. doi:10.1109/ISSNIP.2014.6827678.
- [13] P. Nawrocki, M. Jakubowski, T. Godzik, Notification methods in wireless systems, *Computer Science* 17 (4) (2017) 519.
- [14] Real Time Logic, SMQ Protocol Specifications, Available online: <https://realtimelogic.com/downloads/docs/SMQ-specification.pdf>.
- [15] A. Al-Fuqaha, M. Guizani, M. Mohammadi, M. Aledhari, M. Ayyash, Internet of things: A survey on enabling technologies, protocols, and applications, *IEEE Communications Surveys Tutorials* 17 (4) (2015) 2347–2376. doi:10.1109/COMST.2015.2444095.
- [16] D. H. Mun, M. L. Dinh, Y. W. Kwon, An assessment of internet of things protocols for resource-constrained applications, in: 2016 IEEE 40th Annual Computer Software and Applications Conference (COMPSAC), Vol. 1, 2016, pp. 555–560. doi:10.1109/COMPSAC.2016.51.
- [17] L. D. Xu, W. He, S. Li, Internet of Things in Industries: A Survey, *IEEE Transactions on Industrial Informatics* 10 (4). doi:10.1109/TII.2014.2300753.
- [18] Y. Feng, Q. Li, The distributed UDDI system model based on service oriented architecture, in: 2016 7th IEEE International Conference on Software Engineering and Service Science (ICSESS), 2016, pp. 585–589. doi:10.1109/ICSESS.2016.7883138.
- [19] A. Furno, E. Zimeo, Self-scaling cooperative discovery of service compositions in unstructured p2p networks, *Journal of Parallel and Distributed Computing* 74 (10) (2014) 2994–3025. doi:http://dx.doi.org/10.1016/j.jpdc.2014.06.006.
- [20] M. Kaouan, D. Bouchiha, S. M. Benslimane, Shared-repository based approach for storing and discovering web services, *Procedia Computer Science* 73 (2015) 56–65, international Conference on Advanced Wireless Information and Communication Technologies (AWICT 2015). doi:http://dx.doi.org/10.1016/j.procs.2015.12.049.
- [21] B. Yuan, L. Liu, N. Antonopoulos, A self-organized architecture for efficient service discovery in future peer-to-peer online social networks, in: 2016 IEEE Symposium on Service-Oriented System Engineering (SOSE), 2016, pp. 415–422. doi:10.1109/SOSE.2016.57.



- [22] J. Granjal, E. Monteiro, et al., Security for the Internet of Things: A Survey of Existing Protocols and Open Research Issues, *IEEE Communications Surveys and Tutorials* doi:10.1109/COMST.2015.2388550.
- [23] S. Raza, P. Misra, Z. He, T. Voigt, Building the Internet of Things with Bluetooth Smart, *Ad Hoc Networks* 57 (2017) 19–31, special Issue on Internet of Things and Smart Cities: security, privacy and new technologies. doi:http://dx.doi.org/10.1016/j.adhoc.2016.08.012.
- [24] Microsoft, Bluetooth Version and Profile Support in Previous Windows Versions, available online: <https://msdn.microsoft.com/en-us/library/windows/hardware/mt734149%28v=vs.85%29.aspx> (2017).
- [25] S. Kent, K. Seo, RFC 4301 - Security Architecture for the Internet Protocol, available online: <https://www.rfc-editor.org/rfc/rfc4301.txt> (Dec 2005).
- [26] T. Dierks, C. Allen, RFC 2246 - The TLS Protocol Version 1.0, available online: <https://www.ietf.org/rfc/rfc2246.txt> (Jan 1999).
- [27] E. Rescorla, N. Modadugu, RFC 6347 - Datagram Transport Layer Security Version 1.2, available online: <https://tools.ietf.org/html/rfc6347> (Jan 2012).
- [28] Y. Wang, T. T. Gamage, C. H. Hauser, Security implications of transport layer protocols in power grid synchrophasor data communication, *IEEE Transactions on Smart Grid* 7 (2) (2016) 807–816. doi:10.1109/TSG.2015.2499766.
- [29] S. Raza, S. Duquennoy, J. Höglund, U. Roedig, T. Voigt, Secure communication for the Internet of Things: a comparison of link-layer security and IPsec for 6LoWPAN, *Security and Communication Networks* 7 (12) (2014) 2654–2668. doi:10.1002/sec.406.
- [30] A. Rahman, E. Dijk, et al., RFC7390, Group Communication for the Constrained Application Protocol, available online: <https://www.rfc-editor.org/rfc/rfc7390.txt> (Oct 2014).
- [31] Isam Ishaq, Jeroen Hoebeke, et al., Experimental Evaluation of Unicast and Multicast CoAP Group Communication (July 2016).
- [32] I. Ishaq, J. Hoebeke, I. Moerman, P. Demeester, Observing CoAP groups efficiently, *Ad Hoc Networks* 37 (2016) 368–388. doi:http://dx.doi.org/10.1016/j.adhoc.2015.08.030.
- [33] S. H. Shaheen, M. Yousaf, Security Analysis of DTLS Structure and its Application to Secure Multicast Communication, in: *12th Int. Conf. on Frontiers of Information Technology*, 2014, pp. 165–169. doi:10.1109/FIT.2014.39.
- [34] Keoh, S. L., Kumar, S. S., Tschofenig, H., Securing the Internet of Things: A Standardization Perspective (June 2014). doi:10.1109/JIOT.2014.2323395.

- [35] V. Vujovic, M. Maksimovic, Raspberry Pi as a Wireless Sensor node: Performances and constraints, Proc. of 37th International Convention on Information and Communication Technology, Electronics and Microelectronics (2014) 1013–1018doi:10.1109/MIPRO.2014.6859717.
- [36] Raspberry foundation, RaspberryPI 3 Model B, available online: <https://www.raspberrypi.org/products/raspberry-pi-3-model-b/> (2016).
- [37] J. W. Lee, H. Schulzrinne, W. Kellerer, et al., z2z: Discovering Zeroconf Services Beyond Local Link, in: 2007 IEEE Globecom Workshops, 2007, pp. 1–7. doi:10.1109/GLOCOMW.2007.4437805.
- [38] J. Verstrynge, JXSE v2.7 The JXTA Java Standard Edition Implementation Programmer’s Guide, available online: <https://jxse.kenai.com/Tutorials/> (March 2011).
- [39] B. Traversat, M. Abdelaziz, E. Pouyoul, Project JXTA: A loosely-consistent DHT rendezvous walker, Sun Microsystem (2003).
- [40] JXTA 2.0 Protocol Specifications, available online: <https://jxta.kenai.com/Specifications/JXTAProtocols2.0.pdf>.
- [41] A. Ghosh, T. Givargis, Source routing made practical in embedded networks, Proceedings of 18th International Conference on Computer Communications and Networks (2009) 1–6doi:10.1109/ICCCN.2009.5235356.
- [42] I. Stoica, R. Morris, D. Kanger, et al., Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications, IEEE ACM Transactions on Networking 11 (1) (2003) 149–160. doi:10.1145/383059.383071.
- [43] M. Kirsche, R. Klauck, Unify to bridge gaps: Bringing XMPP into the Internet of Things, Proc. of IEEE Int. Conf. on Pervasive Computing and Communications Workshops (2012) 455–458doi:10.1109/PerComW.2012.6197534.
- [44] X. Cheng, G. Dang, The P2P communication technology research based on Internet of things, IEEE Workshop on Adv. Research and Technology in Industry Applicationsdoi:10.1109/WARTIA.2014.6976225.
- [45] M. Domingo-Prieto, J. Arnedo-Moreno, X. Vilajosana-Guillen, jxSensor: a sensor network integration layer for JXTA, Proc. of 15th Int. Conf. on Network Based Information Systemsdoi:10.1109/NBiS.2012.125.
- [46] F. Battaglia, G. Iannizzotto, L. L. Bello, JxActinium: A Runtime Manager for Secure REST-ful CoAP Applications Working over JXTA, in: Proc. of ACM Symposium on Applied Computing, SAC ’16, ACM, 2016. doi:10.1145/2851613.2851808.
- [47] P. Evensen, H. Meling, SenseWrap: A service oriented middleware with sensor virtualization and self-configuration, in: 2009 5th Int. Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), 2009, pp. 261–266. doi:10.1109/ISSNIP.2009.5416827.

- [48] J. E. Kim, G. Boulos, J. Yackovich, et al., Seamless Integration of Heterogeneous Devices and Access Control in Smart Homes, in: 2012 8th International Conference on Intelligent Environments (IE), 2012, pp. 206–213. doi:10.1109/IE.2012.57.
- [49] I. Mashal, O. Alsaryrah, T.-Y. Chung, C.-Z. Yang, W.-H. Kuo, D. P. Agrawal, Choices for interaction with things on Internet and underlying issues, *Ad Hoc Networks* 28 (2015) 68–90. doi:http://dx.doi.org/10.1016/j.adhoc.2014.12.006.
- [50] Apache ActiveMQ, available online: <http://activemq.apache.org/>.
- [51] Node-RED , available online: <https://nodered.org/>.
- [52] M. Blackstock, R. Lea, Toward a Distributed Data Flow Platform for the Web of Things (Distributed Node-RED), in: Proceedings of the 5th International Workshop on Web of Things, WoT '14, ACM, New York, NY, USA, 2014, pp. 34–39. doi:10.1145/2684432.2684439.
- [53] M. Blackstock, R. Lea, IoT mashups with the WoTKit, in: 2012 3rd IEEE International Conference on the Internet of Things, 2012, pp. 159–166. doi:10.1109/IOT.2012.6402318.
- [54] M. Blackstock, R. Lea, FRED: A Hosted Data Flow Platform for the IoT, in: Proceedings of the 1st International Workshop on Mashups of Things and APIs, MOTA '16, ACM, New York, NY, USA, 2016, pp. 2:1–2:5. doi:10.1145/3007203.3007214.
- [55] EVERYTHING, available online: <https://evrythng.com/>.
- [56] Evrythng THNGHUB, available online: <https://developers.evrythng.com/docs/thnghub-dal-specification>.
- [57] Bluetooth SIG, Bluetooth Specification: HTTP Proxy Service, available online: <https://www.bluetooth.com/specifications/adopted-specifications>.
- [58] Bluetooth SIG, Bluetooth Specification: Internet Protocol Support Profile, available online: <https://www.bluetooth.com/specifications/adopted-specifications>.
- [59] Cypress, CYW43438 Single-Chip IEEE 802.11ac b/g/n MAC/Baseband/Radio with Integrated Bluetooth 4.1 datasheet, available online: <http://www.cypress.com/documentation/datasheets/cyw43438-single-chip-ieee-80211ac-bgn-macbasebandradio-integrated-bluetooth>.
- [60] K. Abboud, H. A. Omar, W. Zhuang, Interworking of DSRC and Cellular Network Technologies for V2X Communications: A Survey, *IEEE Transactions on Vehicular Technology* 65 (12) (2016) 9457–9470.
- [61] C. Campolo, A. Molinaro, R. Scopigno, From today's VANETs to tomorrow's planning and the bets for the day after, *Vehicular Communications* 2 (3) (2015) 158–171. doi:http://dx.doi.org/10.1016/j.vehcom.2015.06.002.

- [62] X. Lin, J. G. Andrews, A. Ghosh, R. Ratasuk, An overview of 3GPP device-to-device proximity services, *IEEE Communications Magazine* 52 (4) (2014) 40–48. doi:10.1109/MCOM.2014.6807945.
- [63] 3GPP, Technical Specification Group SA. Feasibility Study for Proximity Service (ProSe) (Release 12) TR22.803 v12.2.0 (June 2013).
- [64] C. Hoymann, D. Astely, M. Stattin, G. Wikstrom, J. F. Cheng, A. Hoglund, M. Frenne, R. Blasco, J. Huschke, F. Gunnarsson, LTE release 14 outlook, *IEEE Communications Magazine* 54 (6) (2016) 44–49. doi:10.1109/MCOM.2016.7497765.
- [65] BlueCove JSR-82 implementation, available online: <http://bluecove.org/>.
- [66] JSR-82 Java API for Bluetooth Specifications 1.0 Final release, available online: <http://download.oracle.com/otndocs/jcp/7851-bluetooth-1.0-fr-spec-oth-JSpec/>.
- [67] L. C. Fernandes, J. R. Souza, et al., Carina intelligent robotic car: Architectural design and applications, *Journal of Systems Architecture* 60 (4) (2014) 372–392. doi:<https://doi.org/10.1016/j.sysarc.2013.12.003>.
- [68] BlueZ: Bluetooth stack for Linux kernel, available online: <http://www.bluez.org/>.
- [69] M. T. Zia, M. U. Farooq, et al., Seamless Communication over Heterogeneous Interfaces in Mobile Ad hoc Networks, *IFIP Int. Conf. on Wireless and Optical Comm. Networks* doi:10.1109/WOCN.2009.5010564.
- [70] A. Sikora, P. Digeser, et al., Model based development of a TinyOS-based Wireless M-Bus implementation, in: 2012 IEEE 1st Int. Symp. on Wireless Systems (IDAACS-SWS), 2012, pp. 91–94. doi:10.1109/IDAACS-SWS.2012.6377640.
- [71] O. Hahm, E. Baccelli, H. Petersen, N. Tsiftes, Operating systems for low-end devices in the internet of things: A survey, *IEEE Internet of Things Journal* 3 (5) (2016) 720–734. doi:10.1109/JIOT.2015.2505901.
- [72] M. Zhao, A. Kumar, et al., A comprehensive study of RPL and P2P-RPL routing protocols: Implementation, challenges and opportunities, *Peer-to-Peer Networking and Applications* (2016) 1–25 doi:10.1007/s12083-016-0475-y.
- [73] H. Fotouhi, D. Moreira, et al., mRPL+: A mobility management framework in RPL/6LoWPAN, *Elsevier Computer Communications* (2017) – doi:<http://dx.doi.org/10.1016/j.comcom.2017.01.020>.
- [74] A. Ludovici, A. Calveras, J. Casademont, Forwarding Techniques for IP Fragmented Packets in a Real 6LoWPAN Network, *Sensors* 11 (1) (2011) 992–1008. doi:10.3390/s110100992.
- [75] C. Nocentini, P. Crescenzi, L. Lanzi, Performance evaluation of a Chord based JXTA implementation, *Proceedings of First Int. Conference on Advances in P2P Systems* (2009) 7–12 doi:10.1109/AP2PS.2009.9.

- [76] N. Freed, N. Borenstein, et al., RFC 2045 - Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies (1996).
- [77] D. Peintner, H. Kosch, J. Heuer, Efficient XML Interchange for rich internet applications, Proc. of IEEE Int. Conf. on Multimedia and Expo doi:10.1109/ICME.2009.5202458.
- [78] G. Jaiswal, M. Mishra, Why use Efficient XML Interchange instead of Fast Infoset, IEEE 3rd International Conference on Advance Computing (2013) 925–930 doi:10.1109/IAAdCC.2013.6514350.
- [79] J. Arnedo-Moreno, et al., A Security Layer for JXTA Core Protocols, in: 2009 Int. Conf. on Complex, Intelligent and Software Intensive Systems, 2009, pp. 463–468. doi:10.1109/CISIS.2009.24.
- [80] M. Kovatsch, M. Lanter, Z. Shelby, Californium: Scalable cloud services for the Internet of Things with CoAP, Proc. of the 2014 Int. Conf. on the Internet of Things (2014) 1–6 doi:10.1109/IOT.2014.7030106.
- [81] C. Bormann, Z. Shelby, Blockwise transfers in CoAP, draft-ietf-core-block-20 (2013).
- [82] A. Baig, IPv6 campus network deployment guidelines for DNS, Web server, Proxy server and Wi-Fi, in: 2016 26th International Telecommunication Networks and Applications Conference (ITNAC), 2016, pp. 237–242.
- [83] P. K. Kamma, C. R. Palla, U. R. Nelakuditi, R. S. Yarrabothu, Design and implementation of 6LoWPAN border router, in: 2016 Thirteenth International Conference on Wireless and Optical Communications Networks (WOCN), 2016, pp. 1–5. doi:10.1109/WOCN.2016.7759025.
- [84] S. Zander, L. L. Andrew, G. Armitage, G. Huston, G. Michaelson, Investigating the IPv6 Teredo Tunnelling Capability and Performance of Internet Clients, SIGCOMM Computer Communication Review 42 (5) (2012) 13–20. doi:10.1145/2378956.2378959.
- [85] G. Antoniu, P. Hatcher, M. Jan, et al., Performance Evaluation of JXTA Communication Layers, Workshop on Global and Peer-to-Peer Computing (GP2PC'2005) doi:10.1109/CCGRID.2005.1558562.
- [86] B. Sun, C.-C. Li, K. Wu, Y. Xiao, A lightweight secure protocol for wireless sensor networks, Computer Communications 29 (13) (2006) 2556–2568, wireless Sensor Networks and Wired/Wireless Internet Communications. doi:https://doi.org/10.1016/j.comcom.2006.02.006.
- [87] STM32W RF Control Kit Specifications, available online: <http://www.st.com>.
- [88] C. Pham, Communication Performances of IEEE 802.15.4 Wireless Sensor Motes for Data-intensive Applications, J. Netw. Comput. Appl. 46 (C) (2014) 48–59. doi:10.1109/WD.2013.6686516.

- [89] F. Battaglia, G. Iannizzotto, L. L. Bello, A person authentication system based on RFID tags and a cascade of face recognition algorithms, *IEEE Trans. on Circuits and Systems for Video Technology* PP (99) (2016) 1–1. doi:10.1109/TCSVT.2016.2527299.
- [90] Raspberry foundation, RaspberryPI Zero, available online: <https://www.raspberrypi.org/products/pi-zero/> (2016).
- [91] M. Ibrahim, A. Elgamri, S. Babiker, et al., Internet of things based smart environmental monitoring using the Raspberry-Pi computer, in: 2015 Fifth International Conference on Digital Information Processing and Communications (ICDIPC), 2015, pp. 159–164. doi:10.1109/ICDIPC.2015.7323023.
- [92] Oracle, OpenJDK for Android, available online: <http://openjdk.java.net/projects/mobile/android.html>.
- [93] ROMOSS Solo 5 Power Bank, available online: <https://www.romoss.com/products/power-banks/romoss-solo-5-10000mah-power-bank/>.
- [94] G. Patti, G. Alderisi, L. L. Bello, Performance assessment of the PRIME MAC layer protocol, *IEEE 11th Int. Conference on Industrial Informatics (INDIN)* (2013) 158–164doi:10.1109/INDIN.2013.6622875.