Distributed Crowd-Sensing Infrastructure for Personalized Dynamic IoT Spaces

Peter Rothenpieler, Bashar Altakrouri, Oliver Kleine, Lukas Ruge
Institute of Telematics, University of Lübeck
Ratzeburger Allee 160, Lübeck, Germany
{rothenpieler|altakrouri|kleine|ruge}@itm.uni-luebeck.de

ABSTRACT

In this paper, we describe a distributed crowd-sensing infrastructure that integrates and bridges small scale personalized ad-hoc Internet of Things (IoT) spaces (consisting of personal interconnected smart devices, sensors and actuators dynamically deployed at runtime) to large scale IoT spaces. While a lot of innovation takes place on large scale IoT infrastructures, we focus on a personalized IoT infrastructure that allows user level control and management of personally owned IoT resources. Our approach uses a peer to peer (P2P) network together with distributed discovery- and directory-services, without the need for centralized infrastructure. The contribution of this paper is twofold: Firstly, we present an Android-based solution called Ambient Bridge that exposes a user-selected subset of the build-in sensors and actuators of a smart device as CoAP (Constrained Application Protocol) web services. Moreover, it is used to dynamically integrate external sensors and actuators at runtime that are normally only accessible via proprietary or non-networked interfaces. Secondly, we present a directory service and distributed semantic search engine called the Smart Service Proxy (SSP). The SSP allows application developers to search for sensors and actuators using SPARQL queries, which are automatically distributed between and processed by the cooperating SSPs.

Categories and Subject Descriptors

C.2 [**Computer-Communication Networks**]: Distributed Systems; D.2.2.c [**Software Engineering**]: Design Tools and Techniques—*Distributed/Internet based software engineering tools and techniques*

General Terms

Distributed Crowd-Sensing

Keywords

CoAP, Internet of Things, Distributed Fusion, Smart Service Proxy, Ambient Dynamix

Urb-IoT 2014, October 27-28, Rome, Italy

DOI 10.4108/icst.urb-iot.2014.257310

1. INTRODUCTION

Advancements in IoT technologies, systems, and infrastructures have led to an unprecedented expanding hybrid world of interconnected smart objects and devices to realize new context-aware services in a seamlessly integrated physical and virtual world. Personal devices and things owned by users have become a prime axis of attention to the IoT community due to their flourishing richness of computation and sensing capabilities. Hence, the community strives to utilize personal devices and things in many application domains such as crowd context sensing in urban cities (e.g., Smart Santander [5]). Moreover, there has been enormous effort towards standard communication protocols for IoT in Wireless Sensor Networks (WSNs) such as 6LoWPAN [6] and the Constrained Application Protocol (CoAP) [7].

Nonetheless, the richness of user devices and things come with a set of open issues and challenges [3] specially due to the heterogeneity and distributivity (i.e., a variety of devices with various capabilities), dynamic media mobility (i.e., devices may join or leave the ambient space at any time), and user mobility (i.e., attending to the user's needs). While large scale IoT infrastructure solutions focus on holistic service delivery and aggregating user generated context, we believe that more effort should be focused on personalized IoT where a better consideration of user IoT resources is possible and better user control are granted.

Our approach uses a peer to peer (P2P) network together with distributed discovery- and directory-services, without the need for a centralized server. Our infrastructure utilizes CoAP [7] to exchange information and to offer services to external users and applications. To allow backwards compatibility, our infrastructure supports transparent protocol conversion from IPv6 to IPv4 and from CoAP to HTTP (and vice versa). In our paper, we substantiate the following main contributions:

- The Ambient Bridge is an Android-based solution and exposes a user-selected subset of the build-in sensors and actuators of a smartphone as CoAP web services. This bridge can also be used to dynamically integrate external sensors and actuators at runtime that are normally only accessible via proprietary or non-networked interfaces. The bridge is built on the Dynamix context modeling framework which allows an in-situ dynamic and ad-hoc integration of context resources on mobile devices (presented in section 2.1).

- The Smart Service Proxy (SSP) which allows application developers to semantically search for sensors and actuators using SPARQL queries which are automatically distributed between and processed by the cooperating SSPs (presented in section 2.2).

2. ARCHITECTURE

This section describes the different components of our IoT and crowd-sensing infrastructure. We begin with a description of the Ambient Bridge in section 2.1 that allows the flexible integration of the sensors and actuators of smartphones as CoAP Web services at runtime. Section 2.2 describes the SSP which acts as a semantic search engine and discovery service and allows the federation of multiple IoT spaces.

2.1 Ambient Bridge

Smartphones contain a large set of sensors which may generate relevant context about the users or their environment. Phones can also use a vast array of wireless technologies to connect to further devices in their vicinity to generate even more contextual data. Applications can use such data to adapt their behavior based on environmental data or on personalization. Phones and connected peripheral devices also act as actuators for applications, to interact with the user in a contextualized manner. As a component of the crowd sensing platform, we developed an Android application called Ambient Bridge which makes the contextual data from the phone and connected devices available as CoAP Web services.

The Ambient Bridge is build atop Ambient Dynamix[1] [1], an Android middleware framework that allows OSGi bundles to be downloaded and installed during runtime on the behest of an application (app) installed on the mobile device. These bundles, called Dynamix plug-ins in the parlance of the framework, can be used to generate high level contextual information from sensors embedded on the phone or from connected resources. Dynamix plug-ins can be written by domain experts and supplied to a repository from which they may be fetched on demand. The contextual information generated by the plug-ins can be accessed by the apps in several formats, e.g. as Java-Objects or in a textual representation. Using Ambient Dynamix, app-developers can access ready-made contextual information, designed by experts, instead of implementing algorithms to generate such information. Aside from generating contextual information, Ambient Dynamix can also use plug-ins to provide applications on the smartphone access to actuators, either on the phone or peripheral devices.

The Ambient Bridge allows access to the context and actuators provided via Dynamix with a CoAP and a HTTP server. This allows applications on arbitrary devices to use the generated context provided by Dynamix. Using the DNS Service Discovery [2] implementation available in Android (4.0 and higher), the Ambient Bridge registers in a local network to be discoverable. After being discovered, any system may use the Well-Known-Resource of the CoAP-Server to discover the available sensors and actuators. Using GET-Requests any device in the local network can request an overview of available context types generated by the smartphone, as well as currently not available context-types which may be made available by installing Dynamix plug-ins. Using POST-Requests, users can request the installation of a context plug-in or request access to a context type.

The Ambient Bridge may further register itself at the Smart Service Proxy configured by the user, which allows its integration into the large-scale IoT space. This allows the use of their sensors and actuators of the smartphone outside the local network in a larger crowd-sensing context which is described in more detail in the following section. To protect the privacy of the user, the Ambient Bridge does not provide any context type without the users' authorization. Given such an authorization, a new CoAP-Resource is generated and the context is made available. Users can revoke such access at any time.

Using contextual data generated by a mobile device, such as a users' personal smartphone, provides several advantages over the singular use of sensors and actuators already provided in the smart space. The user may provide additional data that was previously not available and thus provides a broader set of data to any application. Contextual data generated by the user may also be personalized, due to the fact that the user can manually configure many Dynamix plug-ins to personalize the provided data and that context plug-ins can include learning algorithms that adapt the provided context over time.

The list of plug-ins available for Ambient Dynamix is already large and continually growing, providing information on noise- and light-level, location, recognition of music playing in an environment or speech recognition, step counting, a devices battery level, access to the NFC capabilities of the phone, the ability to control room lighting, Heart Rate or Air Quality data and much more[2].

2.2 Smart Service Proxy

The Smart Service Proxy (SSP) is the middle box to connect the domain of sensors and actuators with the Internet. Many sensors are attached to platforms (such as smartphones) that suffer from several resource constraints like available limited energy, memory, bandwidth, connectivity, or computational power. In the following, we will refer to sensors or actuators that are attached to resource constrained platforms as sensors. However, due to those constraints the direct access from the Internet to such sensors does not scale and thus we aim to avoid this whenever possible. That particularly holds for the retrieval of actual measurements as sensors usually do not measure permanently but, e.g., only once a second to save energy.

Another issue with direct access is that sensors often only talk proprietary protocols. Thus, an interested party would have to be able to talk a number of protocols depending on the sensors that produce the data of interest. The SSP unifies the access to sensors and actuators by means of a RESTful HTTP interface. The common HTTP methods GET, POST, PUT, and DELETE provide the interface to e.g. retrieve a sensors measurement or change the status of

[1] http://dynamix.io

[2] http://dynamix.io/category/plugins

an actuator. The SSP takes care of all necessary protocol conversion.

2.2.1 Permanent Availability of Sensor Data

To make sensor data available permanently, despite the sensors sleep intervals, the SSP observes the sensor values and provides them to the Internet on behalf of the sensor. In terms of sensors or actuators running a CoAP server, the procedure is as follows. The sensors register their CoAP Web services at the SSP by sending a registration request to `coap://<ssp-host>/here_i_am` . Upon registration, the SSP starts the observation of the registered Web services using the CoAP observe protocol extension [4]. Due to the observation, the SSP is always aware of the sensors current status and is thus able to answer requests on behalf of the sensor without any need to stress the (possibly resource constraint) sensor platform.

However, since the IoT may potentially consist of billions of sensors and actuators, we need to consider two things: For starters, we want to be able to find relevant pieces of data despite the vast amount of available data. To do so, we make the sensors offer their data in the RDF [8] interchange format. This allows clients to lookup the data of interest not only by sensor name (IP or DNS name) but also by user-defined criteria, such as all temperature sensors that are located in a certain spatial area. The user defines such criteria in the form of a SPARQL query. The second thing to consider is the scalability, which we describe in the following subsection.

2.2.2 Relevant Data Lookup

Our scalability concept supports the use of an arbitrary amount of SSPs, which share a common distributed database, i.e., a distributed RDF triple store that is organized in a P2P overlay network. RDF data is structured in `[<subject> <predicate> <object>]` triples and SPARQL queries are eventually broken down to a number of basic query patterns. A basic query pattern consists of 0, 1, 2, or 3 fixed elements of the triple(s) to be found filled up with variables to form a full triple. E.g., the pattern `[?s <predicate> <object>]` finds all triples having the given predicate and object whereas `[<subject> ?p ?o]` refers to all triples having the given subject.

The storage location of a piece of data within the P2P network is determined by the hash value of the data. With $h(x)$ as the hash function and & as the concatenation operator, the resulting hash values per triple are: $h(subject)$, $h(predicate)$, $h(object)$, $h(subject\&predicate)$, ... $h(subject\&predicate\&object)$. By this means we are able to find triples using any combination of known and unknown triple elements. We would like to note that the use of the `[?s ?p ?o]` query pattern is useless, as this returns all triples.

Thus, the SSP (resp. the network of SSPs) can also be considered a directory service which enables the retrieval of particular data, i.e. triples, using SPARQL queries.

3. CONCLUSION AND FUTURE WORK

In this paper, we have presented a distributed crowd-sensing IoT infrastructure that consists primarily of two components called the Ambient Bridge and the Smart Service Proxy. Firstly, the Ambient Bridge exposes a user-selected subset of the build-in sensors and actuators of a smart device as CoAP web services. Secondly, the Smart Service Proxy acts as a directory service and semantic search engine for sensors and actuators.

4. REFERENCES

[1] D. Carlson and A. Schrader. Dynamix: An open plug-and-play context framework for android. In *Proceedings of the 3rd International Conference on the Internet of Things (IoT2012)*, Wuxi, China, Oct. 2012.

[2] S. Cheshire and M. Krochmal. DNS-Based Service Discovery. RFC 6763, Feb. 2013.

[3] R. K. Ganti, F. Ye, and H. Lei. Mobile crowdsensing: current state and future challenges. *Communications Magazine, IEEE*, 49(11):32–39, 2011.

[4] K. Hartke. Observing Resources in coap. Online version at https://datatracker.ietf.org/doc/draft-ietf-core-observe/, 2014.

[5] J. M. Hernández-Muñoz, J. B. Vercher, L. Muñoz, J. A. Galache, M. Presser, L. A. H. Gómez, and J. Pettersson. *Smart cities at the forefront of the future internet*. Springer, 2011.

[6] J. Hui and P. Thubert. Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks. RFC 6282 (Proposed Standard), Sept. 2011.

[7] Z. Shelby, K. Hartke, and C. Bormann. Constrained Application Protocol (CoAP). draft-ietf-core-coap-18 (Internet-Draft), June 2013.

[8] W3C. Resource description framework (rdf). Online version at http://www.w3.org/RDF/, 2014.