# A Wireless Application Overlay for Ubiquitous Mobile Multimedia Sensing and Interaction

Hanno Wirtz, Jan Rüth, Torsten Zimmermann, Matteo Ceriotti, Klaus Wehrle
Chair of Communication and Distributed Systems (COMSYS)
RWTH Aachen University
{wirtz, rueth, zimmermann, ceriotti, wehrle}@comsys.rwth-aachen.de

## ABSTRACT

Mobile users are continuously surrounded by heterogeneous (mobile) devices that provide a basis for direct multimedia interaction within the immediate physical context. However, current multimedia approaches are restricted to sense the availability of devices and their respective multimedia application content through Internet services that manage the locality of devices and content. Each service thereby manages only the narrow context of each device according to the respective application, preventing devices and approaches from sensing the diversity, spontaneity, and dynamics of mobile contexts and the physical interaction scope.

Orthogonally, we facilitate direct and ubiquitous mobile multimedia sensing and interaction in 802.11 and 802.15.4 within the unrestricted physical context of mobile devices. Removing the need for Internet services or wireless networks, we derive a wireless service overlay over the sensed media in the locality of a device that exposes media to applications. Our technical evaluation shows the feasibility of the underlying uncoordinated, unrestricted network-less communication. We substantiate our design along an implementation of traditional and novel mobile use cases, illustrating the design space made accessible by our approach.

## Categories and Subject Descriptors

C.2.1 [**Computer Communication Networks**]: Network Architecture and Design—*Wireless communication*

## Keywords

Ubiquitous computing, mobile multimedia sensing

## General Terms

Design, Measurement, Performance

## 1. INTRODUCTION

Within their local wireless communication range, mobile users are ubiquitously surrounded by heterogeneous wireless devices and the multimedia application content they offer.

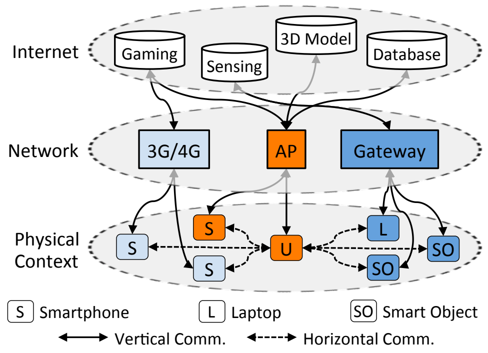

**Figure 1: Mobile multimedia approaches detach application contexts from the physical context of users (U) into restricted, hierarchical networks and sense interaction mediately through Internet services.**

Sensing devices and their applications, i.e., discovering their availability and semantics, then immerses mobile users in a unique and dynamic multimedia context within their *physical wireless context*, as illustrated in Figure 1. Specifically, this context is given by the reachable devices in immediate wireless communication range. A diverse set of novel applications becomes possible through direct interaction between devices in this multimedia context, e.g., mobile gaming [18], media sharing [17], real-time sensing and interaction [13, 16, 17], mobile augmented reality [8, 19, 20, 30], and ubiquitous networking [6]. Notably, the spontaneity and locality of physical contexts offers a rich basis for simultaneous applications.

However, current realizations of multimedia applications take an Internet-based, centralized, per-application approach to mobile multimedia sensing, as illustrated in Figure 1. Specifically, each mobile application senses only its relevant information from the physical context in an isolated fashion, e.g., the GPS position of the device, disregarding all other available information. An application-specific Internet service then creates application context from local information and centrally manages and processes it for all application participants. In the aforementioned applications, services are information [6, 16] or media databases [17], 3D models [19, 20], and gaming [18] or processing [8, 30] infrastructures.

To illustrate the problem, imagine a walkie talkie application that enables mobile users to send and receive audio broadcasts between all peers in communication range, e.g., for informative or emergency purposes. Currently, the application has no means of directly sensing available peers. Instead, it relays the broadcast, via its 3G/4G or 802.11 net-

*MMSys '14,* March 19 - 21 2014, Singapore, Singapore

ACM 978-1-4503-2705-3/14/03 ...$15.00.
http://dx.doi.org/10.1145/2557642.2578224 .

work (cf. Figure 1), through an Internet service. The service manages the locality of devices and determines peers to send the audio file to, recreating the physical context inherently given by the devices in wireless communication range.

Similar, *Ingress* [18] offers mobile locality-oriented gaming, enabling users to fight for real-world landmarks in a virtual Internet-based scenario representation. Hence, the gameplay motivates sensing users and landmarks directly, as users need to be in the vicinity of landmarks and each other to interact. However, users only relay their position, out of the available physical context information, to an Internet service that facilitates gameplay interaction and scene representation. In contrast, the physical context of a mobile user inherently provides the locality, landmarks, and other participants that make up the gameplay.

We thus argue that this detachment of application context from the physical context prevents sensing the locality, diversity, and ubiquity of contexts to leverage them for multimedia applications. Specifically, the requirement of 802.11 and 3G/4G carrier networks for users to be part of the network prior to communication i) induces a hierarchy in communication in contrast to the inherently "flat" physical wireless context and ii) restricts the sensing and communication scope to the current network. Instead, ubiquitous multimedia sensing and interaction envisions sensing peers based on locality and multimedia applications. Multimedia applications, in turn, could then account for the ubiquitous character of applications and content, departing from the current centralized model of content discovery and provision.

In this paper, we thus propose a *wireless application overlay* that makes the diversity and ubiquity of sensed devices and multimedia content within physical contexts available to applications. Accounting for the communication requirements of ubiquitous mobile multimedia sensing, We interpret mobile multimedia sensing as discovering multimedia at devices in the vicinity, as a subsequent step to physical sensing of multimedia data, such as sound or video. Hence, we abstract from wireless communication technologies in a network-free communication layer to allow comprehensive and unrestricted sensing of devices and content in current wireless physical contexts. We incorporate both 802.11 and 802.15.4 in this layer, extending our previous work on network-free 802.11 networking [28] and acknowledging the increasing role of smart 802.15.4 objects in Internet of Things (IoT) scenarios. Leveraging this comprehensive scope, we propose interest-based addressing that combines peer and multimedia application content sensing in a single step. From the set of sensed devices and their applications, each device then derives a wireless application overlay over its physical, local wireless context, providing an access and interaction mechanism to mobile multimedia applications.

In this, we do not aim to replace Internet-based, per-application multimedia approaches. Rather, our approach coexists with current communication techniques and augments global Internet-based application scopes with a local and ubiquitous sensing and interaction mechanism.

The rest of the paper is structured as follows. In Section 2, we delimit our design from current commercial and academic approaches and revisit our work on network-less wireless communication. We present our design of comprehensively sensing and unrestrictedly interacting with 802.11 and 802.15.4 devices based on interest-driven addressing in Section 3. Section 4 describes our system implementation for commodity netbook and smartphone devices and evaluates its communication performance. We sketch the resulting design space in novel ubiquitous multimedia use cases in Section 5 and conclude in Section 6.

# 2. BACKGROUND

Several commercial and academic approaches address mobile multimedia networking and interaction. We thus discuss prominent approaches to establish the state of the art and limitations with regard to the vision of ubiquitous and immediate mobile interaction. Furthermore, we briefly revisit our previous work on association-less 802.11 networking in CA-Fi [28] as a building block in our design.

## 2.1 Related Work

We distinguish related commercial and academic approaches by their design structures and components.

**Commercial approaches:** Tailored to a single specific use case, commercial approaches employ a centralized approach of sensing and backend services for simplicity, maintenance, full control over content and traffic, and performance, e.g., in databases [6,19] and delivery or gaming infrastructures [18]. Following the specified use case and context, they disregard additional local or spontaneously arising context information as well as coexisting applications in favor of narrow contexts and optimized communication mediated by carrier networks and Internet infrastructures. In this, optimized and central media provision prohibits the flexible interaction with sensed devices in the surrounding, narrowing down interaction to predefined media contexts. This induces a dependence on uplink capacity, which proves to be a rising bottleneck [7].

We *complement* this proven design by unmediated, unrestricted communication between devices in local communication range in a single broad mechanism. Indeed, we envision information provision and interaction to occur in a flexible, spontaneous, and infrastructure-less manner within mobile and local contexts. We thereby meet both the proliferation of wireless devices capable to create and distribute media content as well as the arising ubiquitous contexts.

Closest to our goals, industry standards Wi-Fi Direct [27] and Bluetooth, via its Service Discovery Protocol (SDP) [3], enable device-to-device discovery in radio range. Both result in hierarchical network structures being established between devices, restricting the communication range. In contrast, we envision communication to remain unrestricted. However, both technologies may be integrated in our design (cf. Section 3.1) as alternatives to 802.11 and 802.15.4 sensing, providing further pillars of direct communication.

**Academic approaches:** Traditionally, multimedia approaches assume a pre-defined network infrastructure as the basis for communication and interaction, spanning infrastructure-bound wireless networks [11] as well as mobile and distributed networks [4,9,15]. Similar, context-driven systems [16,17] exchange and interact with data over existing network infrastructures, negating the immediateness of derived and derivable contexts. Even direct mobile-to-mobile interaction, e.g., in mobile gaming [24,31], typically avoids a communication channel between the participating phones. Inherently, these approaches incorporate the hierarchy of network associations in their design, a trait we strive to complement in our solution, where unrestricted communication constitutes the building block for interaction.

Closest to our vision of ubiquitous connectivity and interaction, heterogeneous networking [1] fuses (hierarchical) cellular communication with local, mobile 802.11 networks such as MANETs. Immediate multimedia applications then enable, e.g., cooperative video streaming [12], compensating for cellular connections of poor quality by locally merging partial downloads of the overall video. Still, local communication is restricted to the scope and maintenance of each 802.11 network and suffers from the associated overhead of 802.11 as well as establishing routing, layer 2 and layer 3 addressing, and service discovery. In contrast, we envision unrestricted ubiquitous communication to include all devices that are sense-able within the physical vicinity. Nonetheless, we design our approach to complement and extend the merits of 802.11 and cellular communication.

## 2.2 CA-Fi – Ubiquitous Network-less Wireless Communication

In earlier work [28], we proposed CA-Fi (Concurrent Association-less Wi-Fi), mitigating the overhead and restrictions of 802.11 for mobile ubiquitous wireless networking. We identified two main requirements of mobile communication and addressed them in our design.

**Minimal Overhead and Coordination:** The main overhead hindering mobile communication stems from maintaining, discovering, selecting, and joining 802.11 networks. Maintaining a permanent 802.11 network furthermore exceeds the energy resources of mobile devices. Also, 802.11 networks inherently prevent communication with non-associated devices, severely restricting the communication scope. CA-Fi thus enables association-less communication by broadcasting and receiving network-less 802.11 frames with data rates up to 30 kB/s while preserving *concurrent* association-based networking for data-intensive approaches. Leveraging wireless overhearing, CA-Fi sends frames on overlapping 802.11 channels, e.g., {2,5,8,11}, to account for unsynchronized or associated receivers in uncoordinated scenarios.

**Flexible Addressing:** CA-Fi disposes of layer 2 and layer 3 addressing along with 802.11 networks, motivating application-based, immediate addressing that eliminates service discovery overhead. CA-Fi directly uses the diversity of mobile applications and their identifiers, e.g., mobile social networking IDs, as input for hash-based addressing using Bloom filters in 802.11 frames. Identifiers of messages are thereby sequentially inserted in the Bloom filter to enable a receiver-side, driver-level check for identifiers of interest with high probability. Frames then carry a bounded number of messages whose identifiers are aggregated in the Bloom filter address field. We thereby support flexible and efficient addressing and aggregation of multiple messages in single frames, saving the energy and space overhead of per-message frames as in 802.11. Applications then register (multiple) identifiers of interest, e.g., a mobile social networking ID or an application name, with CA-Fi to send and receive frames with this identifier. Frames with other identifiers are filtered.

We implemented CA-Fi in the *ath9k* 802.11 driver and evaluated the performance in terms of communication and energy efficiency on commodity wireless netbook devices. Our results show that CA-Fi enables instant and unrestricted communication in the wireless medium while preserving 70 % of association-based 802.11 throughout and improving energy consumption by up to 44 % compared to the 802.11 adhoc mode. Building on the increased communication scope, i.e., the whole transmission range of a device, CA-Fi may serve as a stand-alone ubiquitous communication channel for low-volume transmissions. Approaches requiring higher data volumes may leverage this channel for discovery and network negotiation to spontaneously establish mobile high-bandwidth 802.11 networks. The application design space of CA-Fi thus allows efficient and ubiquitous heterogeneous opportunistic interaction such as in BUBBLE Rap, Floating Content, and MobiClique as well as simultaneous network-based approaches resulting from opportunistic interaction.

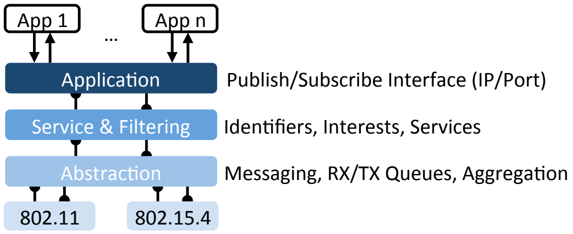

**Figure 2: High-level system design overview.**

Originally, we designed CA-Fi to support opportunistic peer-to-peer networking within urban scenarios. In this paper, we employ CA-Fi as a technical building block towards mobile sensing and communication in a complete system that supports and makes accessible the heterogeneity of multimedia content. We therefore use CA-Fi alongside an 802.15.4 channel and introduce a versatile abstraction layer, semantic and interest-based addressing, as well as versatile visualization and interaction layer.

# 3. A WIRELESS APPLICATION OVERLAY

Figure 2 provides an overview over our layered design. In order to meet the diversity of wireless devices, we introduce an abstraction layer to enable unrestricted, i.e., network-less, and direct communication with 802.11- and 802.15.4-based devices in the vicinity (Section 3.1). Given the resulting broad communication scope and diversity of mobile multimedia approaches, we enable discovery and filtering of applications offered by devices via information- and interest-based addressing (Section 3.2). Applications then utilize the resulting mechanism in a publish/subscribe fashion as provided in an application layer (Section 3.3).

## 3.1 Communication Abstraction Layer

The abstraction layer caters to two functionalities, namely a uniform sensing mechanism that accounts for the specifics and co-existence of 802.11 and 802.15.4 as well as addressing in direct communication. It thereby hides the specifics of underlying communication technologies from the higher layers to provide a uniform interface for direct discovery and communication. As such, it adapts messages from higher layers to the respective indicated communication technology, e.g., with regard to differing frame sizes, and vice versa. It furthermore facilitates communication with the respective driver and operating subsystem of each technology. Currently, our system supports 802.11, to account for personal and mobile communication, and 802.15.4, to include dedicated small-scale devices in, e.g., IoT deployments. Further technologies, e.g., Bluetooth low energy (802.15.1), can in the future be incorporated by adding the respective abstraction.

### 3.1.1 *Sending (and Receiving)*

In 802.11, devices broadcast standard reserved 802.11 frames at the 802.11 base rate of 1 Mbit/s on the layer 2 broadcast address *without any* network identifiers. The resulting maximum payload size is 1500 byte per frame and frames are neither acknowledged nor retransmitted. Time-to-live (TTL) and retransmission (RTx) parameters for messages indicate spatial and temporal dissemination ranges, respectively.

802.15.4 in principle supports direct, i.e., association-less, networking over a broadcast personal area network (PAN) identifier `0xffff` that all devices can choose to receive. Consequently, within this "network", devices in our design address frames to the broadcast address, enabling interest- and information-based semantic addressing which we detail in the next section. Devices then send 802.15.4 frames with a preset maximum data rate of 250 kBit/s and a maximum payload of 100 byte. Please note that we do not change medium access functionality in either technology, allowing unchanged contention-based access to the whole wireless medium in transmission range.

Both 802.11 and 802.15.4 operate in the ISM 2.4 GHz spectrum, leading to destructive cross-interference [23, 32][1]. In the potential absence of non-overlapping channels, our design introduces this interference already at the sending device. We address this issue by implementing a slotted sending scheme that separates 802.11 and 802.15.4 transmissions temporally. The abstraction layer therefore enqueues messages uniformly, regardless of the target interface, and sequentially delivers them to the respective TX queue. Please note that the execution speed of message slotting does not introduce any limitations on the effective sending rate of either interface.

### 3.1.2 *Addressing*

Addressing in direct, i.e., network-less and fully spontaneous, communication necessarily departs from ISO/OSI-compliant addressing. In fact, layer 3 addressing, e.g., via IP addresses, is only viable in managed, restricted networks and layer 2 device addresses only identify devices in discovery but can not express media or services. In contrast, ubiquitous communication requires an emphasis on discovering previously unknown devices and applications spontaneously.

As such, we propose immediate and semantic, two-tiered addressing in direct communication to allow i) indication of interests directly in wireless frames, and ii) filtering of received frames immediately based on the respective interests. Frames therefore carry a first-tier *categorical* identifier that indicates the multimedia use case and broad type of information contained in frames. Examples for category identifiers are " IoT building automation/information", "mobile social networking", "mobile gaming", or "tourist info".

First-tier identifiers then form one input to CA-Fi's Bloom filter address field in 802.11 frames and thereby allow positive filtering of frames by category at the driver level. Analogously, we implement a hash-based identifier field in 802.15.4 frames. However, we account for both the dedicated deployment of 802.15.4 devices as well as the limited frame payload size by neglecting aggregation of identifiers, as in Bloom filters, and choose a single-input hash function with limited-length output. Specifically, we use a `lookup3 Jenkins` hash with an output length of 4 byte.

| Stacked | Type | Desc ID | Latitude | Longitude | Length | Service |
|---|---|---|---|---|---|---|
| 1 bit | 7 bit | 32 bit | 32 bit | 32 bit | 8 bit | ID |

**Figure 3: Service discovery and announcement frame. Service specifics are optional in discovery.**

| Stacked | Type | Request ID | Seq-Nr | Count | Length | Data |
|---|---|---|---|---|---|---|
| 1 bit | 7 bit | 32 bit | 8 bit | 8 bit | 32 bit | |

(a) Request frame.

| Stacked | Type | Response ID | Seq-Nr | Count | Length | Data |
|---|---|---|---|---|---|---|
| 1 bit | 7 bit | 32 bit | 8 bit | 8 bit | 32 bit | |

(b) Response frame.

**Figure 4: Service request and response frames.**

However, such categorical identifiers are too coarse to allow sensible sensing, addressing, and selection of devices and provided applications, requiring an additional identifier (tier). Conversely, directly including specific service identifiers in frames in hashed form would require pre-defined and pre-distributed identifiers for every service and would thus prevent human-readable selection of unknown applications and meaningful discovery. We detail the use of second-tier identifiers for semantic, network-free service addressing in the next section.

## 3.2 Service and Filtering Layer

Devices consume and may offer multiple types of interaction concurrently, examples are mobile gaming, ubiquitous networking, and mobile social networking. We thus regard each as a distinct *service* in handling and addressing, especially for discovery and announcement. A service thus internally represents multimedia interaction with any device as discovered, registered, and exposed to applications in our design. The service and filtering layer provides the necessary handling methods and filters specific services of interest *after* the broad categorical filtering on the abstraction layer.

To this end, services know their interaction functionality (e.g., "elevator") within their broad category ("IoT building automation") and use the hashed value of this functionality as the second-tier identifier, the *describing ID* of a service (Desc ID). In any given category, the relatively small set of meaningful functionalities allows users to perform intuitive discoveries and queries. Services furthermore may listen to multiple (linguistic) representations of a functionality, e.g., "elevator", "elevator control", "lift", "ascensor". We make use of this identifier in sensing (cf. Figure 3), allowing devices to *optionally* specify the desired service functionality without requiring an explicit service identifier. Note that the abstraction layer always prepends the category identifier. In discovery, missing describing IDs are treated as a wildcard within a category. Services that match the category identifier and, if set, the describing ID reply to observed discovery frames with an announcement frame that carries their describing ID, GPS coordinate (if available), and an ID that expressively identifies the service (e.g., "elevator second floor left wing") (cf. Figure 3). Announcement frames have the same structure as discovery frames but differ in their *Type* field and include service details, e.g., location and service ID.

Devices may thereby toggle their announcement behavior based on their current context. For example, an elevator control service may announce itself periodically during working

---

[1] Channels 25 and 26 in 802.15.4 constitute a singular, situational exception as they do not overlap with certain 802.11 channel assignments in North America.

hours and otherwise only respond to discovery frames. Similar, mobile, location-based gaming [18] may toggle announcements by the availability of a GPS signal, i.e., outdoors, in order to reliably determine the location.

After discovery of a service, i.e., reception of an announcement frame, direct communication employs the indicated *service ID*. Specifically, hashing the service ID allows using the resulting value as a semantic address in frames on the abstraction layer to positively filter all subsequent data frames (cf. Figure 4). *Request* and *Response IDs* thereby identify a communication "stream", while sequence numbers (*Seq-Nr*) and a *Stacked* bit allow fragmentation and ordering of payload over multiple frames.

Please note that, upon discovery of a service, subsequent communication may require a high-volume communication channel, e.g., to transmit or stream large multimedia content such as video. By design, our approach supports concurrent association-based 802.11 networking [28] as well as spontaneous, mobile network negotiation and establishment triggered by ubiquitous application sensing. Interaction may then simultaneously leverage ubiquitous sensing and the performance and coordination of purpose-driven networks.

In direct interaction, applications can register to describing IDs and service IDs. In the next section, we briefly describe the application layer that provides an interface to direct communication for (legacy) applications.

## 3.3 Application Layer

Direct, ubiquitous communication trades network structures for unrestricted access to the wireless medium on a per-application basis. It thus departs from the traditional IP-based way of interconnecting devices and, in effect, applications. Therefore, the application layer serves as an interface that abstracts from network-less, direct communication and addressing to enable adoption by higher-layer (legacy) applications.

To this end, the application layer exposes service management in direct communication over a traditional IP/port-based socket interface that applications may address locally and centrally. Simple primitives then enable registration of a service, defining announcement behavior, requesting discovery of services, and subsequent interaction. The publish/subscribe interface provided by the application layer thereby reflects the uncoordinated spontaneity and asynchrony of ubiquitous wireless communication in mobile scenarios, delivering messages to and from applications based on the sensing of communication partners, i.e., devices with matching interests or applications. Services, as exposed by the service layer, thereby become traditionally addressable data sources and recipients, allowing easy adaptation of existing interaction primitives from network-based interaction.

# 4. TECHNICAL EVALUATION

The proposed wireless application overlay operates outside of the restrictions of traditional 802.11 or 802.15.4 networking to achieve unrestricted wireless communication. As such, we embrace a tradeoff between losing coordination functionalities, e.g., as provided by an 802.11 AP, and increasing the communication scope and spontaneity. In this section, we briefly present our proof-of-concept implementation (Section 4.1) and evaluate the real-world communication performance of our proposed design to highlight its feasibility.

**Figure 5: Walkie Talkie app (cf. Section 5.2) on prototype with external 802.11 and 802.15.4 radios.**

We thus first measure the viable communication range with regard to throughput and packet loss to derive a notion of the actual context scope around a device (Section 4.2). With regard to uncoordinated communication scenarios in populated areas, we then evaluate the impact of multiple competing senders on this performance (Section 4.3). Conversely, we quantify the benefit of relaying messages for the benefit of devices outside of communication range (Section 4.4). Last, we address the co-existence of 802.11 and 802.15.4 transmissions in the 2.4 GHz frequency band at the sender (Section 4.5).

## 4.1 System Implementation

We implemented our approach for both Linux netbooks and Android smartphones. The netbooks thereby serve for debugging and performance analysis while smartphones are our target device category with regard to feasibility studies and application development. Specifically, the netbooks are Lenovo S10-3 IdeaPads and a Dell Inspiron 1090 Duo that run Ubuntu 12.10 on top of a 1.5 GHz dual-core CPU and an Atheros AR9285 802.11n wireless card that uses the `ath9k` 802.11 driver. Smartphone devices are Samsung Galaxy Nexus phones running Android 4.2.2 on top of a 1.2 GHz dual-core CPU and a Broadcom BCM4330 802.11n wireless card. Because the Broadcom card uses an 802.11 driver building on closed-source MAC functionality in firmware, we attached an external TP-Link TL-WN822N 802.11n module that includes an Atheros chip running the `carl1970` 802.11 driver (cf. Figure 5). This driver shares the majority of its code with the `ath9k` driver, some adjustments to the Android system and the driver allowed us to run all 802.11 parts on otherwise closed-source smartphone devices. Note that we deactivate the internal wireless card and thus only operate a single 802.11 interface.

We integrate 802.15.4 support by attaching Digi XBee Series 1 radios to the respective devices via a serial-to-USB interface. As Figure 5 shows, adding a second USB module to the smartphone required us to make use of an external USB hub and prevents a meaningful evaluation of the energy characteristics of our approach. This design will be made obsolete once future smartphone devices support 802.15.4 as well as make use of open-source 802.11 drivers.

We conduct communication performance evaluations using the netbook devices because of usability and an insufficient number of smartphones. However, identical hardware in the smartphone setting suggests equivalent results.

## 4.2 Communication Performance

To substantiate our proposal, we quantify both the performance as well as the robustness of 802.11 and 802.15.4

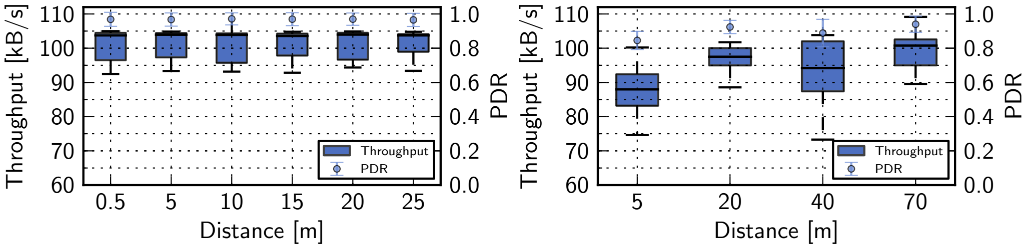

(a) 802.11 indoor.

(b) 802.11 outdoor.

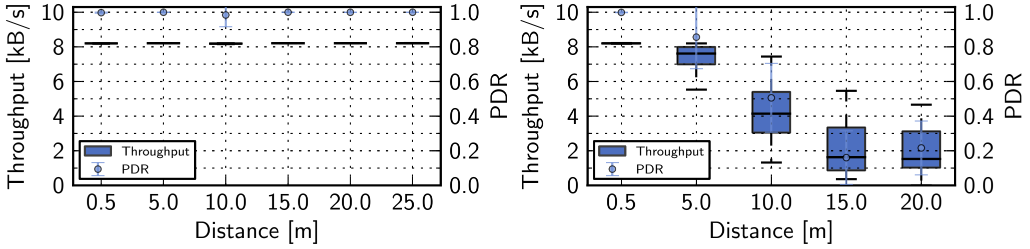

(c) 802.15.4 indoor.

(d) 802.15.4 outdoor.

**Figure 6: Box plot of device-to-device throughput and packet delivery rate (PDR) over distance in associationless 802.11 as well as 802.15.4 communication in indoor and outdoor environments.**

transmissions in ubiquitous, network-less wireless communication. We thus vary the distance between two communicating devices and measure both the achieved throughput as well as the packet delivery rate (PDR). Our evaluation is driven by two questions, namely i) what are the real-world dimensions of the physical context of a mobile device in which viable communication performance is possible within the respective communication technology, and ii) what are deployment and usage scenarios for 802.11 and 802.15.4 appropriate to the observed performance characteristics?

Figure 6 shows a box plot of throughput and the PDR over distance of 30 runs when transmitting 1000 packets of 1500 byte in 802.11 and 1000 packets of 100 byte in 802.15.4 in an indoor office setting as well as outdoors. Associationless 802.11 communication thereby achieves throughput rates of 100 kB/s indoors as well as outdoors, supporting a diverse set of mobile multimedia applications over a range of up to 70 m[2]. Notably, communication over small distances outside shows both the least robustness, as hinted at by a PDR of 0.85, and lowest throughput with regard to the measured distances. This is because the sending device was placed close to our campus building with a high amount of 802.11 traffic, resulting in strong interference around sender and receiver for small distances. In contrast, transmissions over longer distances show a higher PDR as well as higher throughput.

802.15.4 communication achieves constant 8 kB/s in indoor scenarios. In contrast, outdoor communication in 802.15.4 suffers from the increased exposure to interference and loss of positive multi-path propagation, showing an exponential decrease of the achieved throughput. Especially, measurements over 25 m did not produce any usable results and are thus not included in Figure 6(c).

The results support the originally intended usage of 802.15.4 for dedicatedly installed devices that expose singular, low-volume interaction and 802.11 as the basis for personal communication in larger radii. With regard to our design, the discrepancy between indoor and outdoor performance in 802.15.4 further suggests a functional separation. Namely, 802.15.4 appears especially suited to discover singular short-range indoor interaction with "things", e.g., wireless interaction with elevators, doors, signs, and localization mechanisms. In contrast, 802.11 proves capable of sensing and sustaining multiple, heterogeneous multimedia approaches simultaneously within substantial ranges, giving rise to a rich set of exploitable contexts. Subsequent to sensing multimedia content, high-volume approaches may build support for concurrent 802.11 networking to establish association-based 802.11 networks with high data rates spontaneously and purposefully.

## 4.3 Uncoordinated, Ubiquitous Communication

In our design, we propose network-less, ubiquitous, uncoordinated communication without a coordinating entity, such as an 802.11 AP. Multiple senders may thus transmit in the same physical space simultaneously, resulting in interference and degraded communication performance at the receiver. In this section, we thus quantify the impact of uncoordinated competing transmissions on direct communication in 802.11 and 802.15.4. We thereby analyze the feasibility and performance of ubiquitous communication in populated areas where multiple, heterogeneous multimedia approaches concur, e.g., in urban or campus areas. In such areas, a mobile user wants to sense the available devices and subsequently subscribe to or request communication with those of interest. Figure 7 shows the average and standard deviation in packet delivery rates (PDR) over 30 runs at a receiver that observes concurrent, uncoordinated transmissions from equidistantly placed senders over 5 m in 802.11 and 802.15.4. In both 802.11 and 802.15.4, senders strive for maximum throughput, cf. Section 4.2. Note that, by continuously sending at maximum data rates, we evaluate the worst case scenario in which multiple relatively high-bandwidth transmissions compete for access to the wireless medium.

---

[2] Please note that we set the communicating devices to the same 802.11 channel. Transmitting on multiple 802.11 channels to support receivers associated to independent fixed-channel 802.11 networks [28] limits throughput to 30 kB/s.

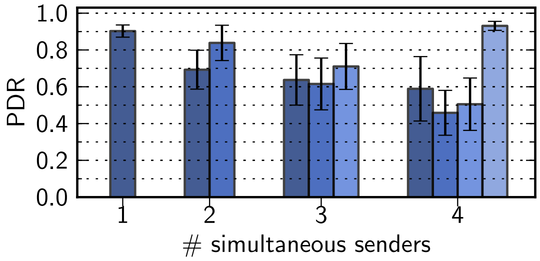

(a) 802.11.

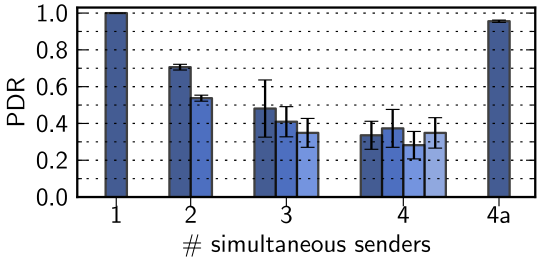

(b) 802.15.4.

**Figure 7: Packet delivery rate (PDR) of competing transmissions as measured at the receiver.**

As illustrated in Figure 7(a), reception rates in 802.11 decrease with an increasing number of simultaneous senders. However, reception rates do not decrease linearly with the number of senders, indicating a good performance of carrier sense in CSMA/CA as well as robust modulation in DSSS when using the 802.11 base rates. Using sequence numbers and stacked bits in frames then enables devices to re-request frames that are missing from opportunistic reception. While we do not explicitly state retransmission mechanisms, per-application strategies can be specified in the service layer as a part of application interaction. Please note, the performance strongly depends on the actual device characteristics such as antenna placement, as hinted at in the results for four senders, in which the Dell netbook dominates the other senders.

In contrast, each reception rate suffers drastically in 802.15.4, cf. Figure 7(b), when observing multiple transmissions. Given an identical evaluation scenario as for 802.11 and a comparable MAC protocol, the decrease seems to indicate a lower performance in carrier sensing. We thus propose 802.15.4 devices to only respond to observed requests to save channel quality for actual data transmissions that arise from such requests. In this configuration, senders transmit announcement messages periodically once per second and only start transmitting payload data upon request. Our additional evaluation result (cf. Figure 7(b) "4a") highlights the feasibility and benefit of this configuration. Approximately, senders in this case achieve a PDR as in the one sender case. With regard to the proposed application of 802.15.4 devices for short-range, dedicated single-purpose application entities, e.g., exposing building functionality, we argue that triggering transmissions, in contrast to continuous sending, actually fits the usage scenarios. For example, communication with an elevator does not require continuous pushing of information but is inherently request-driven.

## 4.4 Benefits of Relaying

In the previous evaluation, we measured the negative impact of concurrent communication in densely populated scenarios. Conversely, mobile devices may exploit this density by cooperatively serving as relays of received transmissions for devices that might be out of range.[3] Devices out of range may then build upon the increased announcement range by moving towards the GPS coordinates indicated in announcements of interest. In this evaluation, we thus strive to quantify the benefit of relaying in terms of transmission performance as well as the negative impact of replicating received transmissions in addition to original ones. To this end, we create a bad link between a communicating pair of devices by placing them at the edge of their transmission range. We then introduce up to two relaying devices and measure both the impact on the packet delivery rate (PDR) and the completion time of a given transmission. Note that we do not include 802.15.4 in this evaluation because of the drastic decrease of communication under competing transmissions, cf. Section 4.3.

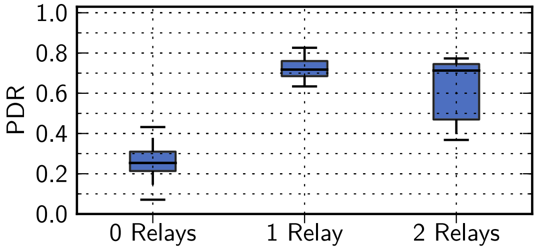

(a) Packet delivery rate (PDR).

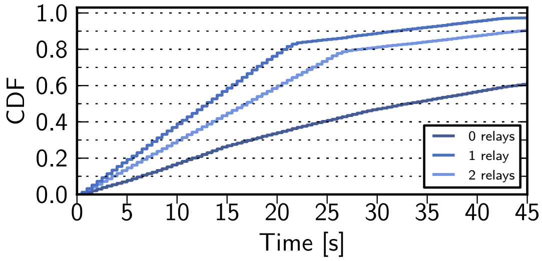

(b) Cumulative distribution function (CDF), indicating transmission completion, over time.

**Figure 8: Benefit of relaying transmissions in 802.11 with regard to the number of relay devices.**

Figure 8 shows the impact on both metrics for 802.11 transmissions of 1000 packets of 1500 byte each over 30 runs. Both relaying setups show an increase of the PDR at the receiving device, although the higher variation when using two relays hints at a slight increase in collisions and layer 2 retransmissions. Similarly, both setups speed up the delivery progress measured over time, with exclusively using one relay performing moderately better than using two relays. We derive two results from this measurements. First, relaying provides a viable method of extending the range of 802.11 services, which might prove especially useful for announcement messages that indicate a geographic location, enabling devices to move in range of the announced service. Relaying transmissions may thereby be incorporated into the communication queue at a device by aggregating transmissions in CA-Fi based on their size or using a hierarchical sorting, with relayed transmissions being processed later. Second, multiple

---

[3] Relaying transmissions over $n$ hops may be indicated by setting the TTL field in CA-Fi frames accordingly.

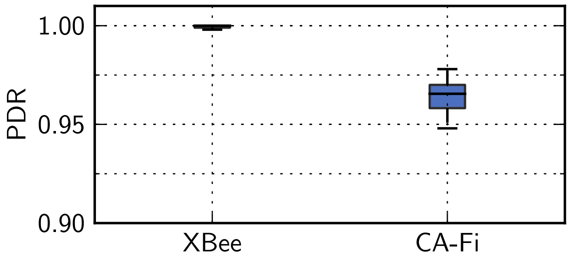

**Figure 9: Packet delivery rate (PDR) of simultaneous, sender-side 802.15.4 and 802.11 transmissions as observed by separate receivers.**

relays introduce an overhead of redundant transmissions. To account for this, we propose to separate relaying devices spatially, with devices that overhear a relayed transmission, by way of the TTL field in CA-Fi frames, ceasing to act as a relay themselves. Devices that do not observe another relay start acting as a relay for designated transmissions.

## 4.5 Co-existence of 802.11 and 802.15.4

In our design, we propose to simultaneously use both an 802.11 and 802.15.4 radio to sample and interact with the diversity of services that build on the respective characteristics of these wireless access technologies. Both 802.11 and 802.15.4 thereby operate in the 2.4 GHz ISM band, resulting in interference between simultaneous transmissions. Following the results of earlier measurements of this interference between separated devices [23], we thus quantify efficiency of our slotted sending scheme in minimizing the two-way impact of concurrent 802.11 and 802.15.4 transmissions. We therefore place the netbook within 5 m of an 802.11 and 802.15.4 receiver and transmit 1000 packets of 1500 byte over 802.11 and 1000 packets of 100 byte over 802.15.4.

As illustrated in Figure 9, coexisting transmissions do not negatively influence each other when temporally slotting the respective transmissions. In contrast, simultaneous transmissions, as shown in [23], introduce a negative impact on the communication performance, motivating cross-technology coordination mechanisms as for example in [32]. Conversely, we enable and deduce the feasibility of a sender-side multi-radio design that adheres to a shared spectrum by temporally separating competing transmissions.

# 5. USE CASE EVALUATION

We envision ubiquitous direct multi-media interaction to enable a wide scope of applications that are currently restricted to the limitations of single network infrastructures. In this section, we thus strive to illustrate the resulting design and application space along a set of example use case implementations that directly build on and benefit from the unrestricted communication scope. Please note that all applications are developed for the Galaxy Nexus phone as smartphones constitute our target device category. Screenshots thus exclusively show the respective Android applications.

## 5.1 Visualization and Interaction

Multimedia sensing in the full range of 802.11 and 802.15.4 results in an unbounded number of concurrently active and observable interaction opportunities within diverse applications. Visualizing information that is filtered according to the interests and applications of the user then is the first step

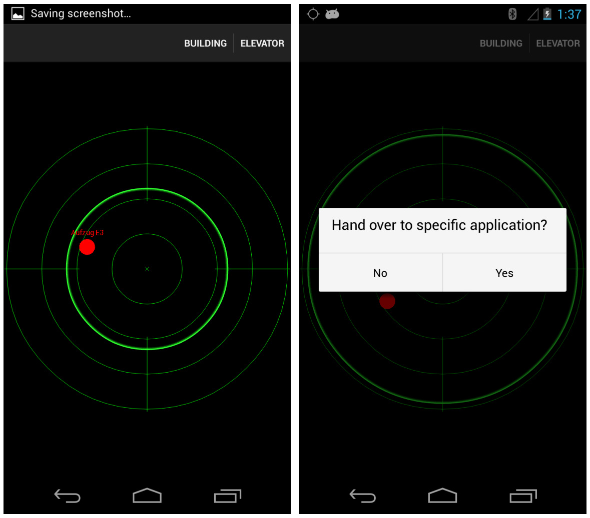

Hand over to specific application?

(a) Visualizing sensed devices/applications.

(b) Starting point for specific applications.

**Figure 10: Device and application visualization and interaction in the wireless application overlay.**

in making applications and devices in mobile contexts accessible to multimedia interaction approaches. In this section, we show our prototypical overlay visualization approach as a launching point for specific multimedia approaches. We further illustrate the design space of such approaches by way of a number of examples as an alternative to network driven and Internet-based implementations.

### *5.1.1 Overlay Virtualization*

Visualizing the result of multimedia sensing in the physical vicinity of a user should reflect their location relative to the user, i.e., projected onto the surroundings. As an intuitive realization, Figure 10(a) shows a radar-like visualization of the sensed devices, equally placing 802.11 and 802.15.4 sources around the user. Note that devices in this visualization example are already filtered by their category, namely building automation, as well as their specific type, namely elevator services (cf. indicators in upper right).

As described in Section 3.2, devices indicate their GPS position in their announcement frames to enable static localization in the overlay representation. Using the GPS position of the mobile device then allows representing distances to sources, while deriving the orientation of the smartphone via the magnetometer sensor allows reflecting the real-world positions. Newly sensed devices, i.e., that appear in transmission range, can thus be made accessible to the user or specific applications without service discovery or 802.11 and 802.15.4 networking overhead. Note that GPS positioning is only feasible in outdoor scenarios; we address this issue by supporting incorporation of building maps for indoor localization of interactive objects (cf. Section 5.1.3).

To trigger interaction in specific local applications, the generic overlay visualization may serve as a launching point. As such, the design and functionality of the overlay may be kept generic while applications realize their specific functionality autonomously. Note that, in contrast to traditional hierarchical approaches, applications have, at any point in time, access to the full observed context of the mobile device and may further be executed simultaneously since sensed

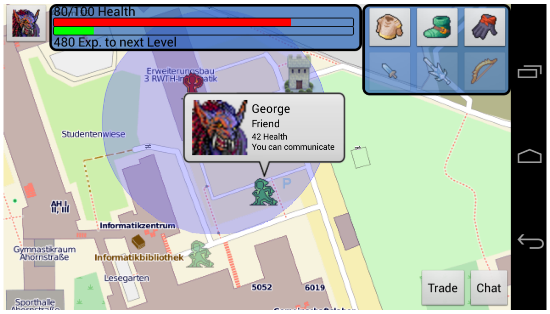

(a) Friendly interaction.

(b) Enemy interaction.

**Figure 11: Mobile gaming in immediate physical vicinity arising from interaction sensing.**

input is processed and presented uniformly in the abstraction layer. In the example illustrated in Figure 10(a), sensing the elevator in our office building (E3) triggers a handover to an application that exposes task-specific functionality. Applications may thereby register for specific interaction by way of intent filters in Android, translating the category and identifier structure into hierarchical schemes or MIME types.

### *5.1.2 Mobile Gaming*

As a striking example for detached multimedia interaction, Ingress [18] rebuilds real-world "landmarks" in a fully virtual, Internet-based game scenario. Once players are in the vicinity of landmarks in the real world, they can fight for this landmark in the virtual world. Mitigating the requirement for Internet communication over cellular networks, mobile gaming building on direct interaction sensing relocates the gameplay in the physical vicinity of players and landmarks and builds on immediate interaction between them.

As an example, we fashioned an outdoor-oriented Android-based mobile game after Ingress, separating players in two fractions. We make use of the `osmdroid` framework to map gameplay to real-world locations and to correctly place landmarks. Interaction between players in transmission range, as illustrated in Figure 11, then translates into combat, communication, or trade actions within the game. As the main difference to Ingress, landmarks would be equipped with an 802.11-capable device and partake directly in the game. Given the proliferation of 802.11 devices and the low communication requirements of the game, we regard this assumption as feasible. We disregard 802.15.4 here because of its low range. Enabling immediate interaction between players then enables new avenues of mobile gaming, for example collaborative or educative gaming in locations of interest.

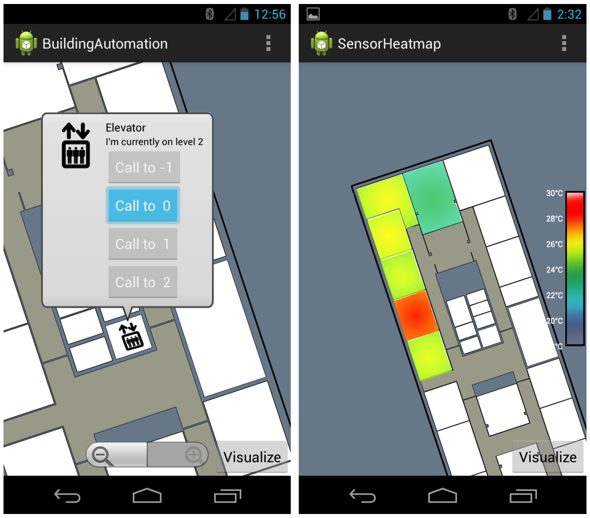

(a) Elevator interaction. (b) Temperature monitor.

**Figure 12: Use cases in building automation.**

### *5.1.3 Building Automation and Interaction*

Building automation approaches, e.g., to monitor temperature or energy expenditure, have been proposed, among others, as a main usage scenario for the Internet of Things (IoT). Deployed smart objects thereby connect building appliances to monitoring and control applications that run on smartphones or tablets. Communication between smart objects typically occurs over low-power 802.15.4 links while data and functionality is exported over a dedicated 802.11 or wired gateway device.

Building on sensing devices directly over 802.15.4, we enable users to immediately interact with IoT appliances. Accounting for the indoor usage scenario, we envision building managers to provide a simple map to (visiting) users, e.g., by way of a download pointed to by a QR code at the entrance. This map then contains the location and ID of interacting smart objects in the building. Objects then announce their ID (upon request), category, and type, allowing placement on the map and visualization, removing the need for GPS localization. When in range, users can interact with objects.

As examples, we implemented the aforementioned elevator interaction app, allowing users to call an elevator once they are in range, and a temperature monitor app, visualizing the temperature in our offices. Figure 12 shows both applications in use, calling the elevator to the current level and visualizing single temperature values by offices in a heat map. The temperature sensor, part of a TelosB sensor node [22], in our office was directly exposed to sunlight to establish an offset to neighboring offices. Note that we were not granted actual write access to the control system of the elevator and thus represent the elevator by a stand-alone sensor node placed at the elevator. Further straightforward appliances that benefit from instant local sensing are, e.g., automatic door openers and lighting control systems.

## 5.2 Speech Transmission

An intuitive multimedia approach that benefits from the increased sensing scope and immediate interaction in the overlay is audio transmissions, specifically transmitting human speech. Traditionally, audio is recorded, transcoded, and

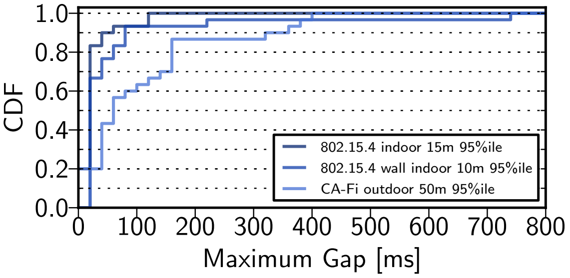

**Figure 13: CDF of maximum audio gaps in WalkieTalkie speech transmissions over 802.15.4 in indoor line-of-sight and office scenarios as well as 802.11 (CA-Fi). CDF functions show the 95th percentile.**

sent, continuously or discretely, through a central Internet service that manages device locations and availability. While this caters to the requirements of coordinated one-to-one communication, use cases such as emergency transmissions, location-based audio sharing [17], and wireless advertisements benefit from spontaneous audio transmissions between unknown participants.

Revisiting our earlier example, we implement a walkie talkie application that allows push-to-talk broadcast communication with nearby devices that subscribed to this service. Figure 5 shows a screenshot of the application. Making use of the smartphone's microphone, the application transcodes human speech using the `Speex` codec [29]. Transcoded speech is then sent in 20 ms blocks via 802.11 or 802.15.4, as indicated by the application. Analogously, the application may send speech recordings, supporting close-range streaming of wireless advertisements or object information, e.g., in museums or a meeting room.

The lack of a managed communication infrastructure, e.g., an 802.11 network, thereby raises the question of the robustness of transmissions. Specific to the walkie talkie application, we need to evaluate whether uncoordinated communication supports a viable quality of service in audio transmissions. We therefore measure the maximum observed gaps in ms, caused by lost speech fragments in frames, within 30 transmissions over 802.15.4 in both an unobstructed indoor scenario and an obstructed one, i.e., through a wall between two offices. For a longer range evaluation, we perform the same measurement over a distance of 50 m using 802.11.

To judge whether these gaps impair the audio service, we assume a minimal value of 150 ms for *observable* disruptions, as recommended by the ITU [10] for one-way audio transmissions. Above this, disruptions up to 400 ms are observed but are still acceptable, while larger disruptions render the service unusable. As Figure 13 shows, transmissions over 802.15.4 are highly robust and, in the majority, do not introduce gaps larger than 100 ms in obstructed and unobstructed scenarios. While the depicted 802.11 transmissions are more affected by range as well as interference, they still provide viable quality of service, with 70 % suffering from a disruption below 150 ms and 95 % below 400 ms. Note that by regarding the *maximum* observable gap, we make a worst case evaluation. The *average* gap was 29 ms in unobstructed 802.15.4 transmissions, 30 ms in obstructed 802.15.4, and 38 ms in 802.11, highlighting robustness of typical transmissions..

Audio transmissions are a fitting example for complementing multimedia approaches that benefit from and build on different types of media. For example, an audio stream may augment mobile gaming or add to recognition approaches, e.g., mobile computer vision as shown in the next section.

## 5.3 Augmented Reality & Computer Vision

Augmented reality (AR) approaches [8, 19, 30] have the potential to seamlessly integrate computer-aided information and interaction with real-world observable contexts. Typically, this requires the computer, e.g., a smartphone, to recognize objects in the current scenery in order to augment them with information or interaction points. In this, the biggest challenge for ubiquitous computer vision (CV) is the dependence on databases that already hold images of the respective objects to facilitate comparison. The actual availability of such databases thereby restricts AR approaches to few selected locations.

We thus propose a novel way of realizing augmented reality approaches through object recognition. Building on immediate sensing and communication in the wireless application overlay, objects directly provide i) ORB [25] key point descriptors to facilitate recognition within the (smartphone) camera picture, and ii) information for immediate annotated augmentation as well as pointers to further information sources, e.g., an Internet URL. Upon reception of descriptors, mobile devices may then recognize the respective object in the camera picture and augment it with the associated information. Such direct provision of relevant descriptors alleviates the need for vast image databases and powerful processing infrastructures as well as mitigates the delay introduced by the associated communication steps. As illustrated in Figure 14, mobile and spontaneous AR applications become possible, e.g., annotating buildings with descriptive information for localization and navigation (cf. Figure 14(a)) or augmenting persons with social networking information (cf. Figure 14(b)).

Transmitting descriptors directly between 802.11 devices[4] introduces a limit on the number of providable descriptors. Also, performing CV recognition on mobile devices requires approaches with appropriate processing power demands. We address this issue in two ways. First, we use ORB descriptors for their comparably small size of 32 byte per descriptor instead of, e.g., SIFT (512 byte), as well as their relatively low processing demands[5]. Given a maximum payload of 1500 byte per frame, we thus are able to transmit up to 45 descriptors in a single frame, depending on the associated annotation data as well as frame headers. Second, we analyze the number of descriptors required for a satisfying recognition rate of everyday objects to substantiate the feasibility of our approach. To this end, we take a representative picture of both our office building (cf. Figure 14(a)) and one of our colleagues (cf. Figure 14(b)) and vary the number of compared descriptors, depending on the source picture. We then evaluate the true positive (successful detection) and true negative rate (correct non-detection) over the number of descriptors by detecting the source picture in a prominent data set containing building [5] and person images [21].

As Figure 15(a) and 15(b) show, rising numbers of descriptors increase the true positive rate in both building and

---

[4] We disregard 802.15.4 here because the low range requires recipients to stand before the object in the first place.

[5] Wang et al. [26] propose visual fingerprints that build on spatiograms and wavelets. In contrast to our work, they do not address actual communication between devices and do not achieve a real-time recognition system.

Computer Science Building Complex E3
@(50.77915,6.060259)

(a) Localization and object annotation.

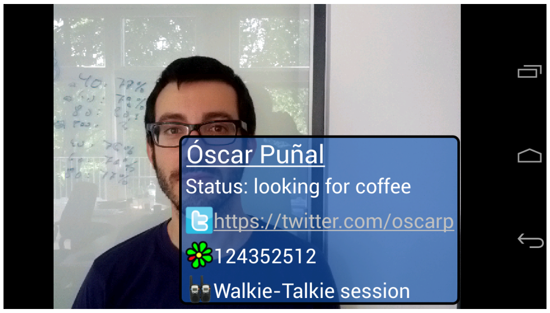

Óscar Puñal
Status: looking for coffee
https://twitter.com/oscarp
124352512
Walkie-Talkie session

(b) Mobile social networking.

**Figure 14: Example use cases for mobile computer vision facilitated by providing ORB [25] descriptors in the wireless application overlay.**

face recognition[6]. Supporting our approach of transmitting a limited number of descriptors, comparisons using only 80 and 83 descriptors, respectively, already allow true positive rates of 77 % for building images and 83 % for face images. Additionally, the true negative rate amounts to 99 % and 98 % providing robustness against false detections.

From these results, we deduce the feasibility of facilitating computer vision within direct transmission of only two 802.11 frames, thereby providing a building block for mobile augmented reality. Object managers may thereby simply extract significant descriptors offline from example pictures of the object and provide these to mobile devices upon the particular service request. Descriptor extraction in our example settings was instant and may be performed over a number of pictures to determine the largest set of stable, i.e., reoccurring, descriptors for an object. Given a reasonable sending frequency, e.g., 2 /s, and number of retransmissions, we argue that the requirement of receiving two 802.11 frames is negligible given the range of 802.11 transmissions. A mobile user may thus be alerted of sensed objects, i.e., buildings, persons, landmarks, etc., and can use its smartphone to both find the object as well as display annotating information.

In terms of practical applicability, detecting a specified object continuously, i.e., in each of the typically 30 captured camera frames per second, proved infeasible due to the repeated processing overhead. We thus reduced the frequency of processed frames to every 5th frame, resulting in 6 detection steps per second and allowing for a usable system. However, these results are unoptimized and specific to the Galaxy Nexus we used as an experimentation device.

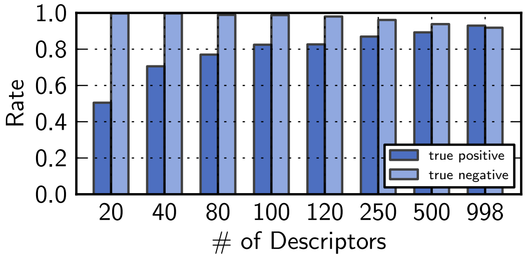

(a) Building recognition in Zürich dataset [5].

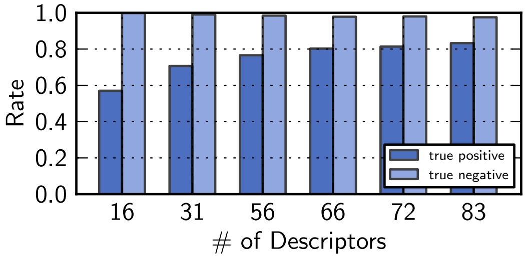

(b) Face recognition in TU Graz dataset [21].

**Figure 15: True positive and negative rate in image recognition over number of transmitted descriptors.**

# 6. CONCLUSION

In this paper, we present a design for ubiquitous multimedia sensing within physical device contexts and the broad diversity of wireless devices and interaction opportunities. We achieve a local wireless application overlay building on this sensing through association-less 802.11 and 802.15.4 networking as well as interest- and application-based addressing. All wirelessly reachable services thus become accessible for multimedia interaction within the dynamics of ubiquitous wireless communication, thereby complementing network- and Internet-based communication.

Our evaluation supports the feasibility of our design, allowing a throughput of up to 100 kB/s in 802.11 and 8 kB/s in 802.15.4 as well as high packet delivery rates. The range and robustness of the respective technologies thereby suggest 802.11 for mobile, heterogeneous interaction in indoor and outdoor scenarios, while 802.15.4 proves especially suitable for short-range interaction with dedicated objects, e.g., in building augmentation. Relaying 802.11 messages in the ubiquitous communication scope then increases the range and motivates a local relay assignment mechanism to avoid interference. Analogous, slotted sending avoids interference between 802.11 and 802.15.4 at the sender. Further efforts will target a fully integrated implementation of our approach as well as an evaluation of energy characteristics.

## 6.1 Limitations

Multimedia sensing through central coordination entities, such as [17, 18], features a consistent management of the availability and application state of mobile participants. Most notably, this is important in mobile gaming to ensure fairness.

Our emphasis on spontaneous and ubiquitous interaction without central entities prevents such coordination as no entity holds the full set of application states. Consistency must therefore rely on distributed checks, for example Lamport timestamps to check event ordering and distributed integrity checks to prevent tampering with or forging of states. In

---

[6] ORB only extracted 83 descriptors from the face picture.

the latter case, we are interested in how approaches from distributed systems research, e.g., P2P and micro payment systems, can be integrated in our design.

Inspired by mobile gaming, recording the interactions that are performed between participants and condensing them into "save states" presents another possibility. Save states could then be transferred to an Internet-based database or entity, thus combining direct interaction with central coordination. Especially in approaches that allow user generation of content, e.g., user-designated landmarks in mobile gaming or ubiquitous interaction points, this combination could facilitate user contributions and involvement.

## 6.2 Outlook

Multiple direct interaction approaches become possible in our design, e.g., mobile gaming and ubiquitous networking. Especially, novel realizations of traditional approaches such as speech transmission and mobile computer vision for augmented reality show the potential of ubiquitous sensing of multimedia content and motivate further cross-community multimedia and networking research. In addition, we strive to explore the full potential of direct, ubiquitous interaction within everyday scenarios such as outdoors on campus or in the city and indoors to assess the characteristics of possible interaction encounters.

Last, sensing and interaction within the confined range of wireless transmissions may realize the notion of humanly graspable privacy-aware interaction through proximity and locality. This might then alleviate the difficulties of establishing and ensuring trust and authentication within the global dissemination scope of the Internet [14]. Similar, locality-based key establishment mechanisms [2] hint at the possibility of spontaneously establishing secure communication within the limits of each node's context and interaction range. While not the focus and scope of this paper, we feel that the growing awareness of privacy concerns and need for solutions motivates further work in this direction.

## Acknowledgements

This work was funded by the DFG Cluster of Excellence on Ultra High-Speed Mobile Information and Communication (UMIC). We are grateful to our colleague (cf. Figure 14(b)) for serving as a test subject in our face recognition evaluation.

## 7. REFERENCES

[1] J. Akhtman and L. Hanzo. Heterogeneous networking: An enabling paradigm for ubiquitous wireless communications [point of view]. *Proceedings of the IEEE*, 98(2), 2010.

[2] K. Argyraki, S. Diggavi, M. Duarte, C. Fragouli, M. Gatzianas, and P. Kostopoulos. Creating Secrets out of Erasures. In *MobiCom*, 2013.

[3] Bluetooth SIG. Bluetooth Service Discovery. `https://www.bluetooth.org/en-us/specification/assigned-numbers/service-discovery`.

[4] D. Brunelli, M. Maggiorotti, L. Benini, and F. L. Bellifemine. Analysis of audio streaming capability of zigbee networks. In *EWSN*, 2008.

[5] Computer Vision Group, ETH Zürich. Zurich Buildings Database. `http://www.vision.ee.ethz.ch/datasets/index.en.html`.

[6] GOOGLE. Google Glass. `http://www.google.com/glass/start/`.

[7] B. Han, P. Hui, V. Kumar, M. Marathe, J. Shao, and A. Srinivasan. Mobile data offloading through opportunistic communications and social participation. *Mobile Computing, IEEE Transactions on*, 2012.

[8] A. Henrysson and M. Ollila. Umar: Ubiquitous mobile augmented reality. In *MUM*, 2004.

[9] E. Huang, W. Hu, J. Crowcroft, and I. Wassell. Towards commercial mobile ad hoc network applications: a radio dispatch system. In *MobiHoc*, 2005.

[10] ITU. Recommendation G.114 : One-way transmission time. `http://www.itu.int/rec/T-REC-G.114`.

[11] S. Jakubczak and D. Katabi. A cross-layer design for scalable mobile video. In *MobiCom*, 2011.

[12] L. Keller, A. Le, B. Cici, H. Seferoglu, C. Fragouli, and A. Markopoulou. Microcast: cooperative video streaming on smartphones. In *MobiSys*, 2012.

[13] M. Kranz, P. Holleis, and A. Schmidt. Embedded interaction: Interacting with the internet of things. *Internet Computing, IEEE*, 14(2), 2010.

[14] M. Langheinrich. Privacy by design - principles of privacy-aware ubiquitous systems. In *UbiComp*, 2001.

[15] M. Lindeberg, S. Kristiansen, T. Plagemann, and V. Goebel. Challenges and techniques for video streaming over mobile ad hoc networks. *Multimedia Systems*, 17(1), 2011.

[16] E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng, and A. T. Campbell. Sensing meets mobile social networks: the design, implementation and evaluation of the cenceme application. In *SenSys*, 2008.

[17] E. Miluzzo, M. Papandrea, N. D. Lane, A. M. Sarroff, S. Giordano, and A. T. Campbell. Tapping into the vibe of the city using vibn, a continuous sensing application for smartphones. In *SCI*, 2011.

[18] NianticLabs@Google. Ingress. The game. `http://www.ingress.com`.

[19] NOKIA. LiveSight: immersive experiences you can act on. `http://conversations.nokia.com/2012/11/13/livesight-immersive-experiences-you-can-act-on/`.

[20] NOKIA. New nokia city lens for wp8 and wp7. `http://conversations.nokia.com/2012/11/28/new-nokia-city-lens-for-wp8-and-wp7/`.

[21] Opelt, Andreas and Pinz, Axel. Image database for object categorization experiments. `http://www.emt.tugraz.at/~pinz/data/GRAZ_01/`.

[22] J. Polastre, R. Szewczyk, and D. Culler. Telos: enabling ultra-low power wireless research. In *IPSN*, 2005.

[23] S. Pollin, I. Tan, B. Hodge, C. Chun, and A. Bahai. Harmful coexistence between 802.15.4 and 802.11: A measurement-based study. In *CrownCom*, 2008.

[24] J. Qiu, D. Chu, X. Meng, and T. Moscibroda. On the feasibility of real-time phone-to-phone 3d localization. In *SenSys*, 2011.

[25] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski. Orb: An efficient alternative to sift or surf. In *ICCV*, 2011.

[26] H. Wang, X. Bao, R. R. Choudhury, and S. Nelakuditi. Insight: recognizing humans without face recognition. In *HotMobile*, 2013.

[27] Wi-Fi Alliance. Wi-Fi Direct. `http://www.wi-fi.org/discover-and-learn/wi-fi-direct`.

[28] H. Wirtz, T. Zimmermann, M. Ceriotti, and K. Wehrle. CA-Fi: Ubiquitous Mobile Wireless Networking without 802.11 Overhead and Restrictions. 2013. Under submission, manuscript available to TPC chairs upon request.

[29] xiph open source community. Speex: A free codec for free speech. `http://www.speex.org/`.

[30] W. Yii, W. H. Li, and T. Drummond. Distributed visual processing for augmented reality. In *ISMAR*, 2012.

[31] Z. Zhang, D. Chu, X. Chen, and T. Moscibroda. Swordfight: enabling a new class of phone-to-phone action games on commodity phones. In *MobiSys*, 2012.

[32] Zhang, Xinyu and Shin, Kang G. Enabling coexistence of heterogeneous wireless systems: case for ZigBee and WiFi. In *MobiHoc*, 2011.