Seamless Resource Sharing in Wearable Networks by Application Function Virtualization

Harini Kolamunna, Kanchana Thilakarathna, Diego Perino, Dwight Makaroff, and Aruna Seneviratne

Abstract—The prevalence of smart wearable devices is increasing exponentially and we are witnessing a wide variety of fascinating new services that leverage the capabilities of these wearables. Wearables are truly changing the way mobile computing is deployed and mobile apps are being developed. It is possible to leverage the capabilities such as connectivity, processing, and sensing of wearable devices in an adaptive manner for efficient resource usage and information accuracy within the personal area network. We show that app developers are not yet taking advantage of these cross-device capabilities, however, instead using wearables as passive sensors or simple end displays to provide notifications to the user. We thus design Application Function Virtualization (AFV), an architecture enabling automated dynamic function virtualization and scheduling across devices in a personal area network, simplifying the development of the apps that are adaptive to context changes. AFV provides a simple set of APIs hiding complex architectural tasks from app developers whilst continuously monitoring the user, device, and network context, to enable the adaptive invocation of functions across devices. We show the feasibility of our design by implementing AFV on Android, and the benefits for the user in terms of resource efficiency, especially in saving energy consumption, and quality of experience with multiple use cases.

Index Terms—Smart wearable devices, wearable computing, energy utilization, context monitoring, function virtualization

1 Introduction

Smart wearable devices such as smartphones, tablets, smartwatches and fitness bands enable mobile users to form Personal Area Networks (PANs). Some of the devices in the PAN are capable of providing the same functionality. For instance, a fitness band, a smartwatch and a smartphone are each likely to have an accelerometer, a gyroscope, and a heart rate monitor. Similarly, a number of devices on the PAN may have direct Internet connectivity, providing multiple network interfaces. Furthermore, some wearable devices will have sufficient computing resources to perform functions such as data encoding, compression and encryption, while others may not.

Previous analysis of several popular wearable health and fitness tracking apps [1] shows that app developers tend not to leverage available resources on other devices. For example, the smartwatch pedometer will still use its own accelerometer when the battery level is low, despite an accelerometer being available on a fully-charged smartphone. The primary reason for relying on local resources only is the app developers reliance on the APIs provided by the target device.[1]

To harness the collective capabilities of PAN devices, developers have to implement each app individually to utilize the distributed device resources, managing the cost of running these functions in each device and communication costs explicitly. This requires developers to have a wider understanding of distributed systems and increases mobile app development complexity. Therefore, an architecture that takes user and device context into account, enabling app developers to utilize all PAN resources easily is needed.

In this paper, we present such an architecture that extends the concept of network function virtualization [2] to device functions. We show the proposed architecture's advantages to users and app developers and make the following contributions:

- Provide an architecture (AFV) that enables wearable/mobile app function virtualization for the development of adaptive wearable/mobile apps.
- Design AFV's inter-device and intra-device communication protocols to minimize the overhead added by the architecture and maximizes the advantages.
- Propose and evaluate a greedy heuristic algorithm for adaptive function allocation across devices considering the available resources and dynamic context of devices in the PAN.
- Demonstrate AFV's ability to adapt to context changes dynamically and demonstrate user and performance benefits of using AFV using a number of apps, and their usage.

- H. Kolamunna and A. Seneviratne are with the School of EE&T, University of New South Wales, Sydney, NSW 2052, Australia, and Data61-CSIRO, Eveleigh, NSW 2015, Australia. E-mail: {h.kolamunna, a.seneviratne}@unsw.edu.au.
- K. Thilakarathna is with the School of Information Technologies, University of Sydney, Sydney, NSW 2006, Australia, and Data61-CSIRO, Eveleigh, NSW 2015, Australia. E-mail: kanchana.thilakarathna@sydney.edu.au.
- D. Perino is with Telefonica Research, Barcelona 28050, Spain. E-mail: diego.perino@telefonica.com.
- D. Makaroff is with the University of Saskatchewan, Saskatoon, SK S7N 5A2, Canada. E-mail: makaroff@cs.usask.ca.

Manuscript received 25 Sept. 2017; revised 1 Apr. 2018; accepted 23 July 2018. Date of publication 1 Aug. 2018; date of current version 2 May 2019. (Corresponding author: Harini Dananjani Kolamunna.)
For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TMC.2018.2861861

1. These APIs abstract a large spectrum of functions (e.g., sensing, communication) actually implemented by the operating system or third party libraries and executed in the device where the app is running.

1536-1233 © 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

Fig. 1. An overview of an personal wearable network.

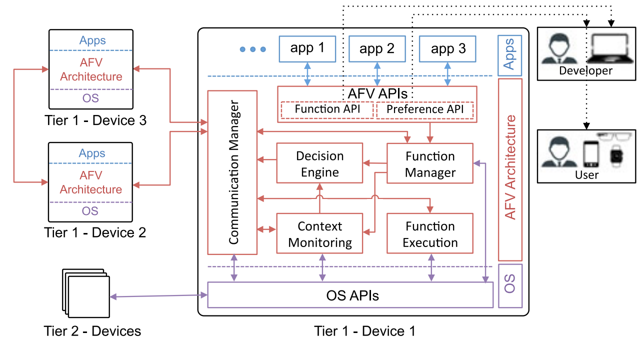

Fig. 2. Overview of the AFV and logical connectivity among devices.

The rest of this paper is organized as follows. We first present the AFV architecture, describing each module of AFV followed by the context-aware adaptive function allocation algorithm. Next, we show the realization of AFV with an Android implementation and experimental calibrations. Performance evaluation via simulation and experiments are presented in detail with the experimented use cases of AFV. Finally, we overview related work and complementary systems, and provide conclusions with future work.

2 AFV: Application Function Virtualization

Current PAN devices can be divided into two broad categories, which we refer to as Tier 1 and Tier 2, depicted in Fig. 1. Tier 1 devices, (e.g., smartglasses, smartwatches, smartphones), are relatively more resourceful than Tier 2 devices. Tier 2 devices, (e.g., smartshirts and bio-patches) simply carry out sensing functions. In contrast to Tier 2 devices, Tier 1 devices have the following additional features: a) availability of heterogeneous long-range network *connectivity* (WiFi, cellular), and b) ability to *process* sensed and received information.

Therefore, Tier 1 devices may be equipped with a rich set of sensors, multiple connectivity interfaces, storage and computing power to perform wide set of functions such as compression, encoding, rendering, intrusion detection, firewall filtering and encryption. Some of these resources are context dependent, e.g., WiFi connectivity will only be available in limited locations, GPS location will not be available indoors and devices may be disconnected if the battery is depleted.

Since Tier 2 devices provide complementary and specific functionality that may be duplicated on available Tier 1 devices, an app does not usually need sensors on all devices to be active simultaneously. Different sensors have differing output quality and resource requirements that make the selection of functionality for optimal user benefit at runtime a challenge. Thus if all the resources available on a PAN can be effectively utilized in an automated fashion, depending on the user/device context, it should be possible to significantly improve both the dimensions of user utility: functional requirements (precision, accuracy), and performance (energy consumption, latency).

As mentioned, app developers do not effectively utilize all the available resources in a PAN, when developing apps. We believe to facilitate the use of PAN wide resources, that it is necessary to provide methods for designers to seamlessly access the resources, taking in to account the context of use. We further believe this is achievable by virtualizing the commonly-used functions distributed across multiple devices on a PAN, and orchestrating the use of the functions depending on the context of use, and the system state, which we refer to as AFV. In so doing, it is possible to accomplish the goals of a) reducing the development effort for app developers; b) reducing the configuration burden of the users; and c) increasing the user quality of experience.

2.1 AFV Architecture

AFV is a collaborative platform, which runs on the Tier 1 devices and interacts with the Tier 2 devices in a PAN as shown in Fig. 1, makes available all potential resources to app developers and/or users seamlessly through APIs.

Fig. 2 provides a schematic view of the AFV architecture. To use AFV, apps register requests for the required functions via the *AFV APIs* explained in Section 2.2. Within AFV, app function registration requests are handled by the *Function Manager* module as described in Section 2.3. The *Context Monitoring* module periodically monitors device and user context as detailed in Section 2.4. One of the Tier 1 devices on a PAN is elected as the *Master Device*. This can either be done by the user or automatically based on a set of criteria such as the lowest ratio between current state of charge and energy usage (i.e., the Tier 1 device that would be least affected by the master tasks if selected).

The automated *Master Device* selection is done in the *Decision Engine* (described in Section 2.5) as a collective process of all the Tier 1 devices. The *Decision Engine* of the selected *Master Device* then determines the optimal mapping of each app function request to a function provided by a device on the PAN. All the other Tier 1 devices transfer context changes to the *Master Device* to be used by the *Decision Engine*. The *Communication Manager* maintains efficient Inter-device and Intra-device communication as detailed in Section 2.6. Finally, the *Function Execution* module performs function invocation on devices as described in Section 2.7. Fig. 3 shows the diagrammatic representation of the flow of events in AFV when a request is made, which is explained throughout this section.

2.2 AFV APIs

AFV provides two main types of APIs: *Function APIs* that are executed during run time and *Preference APIs* that are executed at app start-up.

2.2.1 Function APIs

For each supported function, AFV provides a specific API to the developer. As such, the *AFV APIs* augments the existing APIs provided by the operating system.

As an example, Fig. 4 describes the sensing function enhancements of the *AFV API*. The `onSensorChanged` and `unregisterListener` Android APIs are unchanged,

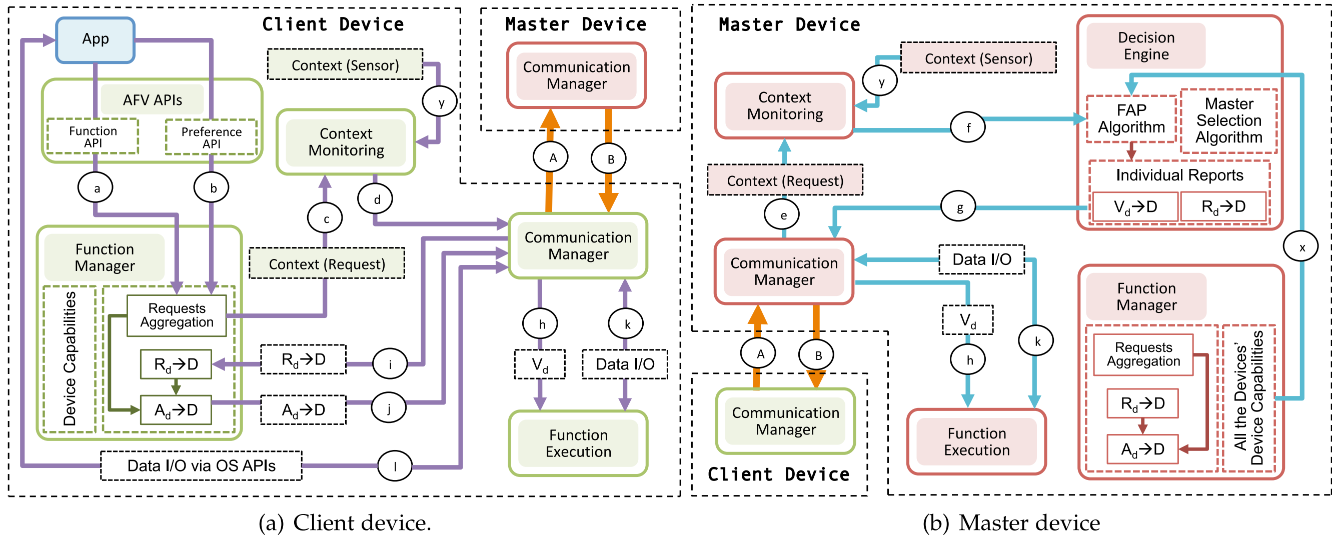

(a) Client device. (b) Master device

Fig. 3. Diagrammatic representation of AFV.

but are reimplemented as `AFVonSensorChanged` and `AFVunregisterListener` for consistency.

The `registerListener` API requires simple modifications. The `AFVregisterListener` call has two additional required parameters: `maxReportLatencyUs` is an input defining the data exchange frequency between apps and devices, and `precision` is an input to specify the required measurement accuracy with respect to absolute correctness, provided as a list of contexts and ranges (e.g., <running, <5, 10 percent>>, <walking, <5, 15 percent>>). In addition, the optional `mapping` parameter that provides a list of (context, device) pairs (e.g., <walking, smartphone>, <sitting, smartwatch>), is used to enable the developer to force the *Decision Engine* to select a particular device where possible.

All *Function APIs* are designed with a similar structure with minimal changes to the existing APIs to reduce the complexity of AFV for the app developers.

2.2.2 Preference APIs

There may be different preferred configurations for the user, app and the device itself when executing an app. These configurations are managed in AFV via the *Preference APIs*. AFV provides three types of *Preference APIs*: app (`setAppPrefs()`), user (`setUserPrefs()`) and device (`setDevicePrefs()`), all with the same structure. An example of `setAppPrefs()` is illustrated in Fig. 5. The configurations (app, user and device) could be conflicting. This is mitigated by certain configurations being mandated, or prevented. In the current implementation the following order of priority is used: Device, User, Application.

Consider a scenario where the developer (via `setAppPrefs()`) may specify a fitness tracking app to synchronize data with an Internet server, whenever data connectivity is available. However, the user may wish to override the app's settings (via `setUserPrefs()`) by configuring the app to synchronize data only when WiFi connectivity is available. For this example, the `preference` input in `setAppPrefs()` is <connectivity, 0>, where "0" denotes "any network", and the `preference` input in `setUserPrefs()` is <connectivity, 1>, where "1" denotes "WiFi network". However, the user *Preference API* is not directly exposed to the user; rather, it is leveraged by the developer to take user preferences as inputs. For instance, the UI of the app could provide an interface with radio buttons or drop-down lists to allow the user to select required or forbidden device/context mappings.

Additionally, each device may have configurations defined by the manufacturer. For example, iOS devices are recommended to operate below 35 °C.[2] These configurations are set by the developer via the `setDevicePrefs()`, so any user specified configurations beyond this is ignored.

All the preferences received via *Preference APIs* and functional requests received via *Function APIs* are first transferred to the *Function Manager*. (cf. ⓐ and ⓑ in Fig. 3a).

Android
```
abstract void onSensorChanged(SensorEvent event)
boolean registerListener(SensorEventListener
        listener, Sensor sensor, int sampling
        PeriodUs, int maxReportLatencyUs)

void unregisterListener(SensorEventListener
        listener, Sensor sensor)
```

AFV
```
abstract void AFVonSensorChanged(AFVSensorEvent
        event)
boolean AFVregisterListener
       (AFVSensorEventListener listener,
        AFVSensor sensor, int samplingPeriodUs,
        int maxReportLatencyUs, List<Context,
        <int,int> > precision, List<Context,
        Device> mapping)

void AFVunregisterListener(AFVSensorEventListe-
       -ner listener, AFVSensor sensor)
```

Fig. 4. Example of function APIs provided by AFV and by android.

2.3 Function Manager

The *Function Manager* is responsible for (i) keeping track of device capabilities in terms of supported functions and their associated costs, and (ii) managing the function registration

2. https://support.apple.com/en-au/HT201678

```
AFV

void setAppPrefs(AFVApplication appName,
       AFVDevice deviceName, List<Context,
       int> preference)
```

Fig. 5. Example of preference APIs provided by AFV.

requests and preferences received from AFV-enabled apps. The *Function Manager* in the *Master Device* additionally stores the supported functions and associated costs for all Tier 1 devices and their paired (passive) Tier 2 devices in the PAN. These stored information is transferred to the *Decision Engine* (cf. ⓧ in Fig. 3b).

A pre-defined list of supported functions of each device is provided with the AFV architecture. The associated cost of each function is composed of two main components; the cost of executing the function and the cost of exchanging inputs/outputs between the requested app and the function running device. Costs could be energy, monetary, latency, etc. The associated costs are either be provided in the list, which is provided with the AFV architecture, or obtained at the initialization using methods available in the devices (e.g., *getPower()* method in Android). Each time a new AFV-enabled Tier 1 device or a Tier 2 device joins the PAN, the *Function Manager* of the Tier 1 device first discovers the supported functions, and then it announces function availability and related costs to the *Master Device* (cf. Section 2.6).

The next key *Function Manager* role is the management of function requests/preferences. The preferences received via *Preference APIs* are considered for the requests where applicable, and apply bound conditions to the requests. For instance, if the user preferred to use WiFi connectivity, the Internet connectivity function requests from that particular app is bound with the condition to use WiFi connectivity only. The *Function Manager* aggregates multiple registration requests for the same function from different AFV-enabled apps, and only invokes the function on one of the devices where possible by considering the preferences (cf. Request Aggregation in Fig. 3a). Any change in the list of registration requests (R_d) is notified as a change of context and transferred to the *Context Monitoring* module of the device (cf. ⓒ in Fig. 3a). Eventually, this is transferred to the *Context Monitoring* module of the *Master Device* (cf. ⓓ, Ⓐ in Fig. 3a, ⓔ in Fig. 3b), and then triggers the *Decision Engine* of the *Master Device* (cf. ⓕ in Fig. 3b).

After each re-evaluation of the function allocation in the *Decision Engine* of the *Master Device*, the *Function Manager* receives the mapping between function registration requests $r \in R_d$ and the device $d \in D$ selected to execute the function ($R_d \mapsto D$) (cf. ⓖ, Ⓑ in in Fig. 3b and ⓘ in Fig. 3a). This is re-mapped to each function requesting AFV-enabled app $a \in A_d$ and the device $d \in D$ selected to execute the function ($A_d \mapsto D$). Finally, the *Function Manager* transfers the $A_d \mapsto D$ mapping to the *Communication Manger* in order to perform the data transfer from/to the AFV-enabled apps (cf. ⓙ in Fig. 3a).

2.4 Context Monitoring

We consider the context monitoring component as an additional virtual function, which runs on a device capable of receiving information from (i) the *Function Manager*, on the changes in registration requests, (ii) sensors, either directly or indirectly (cf. ⓨ in Fig. 3a), and transferring to the *Decision Engine* of the *Master Device*. The context retrieval function operates as needed and only reports changes relevant to the *Decision Engine*.

TABLE 1
Example Set of Context Information and Mapped System Objectives

System objective	Context
(i) Maximizing the Functional Quality	
Precision of fitness tracking	Moving status (i.e., walking or not)
Network throughput	Average link speed of the network
(ii) Energy Utilization	
Extend the battery uptime	Battery level, Default energy usage
(iii) Minimizing the Monetrary Cost	
Not exceeding the cap data level	Each device's connected network

It is possible for context monitoring to be carried out on different devices, depending on the context induced from the sensor values. Each sensed context is directly or indirectly mapped to a system objective function as shown in Table 1. Context monitoring is an essential element of AFV. Several implementations of context monitoring exist in the literature [3], [4]. These can be adapted for AFV. The context monitor evaluates the cost of obtaining the measures and selecting the appropriate mechanism/sensor with which to obtain this information. It is also possible that the context may be obtained from the device's operating system, if those features at that level are enabled [5].

Since AFV is intended to be used across multiple devices and multiple apps, contexts are represented in the common format described as `(name, value)` pairs. Context pairs can be defined per-app and per-device in the specific configuration files. The state of a particular context pair is expressed and stored as an enumerated type or string. The `value` field could be a) a threshold, b) ranges (e.g., moderate temperature could be 20-30 C), or c) a binary value (e.g., the device is either charging or discharging). The context `value` triggers a context change event and orchestrates appropriate flow of events to notify the *Master Device*.

2.5 Decision Engine

Initially, at the system bootstrap, the *Decision Engine* module of each device performs the master selection algorithm and elects the *Master Device*. Then the *Decision Engine* module on the *Master Device* executes the function allocation problem (FAP (cf. Section 3)), based on received context information from all client devices in the PAN. It determines the mapping of app function registration requests to actual function execution across devices.

The *Decision Engine* is triggered for a new assignment of functions by the *Context Monitoring* module on changes of context. First, the decision engine performs the preferred mappings and filters out infeasible function executions by considering the preferences of device, user and app. The preference information is transferred at the context changes. Then, each remaining function registration request is mapped to one of the feasible implementations as described in Section 3. As an example, the maximum acceptable delay for receiving a particular function's data by the app can be specified via the *Perference APIs*. This maximum acceptable delay value is compared with the transmission delay caused by the data transfer rate between devices (e.g., via

Bluetooth) and app data generation rate (e.g., sensing function sampling rate). The lower of these delays determines whether AFV will be used to discover and use the optimal function implementation.

Then, the *Master Device* creates individual messages for each device which has (i) the mapping between function registration requests $r \in R_d$ and the device $d \in D$ selected to execute the function ($R_d \mapsto D$), (ii) the mapping between function executions $v \in V_d$ in the device and each requesting device $d \in D$ ($V_d \mapsto D$). These messages are then transferred to the *Communication Manager* of the *Master Device* in order to communicate to all the devices in the PAN (cf. ⓖ in Fig. 3b).

2.5.1 Objective Functions

Although the overall objective of AFV is to maximize the user quality of experience, there can be specific objectives for individual users. We have categorized the potential user specific objectives into three groups: i) *Quality*, ii) *Energy*, and iii) *Monetary cost*. An example set of context information that is required by the *Decision Engine* to achieve a specific objective is summarized in Table 1. We have given the user and/or app developer the ability to configure the optimization preference (via the UI or *Preference APIs*); by default AFV is configured to optimize *Energy*. In the case of *Quality*, we assume the user prefers quality of service improvements, e.g., maximize the precision of sensing information or network throughput, over energy consumption of devices and monetary cost. Then, in the case of *Energy*, we presume the user prefers to keep all devices in the PAN active for the longest possible time compromising *Quality* and *Monetary cost*. Finally, if the user opts for *Monetary cost*, the *Decision Engine* makes external communication decisions based on cost/byte value provided by the user (via the AFVUI). Achieving one objective does not, however, guarantee the achievement of another objective. In order to support different objectives and function categories, the *Decision Engine* runs an instance of the function allocation problem (FAP) per objective-function category pair (cf. Section 3).

2.6 Communication Manager

This module manages all AFV communications. There are two types of communication modes in AFV: i) Inter-device communication and ii) Intra-device communication.

2.6.1 Inter-Device Communication

Communication among devices in the PAN is performed via Bluetooth or other similar low-powered wireless technologies. We define AFV specific message formats rather than using existing data structures (e.g., Java-defined, csv) in order to minimize the amount of data transferred. In addition, the *Communication Manager* aggregates messages and batches data transfers to further limit communication costs. Inter-device communication in AFV is performed in following situations.

(i) *During the system bootstrap or after any device joins/leaves the PAN.* All the devices broadcast their own capabilities (e.g., via `DataAPI` in Android). An *Initialization Message* is used in this phase. The main purpose of this message is to announce the device to the PAN along with its supported functions.

(ii) *Changes in the context.* In this phase, a *Context (Sensor) Message*, (i.e., changes in sensor data), or *Context (Request) Message*, (i.e., changes in the requests), is transmitted from client devices to master device when there is a change in any of the context information.

(iii) *Once the* Decision Engine *module on the* Master Device *executes the function allocation algorithm.* The *Communication Manager* of the *Master Device* notifies the other devices with the assignments using *Assignment Message*, which contains the $R_d \mapsto D$, and $V_d \mapsto D$ mappings.

(iv) *Data transfer from/to other devices in the PAN.* When the requested function is selected to run on a different device than the requesting device, the required data is transferred using a *Data Message*. In order to remove the additional tail energy after each message transmission, the *Communication Manager* is designed to aggregate data for different requests to a particular device. However, when a function requires any differences in the data transfer frequency (e.g., real-time data transfer requirements for microphone, camera), AFV can either identify the requirements by referring to the request type, (i.e., there are specified functions that are pre-defined as requiring specific data transfer rates), or the user/developer can specify via the *Preference APIs*.

2.6.2 Intra-Device Communication

Intra-device communication involves in data exchanges between the AFV and AFV-enabled apps. The *Data Message* format is used for these transfers. There are four methods available for intra-device communication: *broadcasting*, *sockets streams*, *content provider* (specifically in Android), and *service callbacks*. In *broadcasting*, only the apps registered with a given broadcast listener receive the broadcast data. The *socket streams*, *content provider* and *service callbacks* methods provide one-to-one data transmission instead, where data is directly exchanged between AFV and the targeted app.

Intuitively, the most suitable mechanism for intra-device communication depends on the requested function. For example, in the case of sensing functions where multiple different apps are requesting the same function, the use of *broadcasting* would be efficient. On the other hand, in case of data transfer towards the Internet where apps have different data to transfer, it should be more efficient if data is transferred directly to the app. We experimentally evaluate the energy efficiency of all these methods in Section 5.

2.7 Function Execution

This module is responsible for the invocation of functions on the devices selected by the *Decision Engine*. The *Communication Manager* forwards the required function invocations (cf. ⓗ in Fig. 3b if the selected function invocation device is the *Master Device*). Then, the *Function Execution* module leverages operating system APIs to execute functions. Essentially, *Function Execution* maps the function requests made by *AFV APIs* to operating system APIs and then invokes the operating system APIs accordingly. The data exchanges in between the AFV-enabled apps and AFV is performed via the *Communication Manager* (cf. ⓚ in Fig. 3b).

3 Context-Aware Function Allocation

The objective of the *context-aware function allocation* is to map each function registration request $r_{a,v,d} \in R$ to its chosen implementation, i.e., $R_v \mapsto D_v$, to optimize the total cost of

TABLE 2
Definitions of Notations

Symbol	Definition
D	Set of devices in PAN
A	Set of apps in PAN
V	Set of functions
R	Set of requests
$r_{a,v,d} \in R$	Request for virtual function $v \in V$, at device $d \in D$, for app $a \in A$
D_v	Set of devices providing implementations for a given virtual function v, where $D_v \subset D$
R_v	Set of requests for a given virtual function v, where $R_v \subset R$
F_v	Set of function implementation costs available for a given virtual function v
C_r	Set of communication costs between a given requests and an implementation on a device
$f_{v,d}$	Implementation function cost of v in device d, $f_{v,d} \in F_v$
$c_{r,d}$	Communication costs between a given request and an implementation on device d, where $c_{r,d} \in C_r$
$(m_{r,d})$,$(x_{r,d})$, (y_d)	Binary variables
X, Y	Sets of solutions

executing the all requested functions. All the notations used in this section are detailed in Table 2.

3.1 Function Costs

The function costs are related to the usability objective of the system (e.g., monetary, quality and/or energy), which can be defined either by the app developer or the user. For each objective, there are two types of costs associated with each function request and its implementations; 1) communication costs and 2) implementation costs. If the objective is to optimize the monetary costs, internal communication (e.g., Bluetooth), can be considered as zero. On the other hand, if the objective is to optimize the energy consumption of the devices, communication is not negligible. Usually, local mapping incurs zero cost. For the same function, the implementation cost can be different for multiple devices, for example, the energy cost of activating the WiFi network interface compared to the total battery capacity on the smartphone is lower than on the smartwatch.

3.2 Problem Formulation

We first define a binary variable $m_{r,d}$ where $m_{r,d} = 1$, if function registration request $r_{a,v,d} \in R_v$ can be mapped to implementation on device $d \in D_v$, and $m_{r,d} = 0$ otherwise, depending on the current context of the user and the device. For instance, even if the GPS sensor is implemented on the device $d_1 \in D_v$, it may not be able to map d_1 with any request if the current remaining battery capacity on d_1 is below the threshold. If no function implementation is available for a particular function registration request, we remove that function from the problem formulation. That makes for all considered functions $\sum_{\forall d \in D_v} m_{r,d} = 1; \forall r \in R_v$. Given the set function registrations R_v and function implementations D_v and the associated costs $f_{v,d}$ as input, the optimal FUNCTION ALLOCATION PROBLEM (FAP) can be formulated as follows:

$$\text{Minimize}\left(\sum_{d \in D_v} y_d \cdot f_{v,d} + \sum_{r \in R_v}\sum_{d \in D_v} x_{r,d} \cdot c_{r,d}\right), \quad (1)$$

subject to

$$\begin{array}{lll} 1. & \sum_{d \in D_v} x_{r,d} = 1; & \forall r \in R_v \\ 2. & m_{r,d} \geq x_{r,d}; & \forall r \in R_v, \forall d \in D_v \\ 3. & y_d \geq x_{r,d}; & \forall r \in R_v, \forall d \in D_v \\ 4. & y_d, x_{r,d} \in \{0,1\}; & \forall r \in R_v, \forall d \in D_v. \end{array}$$

The sets of $x_{r,d} \in X$ and $y_d \in Y$ would be the solution of the FAP. $x_{r,d} = 1$ if the function registration request $r \in R_v$ is assigned to the device $d \in D_v$ and $y_d = 1$ if the device d is required to be activated to satisfy certain requests. Only mappable implementations will be assigned and each function registration request will be mapped to an implementation.

3.3 Solution to Function Allocation

When $m_{r,d}$ is given $\forall r \in R_v, \forall d \in D_v$, it is trivial to show that FAP is equivalent to the UNCAPACITATED FACILITY LOCATION (UFL) problem where every function implementation D_v is a facility with $f_{v,d}$ facility opening cost and every function registration request R_v corresponds to a customer associated with $c_{r,d}$ service cost. It immediately follows that FAP is also an *NP-Hard* problem. However, there are a number of approximation algorithms for the well-studied UFL problem. We build on the approximation algorithm proposed by Williamson and Shmoys [6] to take into account the use of valid mappings ($m_{r,d}$) after context aware constraints. The iterative greedy solution to FAP is described in Algorithm 1.

Algorithm 1. FAP(R_v, D_v, m, f, c)

1. $S \leftarrow R_v$
2. $X \leftarrow \emptyset$
3. **while** $S \neq \emptyset$ **do**
4. Select $v \in D_v$ and $P \subseteq S$ s.t. $\forall p \in P : m_{r,d} = 1$ that minimize $\frac{f_{v,d} + \sum_{p \in P} c_{r,d}}{|P|}$
5. $S \leftarrow S - P$; $f_{v,d} = 0$
6. $(R_v \mapsto D_v) \leftarrow (R_v \mapsto D_v) + (P \mapsto v)$
7. **return** assignment $\sigma : R_v \mapsto D_v$

The algorithm iteratively selects a function implementation and the valid registrations associated to it. Assigned registrations are then removed from the problem and the implementation cost set to 0. At each iteration, an implementation is selected to minimize the total cost of function registrations that will be associated to the implementation. The algorithm can be efficiently realized by maintaining the list of registrations not yet satisfied for each implementation in increasing cost order. FAP is performed for each type of function separately, and aggregate the solutions to get the final solution. This further increases the efficiency of FAP.

4 REALIZATION OF AFV FRAMEWORK

4.1 Implementation

Our prototype realizes AFV as a library that can be compiled into an app and as a stand-alone user-level app. In Section 5, we use AFV lib and AFV app to evaluate the practical feasibility of AFV and user benefits with use cases.

For simplicity and without loss of generality, it is assumed that all Tier 1 devices in the PAN run the same OS (e.g., Android). The library, AFV lib, provides access to the *AFV APIs* for the developer once linked into the app binary

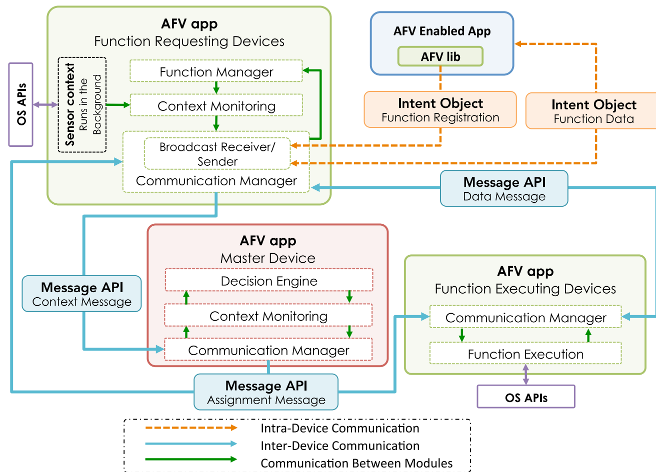

Fig. 6. Message flow in the implementation.

file. To support AFV services to multiple apps in parallel at user-level, the other components of AFV are implemented in a standalone app, AFV app. At the time of installation of an AFV enabled app, it checks whether the AFV app is already installed. If not, it initiates the installation of the AFV app.

The communication between AFV-enabled devices is implemented using `MessageAPI` and `DataAPI` as Android `Services`. Context monitoring is also implemented as an Android `Service`. A *Master Device* is selected among the devices in the network and it runs FAP algorithm described in Section 3 to check for optimal function placement.

Device arrival/dxeparture is monitored using the `onCapabilityChanged` method. When an AFV-enabled device joins/leaves the PAN, all AFV-enabled devices send *Initialization Messages* via `DataAPI` to all the connected devices. In addition, *Context Monitoring* runs in the background and collects information such as battery level, charging status, motion status, connected network, and link speed. When a context of a device changes, it reports the change to the *Master Device* by sending a *Context (Sensor) Message* via `MessageAPI`. This context change will trigger the *Decision Engine* of the *Master Device* to make new decisions.

Also, as shown in Fig. 6, once an app registers a function (e.g., `AFVregisterListener(this, Sensor.TYPE_ACCELEROMETER), AFVHTTPGet`), AFV lib sends a broadcast with an `Intent` about the function registration to AFV app. The `BroadcastReceiver` in the AFV app transfers the request to the *Function Manager*. Then, it is identified as a context change and reports to the *Master Device* by sending a *Context (Request) Message* via `MessageAPI`. Then *Master Device*'s *Decision Engine* checks for the optimal placement of the function.

The *Decision Engine* implements the FAP algorithm by using efficient ordered data structures (i.e., `TreeMultimaps` and `ArrayLists`). *Function Execution* registers the function using Android APIs and returns data in a *self-defined data format* to the *Communication Manager* (*Data Message*). The *Communication Manager* aggregates all the *Data Messages* per device and sends them via `MessageAPI`. Once the *Data Message* is received by a device, the *Communication Manager* broadcasts the data stream. Each app that has a registered `listener` for the function will receive the data stream.

4.2 Experimental Calibrations

4.2.1 Energy Costs of Functions

For the FAP algorithm to allocate functions optimally, the *Function Manager* should contain a list of supported functions and their associated costs. We measured energy consumption of several app functions for sensing, communication,[3] and processing. We use a smartphone running Android 6.0 and LG Watch Urbane running Android 5.1.1 with 2300 mAh and 410 mAh battery capacities respectively for all experiments. The functions are reported in Table 3.

The energy consumption for each function is obtained with a Monsoon power monitor[4] directly connected to each device via USB. Energy usage is obtained by integrating the instantaneous power values calculated using current and voltage measurements from the USB interface sampled at 0.2 ms time intervals. The experiment energy usage is computed by deducting the fixed energy of the background processes from the total energy consumed. Sensing energy is measured for multiple sampling frequencies that are offered by Android by default.[5] These energy cost values are used in the experimental validation of AFV architecture and to show user benefits.

4.2.2 Intra-Device Communication Modes

As mentioned in Section 2.6, there are four potential modes for intra-device communication. We experimentally evaluate the energy efficiency of four methods (*broadcasting*, *socket streams*, *content provider*, *service call-backs*) in the case of AFV message exchanges. Fig. 7a shows the results for different intra-device communication methods in transferring 2.5 KB of data from AFV to an app. The measurements are done in an Android smartphone. As expected, intra-device communication has much lower energy usage than inter-device communication (cf. Table 3). In this particular case, smartphone inter-device communication consumes ~23 mJ without tail energy (~650 mJ with tail energy) and intra-device communication consumes ~1 mJ-6 mJ.

Moreover, we observe that *broadcasting* requires the minimum energy for transactions even in case of a single app. Therefore, we further experiment with the case where up to eight external apps request the same data from the AFV architecture. Fig. 7b illustrates that energy consumption of *broadcasting* to eight apps is still lower than any other method of intra-app communication for a single one app as shown in Fig. 7a. We thus select *broadcasting* as the intra-device communication mode in AFV.

5 Performance Evaluation

We first evaluate the performance of the FAP algorithm with data driven simulations. Then, we present the simulation and experimental results of how AFV prolongs the system uptime. Finally, we show the benefits of AFV, and validating the design objectives based on the experimental use cases.

5.1 Evaluation of the FAP Algorithm

5.1.1 Evaluation Methodology

We developed a custom simulator to analyze the effectiveness of the *Decision Engine*, in particular, the FAP

3. Only WiFi *receive* measured, *transmit* is between 20 and 30 percent higher [7].
4. https://www.msoon.com/LabEquipment/PowerMonitor/
5. http://developer.android.com/reference/android/hardware/SensorManager.html

TABLE 3
Energy Cost Associated to Each Function

Function	Energy cost							
	Smartphone				Smartwatch			
Sensing	Speed [mJ/s]							
	NORMAL	UI	GAME	FASTEST	NORMAL	UI	GAME	FASTEST
Accelerometer	5.01	13.28	34.46	77.71	9.52	24.74	57.61	168.4
Gyroscope	11.71	20.33	36.44	80.15	16.23	33.34	60.44	182
Magnetometer	8.12	15.45	28.46	28.28	17.04	30.21	57.82	79.73
Connectivity	Per Byte [mJ/B]	High Power Idle [mJ]		Low Power Idle [mJ]	Per Byte [mJ/B]	High Power Idle [mJ]		Low Power Idle [mJ]
Bluetooth	0.0095	305		300	0.0024	126.07		64.23
WiFi	0.0005	66		N/A	0.0004	50		N/A
Processing	Per Byte [mJ/B]				Per Byte [mJ/B]			
Compression	0.01				0.0004			
Encoding	0.00026				0.00025			

algorithm. We compared AFV function allocation against three strategies:

- MANUAL: User assignment of functions in a static manner, e.g., MyFitnessCompanion [1].
- ALL: Running functions on all available devices in parallel, which is one of the common strategies in today's wearable apps, e.g., UP [1].
- OPTIMAL: Function allocation, using the optimization problem solver Gurobi.[6]

We assume that costs of executing a function on devices is normally distributed, with a standard deviation $\sigma = 0.1 \times \mu$ where μ is the average value. We change μ to obtain multiple cost values to evaluate the performance of FAP.

5.1.2 Efficiency and Robustness of the FAP Algorithm

Fig. 8a shows the cost reduction obtained when using the FAP algorithm with respect to MANUAL, ALL and OPTIMAL as a function of the ratio of function implementation cost (F_v) to communication cost (C_r). Communication costs are incurred for any transmission of an individual sensor stream to the device executing the app.

We consider 5 active wearable devices in a PAN. Intuitively, if the communication cost is too high, it is more efficient to execute the function on each device, resulting in parallel apps with no coordination. This is reflected in the region where $F_v/C_r < 1$. Under these conditions ALL performs as well as OPTIMAL and FAP. However, FAP significantly reduces the cost compared to MANUAL selection. As F_v/C_r increases, the significance of the communication cost decreases. Thus, executing a function in all devices becomes inefficient as there is potentially a device with a very low relative function execution cost. Since there is a 1/5 chance of selecting the right device, MANUAL performs comparatively well with a high standard deviation. FAP performs equally well (error is less than 1 percent) compared to OPTIMAL, irrespective of the F_v/C_r value.

Fig. 8b shows that FAP increases its cost savings compared to both MANUAL and ALL along with the number of functions when $F_v/C_r = 1$. Furthermore, FAP accuracy does not vary significantly compared to OPTIMAL (error is about 2-3 percent). Overall, Fig. 8 shows that the FAP algorithm is often able to map the function registration requests to the optimal device for executing the function providing significant cost savings.

5.2 Prolonging System Uptime

We analyze AFV effectiveness by considering system uptime (i.e., time until at least one device drains out its battery) as an example of the quality metric. First, we evaluate the system uptime varying the power status of devices with simulations. Then, we augment simulation results by conducting experiments with real-devices. We consider a smartphone and a smartwatch for both simulations and experiments.

5.2.1 Evaluation with Simulations

Uptime of a device depends on its remaining battery percentage (i.e., State of Charge (SoC)) and current energy usage. To simulate typical user behaviour, we assume the smartphone battery would completely drain in two days linearly and the smartwatch would last only one day. We consider the "sensing accelerometer in FASTEST speed" function and 60-second data synchronization frequency: the *Decision engine* makes decisions to maximize system uptime. We use measurements in Table 3 to derive energy consumption for the functions. As an example, for sense only on smartphone (Accelerometer FASTEST speed) and data synchronization frequency of one minute (70 KB of data), we can get the energy consumption per minute from Table 3 as $f + c$=(77.71*60)+((0.0095*70000)+305+300)=5932mJ.

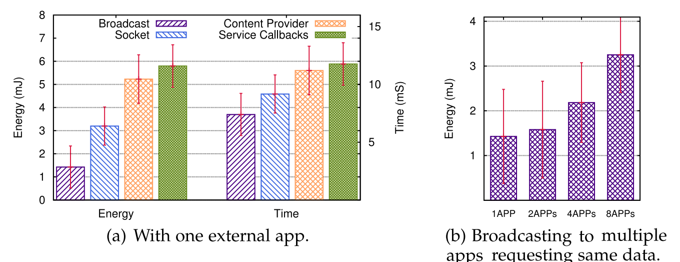

(a) With one external app.
(b) Broadcasting to multiple apps requesting same data.

Fig. 7. Impact of intra-device communication modes.

6. https://www.gurobi.com

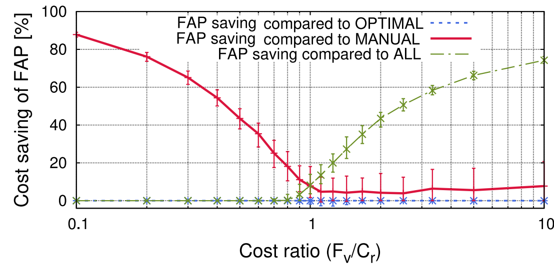

(a) The impact of cost ratio F_v/C_r.

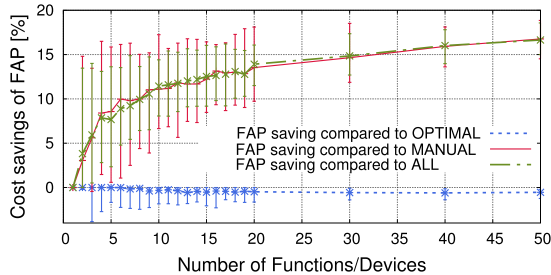

(b) The impact of no. of functions, $F_v/C_r = 1$.

Fig. 8. Efficiency, accuracy, and robustness of FAP algorithm.

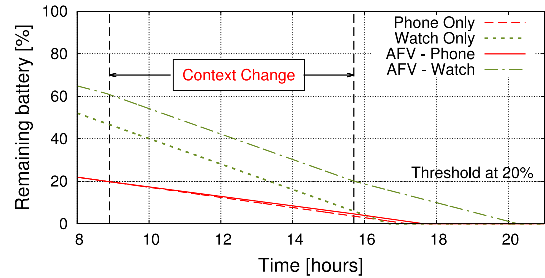

(a) Battery drain profiles

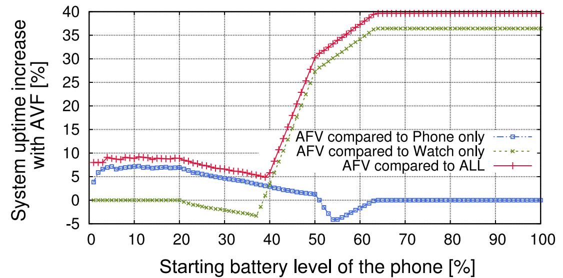

(b) Percentage increase in system uptime.

Fig. 9. The impact of AFV on system uptime.

Fig. 9a illustrates the battery drain profile for MANUAL selection, i.e., sense only on the smartphone or on the smartwatch, and when AFV is running. Due to a lower relative impact on the smartphone, AFV selects the smartphone as the sensing device if the smartphone has sufficient SoC. However, if the smartphone's SoC drops below 20 percent (context change), the *Context Monitor* triggers the *Decision Engine* and sensing switches to the smartwatch if the smartwatch has sufficient SoC (Fig. 9a). To show this context change, we consider the following initial conditions: smartphone-45 percent SoC, smartwatch 100 percent SoC. The smartwatch uptime increases by approximately 2 hours compared to sensing on the smartwatch. The gain for the smartphone is approximately 1/2 hour compared to only sensing on the smartphone.

Since uptime gain is dependent on the initial SoC of devices, in Fig. 9b we change the initial smartphone SoC. If the smartphone's remaining SoC is greater than 60 percent at the beginning, AFV increases the system uptime between 35-40 percent compared to sensing on the smartwatch and on both devices. Due to sufficient battery capacity on the smartphone, AFV selects the smartphone most of the time. As a result, AFV does not increase the uptime compared to sensing only on the smartphone. AFV may marginally reduce the system uptime when the SoC of one or more devices drops below the threshold. To minimize the app energy consumption, while respecting user preferences, the *Decision Engine* selects the only available device or the most energy efficient when both are under the threshold, although this solution may reduce system uptime. This can be observed when initial smartphone SoC is approximately 35 and 55 percent.

5.2.2 Evaluation with Experiments

Next, we quantify the energy consumption of the devices with and without AFV experimentally. We installed AFV on an Android smartphone and a smartwatch. Without AFV, we use counterpart apps that are installed on both devices. We consider the "sensing accelerometer in NORMAL speed" function and 60-second data synchronization frequency. Most current apps select both devices to perform a certain functionality and then exchange data [1]. Therefore, we selected ALL function allocation strategy.

Fig. 10 considers the energy consumption of the function as well as the overheads of AFV such as, communication, context monitoring and running the optimization algorithm. Fig. 10a shows the measured energy consumption for each type of message passing in AFV (cf. Section 2.6). The energy requirement for group formation (*Initialization*) is $(0.6 * (n-1) + 1.8)$J that is much lower compared to the group formation energy in Hemminski et al. [8]. Fig. 10b shows the reduction of battery SoC in smartphone and smartwatch. Using AFV achieves lower energy usage by approximately 3 times for one function request, despite the additional energy consumption of AFV (e.g., *Initialization*, *Context Monitoring*, running the FAP algorithm). In a practical scenario of having 10 function requests within the PAN, the FAP algorithm consumes ~100 mW of power for less than 10 mS of time. Moreover, the battery SoC decreases much faster as the number of functions increases, especially without AFV. Thus more energy is saved with AFV when the number of function requests increases.

5.3 Experimental Validation of AFV-Enabled PAN

We now present experimental results to show quantitative benefits of AFV. As shown in Fig. 11, we implemented an AFV-enabled PAN consisting of 4 devices (3 Tier 1 devices and a Tier 2 device). The Tier 1 devices (smartphone, smartwatch and smartglasses[7]) are connected to each other via Bluetooth and each have Internet connectivity. The Tier 2 device (smartshirt[8]) consists with three different sensor types and paired with the smartphone. We implemented AFV in all three Tier 1 devices. Although the AFV architecture is not installed in the smartshirt (as it is a Tier 2 device), it is considered as a remote sensor that is available in the paired

7. https://developers.google.com/glass/
8. https://www.hexoskin.com/

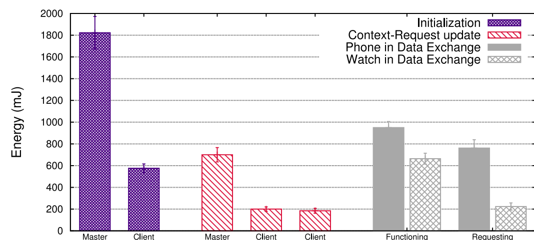

(a) Energy consumption of different message passing phases.

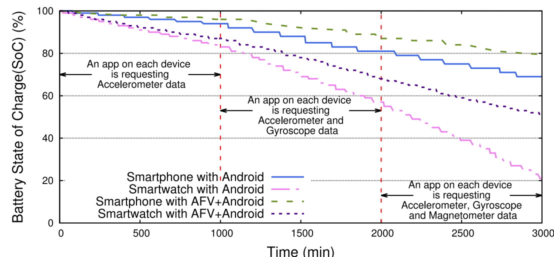

(b) Energy consumption with and without AFV when each device is requesting for different functionalities.

Fig. 10. Energy consumption of the system.

smartphone (Tier 1 device). The smartphone is responsible for pulling data from its remote sensors and feeding it to the AFV via REST APIs. We developed an AFV-enabled fitness tracking app requiring accelerometer and heart rate information to be uploaded to Internet servers that is similar to the previously identified current popular health and fitness apps. The app was installed in the smartphone, smartwatch and smartglasses.

We investigated five scenarios, to investigate the benefit of using AFV. These experiments evaluated three main objectives described in Table 1. For the first objective of achieving the maximum functional quality, we conducted two experiments, one to achieve the best information quality, and the other to achieve the maximum network throughput. To minimize energy utilization, we examined how the usage of AFV extends the device uptime. Finally, we examined the case of minimizing the monetary cost of data usage.

5.3.1 Maximizing the Functional Quality

(1) Maximizing the Precision of Fitness/Health Tracking.

(a) Requesting Accelerometer Data. The AFV-enabled app requires accelerometer data for fitness tracking. The user is wearing smartglasses and the smartwatch, and has a smartphone nearby while standing and exercising. At first, the user is doing head stretching exercises and then moves to body stretching exercises. After a while, the user starts walking, carrying their smartphone in their pocket. We consider that the smartphone provides the best quality information when user is walking, the smartwatch provides the best quality information when user is doing body stretching activities, and the smartglasses provide the best quality information when doing head stretching exercises. This rule is used in the *Decision Engine* in order to feed the app with highest quality data.

At first, while the user is doing head stretching exercise, as there is no walking detected by the smartphone on the table, and no activities are detected by the smartwatch, the smartglass performs the accelerator function and feeds

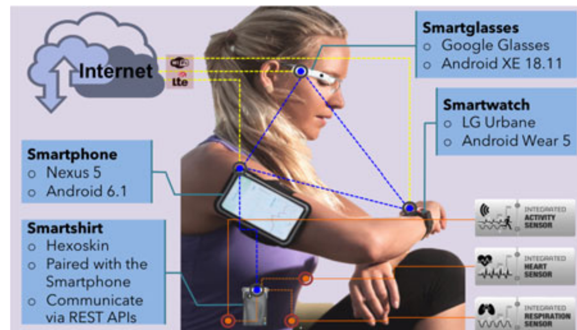

Fig. 11. Experimental setup.

data to the AFV-enabled app. We consider the detection of an activity if the accelerometer (in NORMAL speed) readings are above the specified threshold (we set a threshold to discard the floating values), and also, in order to minimize the erroneous activity detections, the detected activity needs to be continued for a given period of time. This also avoids unnecessary functionality movements between devices for this context, AFV triggers a function placement change only if the new activity continues for a given period of time, in this experiment, 10 seconds. During this time, the previously selected device continues to feed data to the AFV-enabled app. When the user starts doing body stretching exercises, the accelerometer sensing function is moved to the smartwatch and feeds data to the AFV-enabled app. When the smartphone detects that the user is walking, the sensing function moves from the smartwatch to the smartphone. Fig. 12a shows the accelerometer data that is received by the AFV-enabled app installed in the smartphone.

(b) Requesting Heart Rate Data. Next, the AFV-enabled app requests for heart rate (HR) data. Assume that at first, the user has the smartphone and the smartwatch. When the user dons the smartshirt, it is paired with the smartphone.

Assume that the best heart rate measurements are given by the smartshirt as it is specifically designed for sensing. In the absence of the smartshirt, the data is provided by the smartwatch. Therefore, the *Decision Engine* will select the smartshirt whenever it is available. The context change of the availability of the smartshirt via the smartphone triggers the AFV architecture and the *Decision Engine* selects the smartshirt via the smartphone to feed the AFV-enabled app.

Fig. 12b shows the heart rate data received by the AFV-enabled smartphone app every 10 seconds. It illustrates that the smartshirt is feeding stable and accurate HR data when the user is doing the same activity. In this particular case, the smartshirt's data is available for third party app development via the cloud. Therefore, the smartphone is connected to the cloud and retrieves the smartshirt's real-time data and feeds it to AFV. However, managing resources in Tier 2 devices depends on the accessibility provided by the device manufacturers.

(2) Maximizing the Network Throughput. In this use case, AFV virtualizes the Internet connectivity function in order to maximize the network throughput in a heterogeneous environment.

Assume that an AFV-enabled app on the smartphone is configured to upload sensor data from the smartphone to an external server periodically (i.e., per second). We created two WiFi networks with different throughputs to emulate the

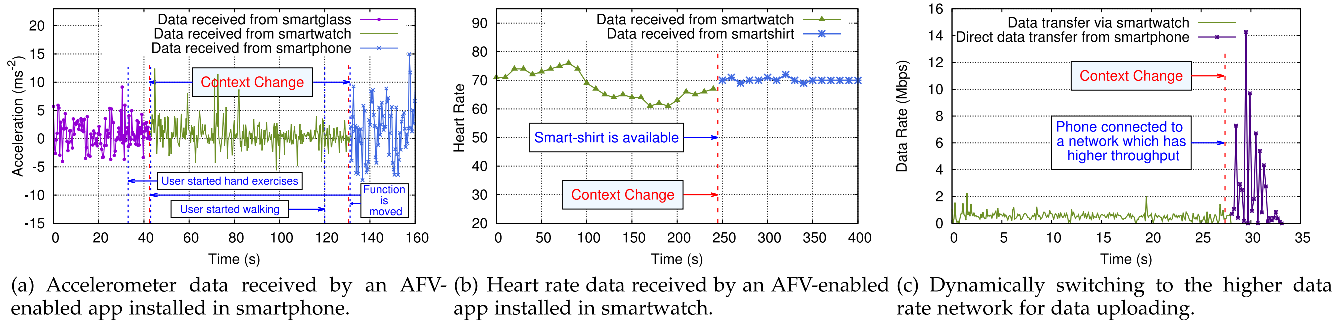

(a) Accelerometer data received by an AFV-enabled app installed in smartphone.

(b) Heart rate data received by an AFV-enabled app installed in smartwatch.

(c) Dynamically switching to the higher data rate network for data uploading.

Fig. 12. Maximizing the functional quality with AFV.

heterogenous network. At first, the smartwatch has Internet connectivity, but not the smartphone. After a while, we enable another higher speed network, to which the smartphone is connected. The device's connectivity to a new network triggers a context change which invokes the *Decision Engine* to select the higher throughput network to upload the file.

At first, when the smartphone does not have direct Internet connectivity, the data from the smartphone is uploaded to the Internet servers by relaying through the smartwatch. After a while, when direct Internet connectivity for the smartphone becomes available, the AFV-enabled app on the smartphone automatically suspends the data transfer via the smartwatch and starts transferring directly to the Internet. The achieved throughput is measured at the access points by using a network analyzer (Wireshark[9]). Fig. 12c depicts the throughput at the access points for data uploads before and after context change.

5.3.2 Extending the Device's Uptime

We use the same AFV-enabled app used previously. In addition, we developed two other apps for the smartphone and smartwatch that have the same functionality but do not use AFV (default Android). In both cases, the app installed in the smartphone requests accelerometer data, and the user preferred the app to get accelerometer data from the smartwatch and transmitted to the smartphone once per minute.

In the case of AFV, when the SoC reaches the threshold (i.e., 20 percent), it triggers a context change. Fig. 13a illustrates the devices' power profiles during the context change. Initially in this experiment, the smartwatch is sensing at normal speed and sending data to the smartphone once in a minute. When the context is changed at $t = 15$ seconds, the smartwatch broadcasts the context change to the *Master Device*, which triggers the *Decision Engine* on the *Master Device* to selects smartphone with a higher SoC for sensing, and informs devices. The smartwatch then stops sensing and the smartphone takes over the sensing function.

The high power peaks of all devices after $t = 15$ seconds is due to the messages received and transmitted by each device, which is followed by high power idle states. The high power idle state is longer for the smartphone (until $t = 26$ seconds) compared to the smartwatch (until $t = 20$ seconds). Fig. 13a shows that the delay of system adaptation to context changes is less than one second as the smartphone starts sensing even before $t = 16$ seconds.

In the default case, the smartwatch keeps running its accelerometer and transfers data to the smartphone until its SoC reaches 0 percent. Fig. 13b shows results for the increased longevity of the smartwatch battery when using AFV. The smartwatch uptime is increased by 5 hours due to the sensing function offloading.

9. https://www.wireshark.org/

5.3.3 Minimizing the Monetary Cost of Data Usage

We virtualize the Internet connectivity function in order to minimize the monetary cost of data transfer. The monetary cost for each data plan is pre-configured at system bootstrap and can be changed at any time via the AFV user interface. Assume that the AFV-enabled app needs to upload files to the Internet. Also assume at first, only the smartphone has Internet connectivity, but the smartphone's data plan has exceeded the available data cap and excess data costs $0.10/MB. After a while, the smartwatch's Internet connectivity, which has not exceeded the data cap, becomes available. Since this data plan has not exceeded the data limit, this plan costs $0.00/MB.

The availability of the additional network connection triggers a context change and the decision is made to use the low-cost network. AFV notifies the device connected to network with lower cost to take over the connectivity function. Fig. 14 shows the data usage of the smartphone and smartwatch before and after the context change.

6 Related Work

6.1 Context Monitoring

There has been substantial work on context awareness and sensing for mobile apps. Always-on sensing can quickly drain battery resources [9], yet continuous context monitoring is essential for proper response to context changes [4]. This suggests a distinction between always-on and continuous that does not degrade app adaptivity nor battery life. These trade-offs are explored in several research projects.

The most comprehensive sensing framework is SeeMon [4]. Their approach leverages the relationship between sensor values and higher level "context" states to minimize the number of sensors and their associated energy costs while continuously recognizing context changes. Another approach to sensing is to use a low-powered sensor processor to save energy. MobileHub [9] provides a framework that determined optimized alerts and submission of sensor data that reduce energy without affecting app semantics. Our context has a limited number of sensors and a small number of devices capable of performing context recognition, therefore, we have simplified the evaluation of sensor readings with a call-back mechanism for each activated sensor to inform the smart device regarding changes to the value of interest.

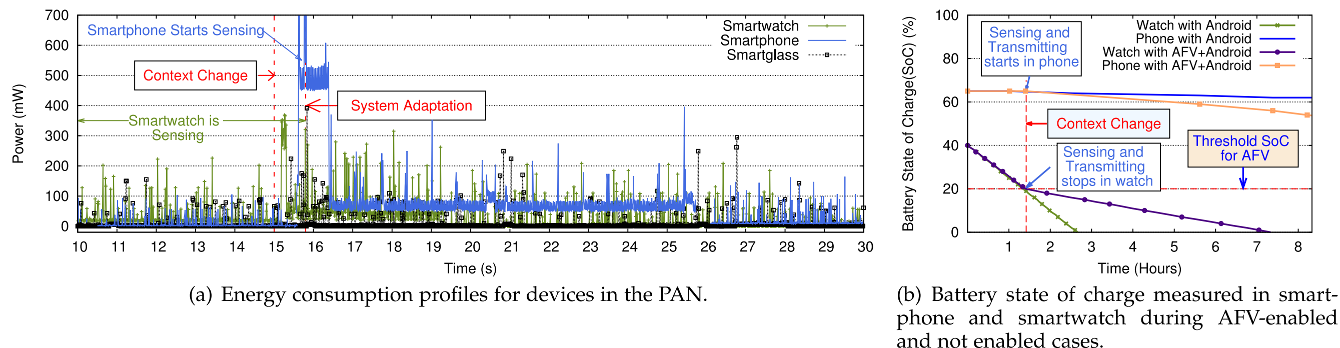

(a) Energy consumption profiles for devices in the PAN.

(b) Battery state of charge measured in smartphone and smartwatch during AFV-enabled and not enabled cases.

Fig. 13. Extending the device uptime with AFV.

6.2 Single-Device Resource Utilization

Adaptive system/framework for the context changes is a key concept in resources utilization. Adaptive systems designs have been in existence for nearly 20 years [10]. Early work provided context based systems development [11], prototype implementations [12], [13], programming language support for existing apps [14], and architectures for system design [15], [16]. Applications on commercially available devices have only recently been deployed, due to the challenges of battery and device form, among other issues [17], [18], [19]. For the purpose of resource utilization, Martins [20] aims to tune the background apps in Android selectively to improve the battery lifetime. They use an OS mechanism to control the frequency of handling background tasks.

CAreDroid [21] is a framework in which to design Android apps to select the most appropriate functions to run for a given app on a single device depending on the context. It takes care of context-monitoring, adaptation decisions and allows the developer to focus on app logic only. Their work provided the inspiration for our focus on distributed system apps for wearable computer network apps. AFV differs in that it provides seamless function placement across devices of a PAN and function sharing across apps. Our optimization engine runs as a lightweight separate process on Tier 1 devices and the adaptation selects which functions from which devices are active at any point in time and what communication strategy will be deployed.

Senergy [22] utilizes the sensing functionality in a way that it reduces the energy usage. This work does not require programmer intervention via the Latency, Accuracy Battery (LAB) abstraction. These 3 components are the main considerations in our framework as well, as they provide a meaningful set of tradeoffs for the user and the developer. The authors develop classifiers to infer context in sensing apps, while we use a simpler sensing strategy, but provide adaptation to achieve app goals.

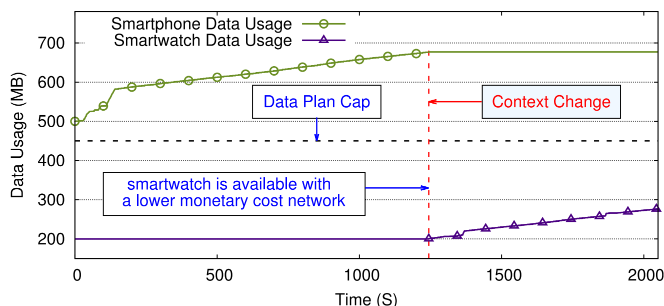

Fig. 14. Dynamically switching to the less costly network for data uploading.

All of the above specified projects consider a single device for the resource utilization. In contrast to this work, AFV targets the utilization of resources in a network of multiple smart-wearable devices in order to achieve the user-selected objective at runtime.

6.3 Multi-Device Resource Utilization

Mechanisms to utilize the resources from multiple devices has been receiving increasing attention, and the adaptive framework designs enable developers to create tasks for multiple devices [23], [24]. There are several studies such as ErdOS [5], CoSense [8], OSone [25] and M+ [26].

ErdOS leverages resources in nearby devices based on user modeling and stated user preferences. It uses a lightweight IPC and network stack to securely broadcast important context information and app data in a user-level communication manager. The implementation in CoSense distributes the sensing tasks between familiar devices that are in close proximity. The group formation is done by the cloud backend once the devices are registered to the cloud backend with their mac address. Once the groups are formed, the data is transferred via local connection.

OSone distributes the functionality of the operating system in a similar fashion to how Barrelfish [27] separates functionality onto different cores. The architecture consists of a kernel node in charge of various host nodes that can be kept simpler. M+ allows cross-device functionality sharing. It uses remote procedure call scheme based on the binder IPC mechanism to utilize app and system functionalities across devices.

Moreover, the work such as in Reptor [28] allows third parties to easily distribute their modifications for a platform without the need to update the entire platform. This provide ease for the open innovation for the multi-device platforms for resources utilization. In contrast to the above systems, Rio [29] presents a system where a device's resources are utilized by remotely accessing them with the help of another device in close proximity.

We implement AFV in a similar fashion with the potential to have multiple controlling nodes over time, depending on remaining resources and app needs. AFV runs on a PAN where the devices are already connected to each other via Bluetooth. Therefore, dynamic group formation is not considered in AFV. Moreover, most of these systems are designed for the optimal usage of the sensing function. However, the AFV framework considers all the available common functions (i.e., sensing, connectivity and computing) in the wearable personal area network. Also, we

consider different optimization objectives than energy-related objectives, such as network quality, functionality precision/quality and also monetary costs.

We follow the philosophies initiated in the early design work on wearables. In particular, Speakeasy [15] motivates the need for domain independent interfaces, mobile code, and user interpretation of semantics. Our representation of context is similar to that provided by Speakeasy and we retain user discernment as well. We are less ambitious in the overall goals as we leverage existing APIs and focus on the adaptive nature of wearable network apps. We do not implement code migration; we provide a system-level extension to code already available on wearable devices. Smailagic and Sieworek [16] provide key design principles/challenges for future wearable apps: user interface models, input/output modalities, matched capability with requirements, and quick interface evaluation methodology. We focus on the third of these challenges to meet the user's needs with the lowest resource utilization.

7 Discussion

AFV is designed with the premise that all devices in a PAN are always connected and managed by the same person. Therefore, we have not considered the option of dynamic device group formation with nearby devices owned by other people. This is primarily to reduce the privacy and security concerns of communicating with untrusted devices of strangers. On the other hand, we assumed that there are no privacy or security risk in communicating or utilizing functions on the trusted devices on the same PAN. However, this assumption may not always be true, as the third-parties such as trackers, intruders and manufacturers have the access and partially control some functionalities of the devices and its data. Therefore, we intend to mitigate the potential threats of information leakage with the PAN by extending AFV with a context-aware security framework incorporating a set of pre-defined and also user defined device access policies. These policies will then be considered during function allocation as another context information. For example, if a device is connected to a public WiFi access point, the device may not be used to virtualize functions by the *Decision Engine*.

In this paper, we have only validated AFV performance for limited functionalities (i.e., sensing and Internet connectivity), although AFV is designed for efficient utilization of many other functions such as compression, encoding and anonymization. We have noticed the potential difference in overheads associated with each function. Therefore, we aim to further strengthen our function allocation algorithm considering available memory and CPU power as additional context. In addition, for each particular use case, we considered a single objective, which are specified in Table 1. However, a user may wish to achieve multiple objectives at the same time. As an example, a user may wish to have the best quality of information while achieving the minimum possible energy consumption. This can be addressed by formulating a multi-objective optimization problem in the *Decision Engine* with different weighting factors for each of the objective. These weighting factors are to be specified by the user via the UI provided by AFV.

Although we developed AFV prototype as a standalone user-level app and a library, it can also be realized by integrating it into the OS as a module (requiring root access permission to the kernel). OS module implementation will be efficient in terms of systems overheads of AFV, but it requires significant development effort as well as reduces the deployability of AFV. However, despite this implementation overhead, we showed that AFV outperforms the vanilla scenario without AFV. Therefore, we aim to further improve the user-level development to release AFV as a software development kit (SDK) for app developers, and also, envisage the implementation in multiple OSs.

8 Conclusion

The majority of devices in a personal area network that consists of multiple smart wearables and hand-held devices have a number of common capabilities and resources. However, current popular mobile and wearable apps do not utilize these resources efficiently that leads to multiple of app function executions in the same personal area network. As a result, the users may not get the best outcome, may incur higher networking cost and may also result in higher energy consumption of devices.

In this paper, we proposed AFV, an architecture that overcomes the above inefficiencies, while reducing the overall energy usage, without adding any latency and minimizing the communication overhead. AFV enables context-aware app function virtualization in a personal area network with a set of APIs that can be easily leveraged by app developers during app development. Our simulation results showed that the proposed function allocation algorithm enables system uptime improvement of up to 35-40 percent compared with typical configurations of current wearable/mobile apps. Then, we showed the viability of the architecture by implementing AFV in Android devices without loss of generality. Finally, we showed the real world applicability of AFV and user benefits via emulating multiple use cases with real devices.

References

[1] H. Kolamunna, Y. Hu, D. Perino, K. Thilakarathna, D. Makaroff, and A. Seneveratne, "AFV : Enabling application function virtualization and scheduling in wearable networks," in *Proc. ACM Int. Joint Conf. Pervasive Ubiquitous Comput.*, Sep. 2016, pp. 981–991.

[2] R. Mijumbi, J. Serrat, J.-L. Gorricho, N. Bouten, F. D. Turck, and R. Boutaba, "Network function virtualization: State-of-the-art and research challenges," *IEEE Commun. Surv. Tut.*, vol. 18, no. 1, pp. 236–262, Jan.–Mar. 2016.

[3] A. Beach, M. Gartrell, X. Xing, R. Han, Q. Lv, S. Mishra, and K. Seada, "Fusing mobile, sensor, and social data to fully enable context-aware computing," in *Proc. 11th Workshop Mobile Comput. Syst. Appl.*, Feb. 2010, pp. 60–65.

[4] S. Kang, J. Lee, H. Jang, H. Lee, Y. Lee, S. Park, T. Park, and J. Song, "SeeMon: Scalable and energy-efficient context monitoring framework for sensor-rich mobile environments," in *Proc. 6th Int. Conf. Mobile Syst. Appl. Serv.*, Jun. 2008, pp. 267–280.

[5] N. Vallina-Rodriguez and J. Crowcroft, "ErdOS: Achieving energy savings in mobile OS," in *Proc. 6th Int. Workshop MobiArch*, Jun. 2011, pp. 37–42.

[6] D. P. Williamson and D. B. Shmoys, *The Design of Approximation Algorithms*, 1st ed. New York, NY, USA: Cambridge Univ. Press, 2011.

[7] Y. Xiao, Y. Cui, P. Savolainen, M. Siekkinen, A. Wang, L. Yang, A. Yl-Jski, and S. Tarkoma, "Modeling energy consumption of data transmission over Wi-Fi," *IEEE Trans. Mobile Comput.*, vol. 13, no. 8, pp. 1760–1773, Aug. 2014.

[8] S. Hemminki, K. Zhao, A. Y. Ding, M. Rannanjärvi, S. Tarkoma, and P. Nurmi, "CoSense: A collaborative sensing platform for mobile devices," in *Proc. 11th ACM Conf. Embedded Netw. Sensor Syst.*, 2013, pp. 34–35.

[9] H. Shen, A. Balasubramanian, A. LaMarca, and D. Wetherall, "Enhancing mobile apps to use sensor hubs without programmer effort," in *Proc. ACM Int. Joint Conf. Pervasive Ubiquitous Comput.*, Sep. 2015, pp. 227–238.

[10] C. Randell, "Wearable computing: A review," Univ. Bristol, Bristol, U.K., Tech. Rep. CSTR-06–004, 2005.

[11] D. Chu, A. Kansal, J. Liu, and F. Zhao, "Mobile apps: It's Time to move up to condOS," in *Proc. 13th USENIX Conf. Hot Topics Operating Syst.*, May 2011, pp. 1–5.

[12] G. Kortuem, Z. Segall, and M. Bauer, "Context-aware, adaptive wearable computers as remote interfaces to 'intelligent' environments," in *Proc. 2nd IEEE Int. Symp. Wearable Comput.*, Oct. 1998, pp. 58–65.

[13] M. Conti, B. Bruno Crispo, E. Fernandes, and Y. Zhauniarovich, "CRêPE : A system for enforcing fine-grained context-related policies on android," *IEEE Trans. Inf. Forensics Security*, vol. 7, no. 5, pp. 1426–1438, Oct. 2012.

[14] P. K. McKinley, S. M. Sadjadi, E. P. Kasten, and R. Kalaskar, "Programming language support for adaptable wearable computing," in *Proc. 6th Int. Symp. Wearable Comput.*, Oct. 2002, pp. 205–212.

[15] W. K. Edwards, M. W. Newman, J. Sedivy, T. Smith, and S. Izadi, "Challenge: Recombinant computing and the speakeasy approach," in *Proc. 8th Annu. Int. Conf. Mobile Comput. Netw.*, Sep. 2002, pp. 279–286.

[16] A. Smailagic and D. Siewiorek, "Application design for wearable and context-aware computers," *IEEE Pervasive Comput.*, vol. 1, no. 4, pp. 20–29, Oct. 2002.

[17] S. Brachmann, "Wearable gadgets: What is the secret to commercial success?" 2014. [Online]. Available: http://www.ipwatchdog.com/2014/11/15/wearable-gadgets-the-secret-to-commercial-success/id=52159/

[18] J. Mischke, "Wearables for professional and industry applications," 2014. [Online]. Available: https://www.wearable-technologies.com/2014/03/wearables-for-professional-and-industry-applications/

[19] R. Rawassizadeh, B. A. Price, and M. Petre, "Wearables: Has the age of smartwatches finally arrived?" *Commun. ACM*, vol. 58, no. 1, pp. 45–47, Dec. 2014.

[20] M. Martins, J. Cappos, and R. Fonseca, "Selectively taming background android apps to improve battery lifetime," in *Proc. USENIX Conf. USENIX Annu. Tech. Conf.*, 2015, pp. 563–576.

[21] S. Elmalaki, L. Wanner, and M. Srivastava, "CAreDroid: Adaptation framework for android context-aware applications," in *Proc. 21st Annu. Int. Conf. Mobile Comput. Netw.*, Sep. 2015, pp. 386–399.

[22] A. Kansal, S. Saponas, A. B. Brush, K. S. McKinley, T. Mytkowicz, and R. Ziola, "The latency, accuracy, and battery (LAB) abstraction: Programmer productivity and energy efficiency for continuous mobile context sensing," in *Proc. ACM SIGPLAN Int. Conf. Object Oriented Program. Syst. Lang. Appl.*, Oct. 2013, pp. 661–676.

[23] T. Kaler, J. P. Lynch, T. Peng, L. Ravindranath, A. Thiagarajan, H. Balakrishnan, and S. Madden, "Code in the air: Simplifying sensing on smartphones," in *Proc. 8th ACM Conf. Embedded Netw. Sensor Syst.*, Nov. 2010, pp. 407–408.

[24] L. Ravindranath, A. Thiagarajan, H. Balakrishnan, and S. Madden, "Code in the air: Simplifying sensing and coordination tasks on smartphones," in *Proc. 12th Workshop Mobile Comput. Syst. Appl.*, Feb. 2012, pp. 4:1–4:6.

[25] B. Pasztor and P. Hui, "OSone: A distributed operating system for energy efficient sensor network," in *Proc. 25th Int. Teletraffic Congress*, Sep. 2013, pp. 1–9.

[26] S. Oh, H. Yoo, D. R. Jeong, D. H. Bui, and I. Shin, "Mobile plus: Multi-device mobile platform for cross-device functionality sharing," in *Proc. 15th Annu. Int. Conf. Mobile Syst. Appl. Serv.*, Jun. 2017, pp. 332–344.

[27] A. Baumann, P. Barham, P.-E. Dagand, T. Harris, R. Isaacs, S. Peter, T. Roscoe, A. Schüpbach, and A. Singhania, "The multikernel: A new OS architecture for scalable multicore systems," in *Proc. ACM SIGOPS 22nd Symp. Operating Syst. Principles*, Oct. 2009, pp. 29–44.

[28] T. Ki, A. Simeonov, B. P. Jain, C. M. Park, K. Sharma, K. Dantu, S. Y. Ko, and L. Ziarek, "Reptor: Enabling API virtualization on android for platform openness," in *Proc. 15th Annu. Int. Conf. Mobile Syst. Appl. Serv.*, Jun. 2017, pp. 399–412.

[29] A. A. Sani, K. Boos, M. H. Yun, and L. Zhong, "Rio : A System Solution for Sharing I/O between Mobile Systems," in *Proc. Annu. Int. Conf. Mobile Syst. Appl. Serv.*, Jun. 2014, pp. 259–272.

Harini Kolamunna received the bachelor's degree in electrical & electronics engineering with a first class honors from the University of Peradeniya, and worked as a research assistant at the National University of Singapore in 2013-2014. She is working toward the PhD degree in the School of Electrical Engineering & Telecommunications, UNSW Australia, and attached to Data61-CSIRO. Her current research interests include wearable/IoT technologies, networking, and communications.

Kanchana Thilakarathna received the PhD degree in electrical engineering and telecommunications from UNSW Australia. He is a lecturer in distributed computing with the School of Information Technologies, The University of Sydney. Previously, he was a research scientist in the Networks Group, Cyber-Physical Systems Research Program at Data61-CSIRO. His research interests include developing technologies for resource allocation, secure communication and authentication, and privacy-preserving data sharing in wearable/mobile/IoT networks.

Diego Perino received the MS degree from Politecnico di Torino and Eurecom institute, and the PhD degree in computer science from the Paris Diderot-Paris 7 University. He is a researcher at Telefonica Research in Barcelona, Spain. His research focuses on design and performance evaluation of networking protocols and systems. He has published several papers at international conferences and in journals, and also led several patents.

Dwight Makaroff received the PhD degree in multimedia systems from the University of British Columbia in 1998. He is a professor with the Computer Science Department, University of Saskatchewan, Canada. His current research interests are distributed systems performance, including network protocols, sensor networks, big data architecture frameworks, wearable systems, and networked games.

Aruna Seneviratne received the PhD degree in electrical engineering from the University of Bath, United Kingdom. He is a professor at UNSW, Australia. He was the foundation professor of telecommunications at UNSW, where he holds the Mahanakorn chair of Telecommunications. His research interests are in mobile technologies, networking and communications, and computer system security.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.