

vSensor: Toward Sensor-rich Mobile Applications

Minsung Jang, HyunJong Lee, Ketan Bhardwaj, Karsten Schwan
{firstname.lastname@cc.gatech.edu}

Georgia Institute of Technology

ABSTRACT

Interacting with sensors presents numerous challenges, particularly for applications running on resource-constrained platforms like smartphones or tablets. This paper suggests vSensor abstractions providing such applications with the ease of use for today's multitude of sensors. vSensor abstractions also interact with nearby and cloud resources for the applications to easily process sensor-acquired data and control actuators in a consistent way, which permits applications to scale their use of sensors by leveraging other nearby and cloud resources. The vSensor abstraction developed in our research seek to provide globally transparent, uniform access to the ambient sensors to unmodified applications as well as those with the vSensor awareness. We implement and evaluate the vSensor abstraction on Android platforms.

1. INTRODUCTION

The Gartner group predicts that the total number of sensors to grow to 26 billion units worldwide by 2020 [2]. This plethora of sensors enables a rich set of applications that can assist users in their daily lives, to help capture user intent and context [5], to enable intuitive and meaningful user interaction, and to assist with specific tasks as in elder care [12], efficient driving [14], etc.

Our work and this paper address the challenges faced by applications that seek to leverage and use the dynamic sets of sensors present on mobile devices and in the environments in which they operate.

Three particular issues faced by such applications are:

- Diverse nature of those sensors, in part because of the fragmented nature of today's sensor ecosystem;
- Computational and data management challenges in interacting with those sensors, particularly for applications running on resource-constrained end devices like smartphones or tablets; and
- Dynamic nature of sensor presence, as users move in and out of their proximity and run applications requiring their dynamic access. A simple example is a smartphone app desiring to interact with sensors in the home when the homeowner pulls into

New Applications for Mobile Devices in a sensor-rich world

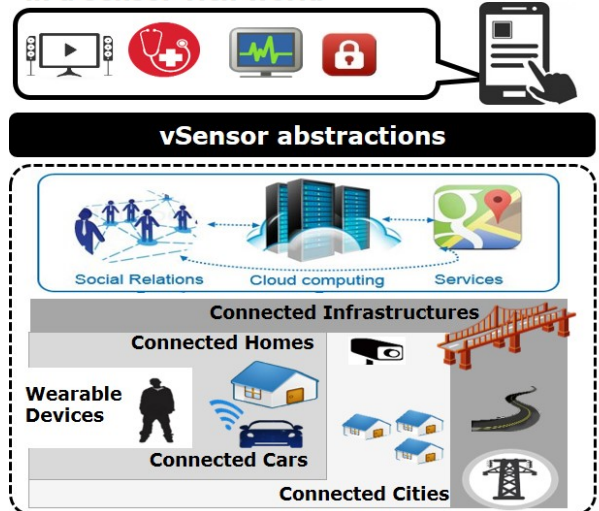


Figure 1: Vision of vSensor in a Sensor-rich world

the garage. Doing so requires the app to dynamically acquire access privileges to those sensors, followed by subsequent phone-sensor interactions that involve extensive communication and processing activities, and developing such app functionality involves understanding and handling the rich variety of sensors present in the home and finding ways to run app functions within the constrained resources available on smartphones.

The object of our research and of this paper are to make it easier for applications to leverage and exploit the on- and off-device sensors present in their current environments. We do so by developing the vSensor (virtual sensors) abstractions that permit applications running on resource-constrained devices like smartphones to interact with and manage the dynamic set of currently accessible sensors. We achieve this by:

- *virtualizing* individual sensors – virtual sensors – to obtain a uniform way to interact with them, thus removing from applications the need to understand in detail how to interface with the many special-

ized and custom sensors that surround them; and

- *externalizing* sensor interactions and processing from the resource-constrained mobile device to nearby and remote cloud resources – vSensor abstractions – to leverage the computational and storage abilities of such always-on facilities for running the complex functionality implemented by sensor-rich applications; and

Figure 1 depicts how vSensor make available to applications the sensors present in their current environments, with application services running on both edge and datacenter cloud systems, ranging end user devices, to nearby computing resources (e.g., a home PC), and/or in the remote cloud.

Our vSensor implementation for Android platforms leverages the representative ‘edge cloud’ infrastructure described in [9], which permits services to run ‘anywhere’, i.e., on devices, nearby machines, and the remote cloud. vSensor extends these underlying functionalities with its new sensor abstractions and access control methods, where sensor-based applications can construct virtual sensors, create services that run on edge cloud (i.e., nearby) or remote cloud resources, and orchestrate their execution as sets of services running on the device and/or in the cloud. This paper’s evaluations of vSensor with concrete applications show them fully interacting with a variety of ambient sensors and resources, without additional overheads compared to the less convenient, direct application-sensor interactions, with the additional ability to scale those interactions to 100s of sensors, actuators, and associated services, even on resource-constrained platforms like smartphones.

The remainder of the paper is organized as follows. Mobile applications’ (apps’) current usage of sensors is explained in Section 2, with vSensor-enabled use cases shown in Section 3. The design and implementation of vSensor are described in Sections 4 and 5, respectively. Experimental results in Section 6 demonstrate our initial implementaion results. Conclusions are in Section 7.

2. MOTIVATION

With the advent of a world of a trillion sensors and mobile devices, a new class of applications has been predicted to emerge rapidly, providing end users with personalized services based on the precise capture of their current contexts and intents. Doing so, however, requires those apps to interact with the numerous sensors present on mobile devices and in users’ current environments, to extract personal user contexts and environmental conditions. Unfortunately, today’s reality is that most mobile applications (apps) hardly use sensors, despite the fact that the devices hosting such apps are themselves sensor-rich. Estimates [3] are that apps

Sensor Count	Apps(top 100)
1	64
2	13
3	3
4	0
5	1
6	0
7	0
8	0
Total(%)	81 (1.62%)

Table 1: The number of apps based on their sensor use. The total number of top 100 apps in each category of the Google Play Store are 5000 on August 21, 2014.

using at least one sensor in Android devices are just 0.5% of all available apps in 2012, with typical devices equipped with six sensors on average.

This section updates such estimates with a comprehensive study of how today’s mobile applications interact with sensors. The intent is to better understand how apps interact with sensors, to motivate the functionality of vSensor intended to benefit them. The study is also a starting point for how vSensor functionality can enable new classes of mobile apps.

2.1 Methodology

Our study uses a tool that downloads apps from the Google Play Store that meet some predefined conditions, based on recently announced techniques like [1]. For such apps, we then analyze how they behave on the Android platform by scrutinizing their Dex bytecodes and manifest files. An initial study used the top 100 apps in each category on the Google Play Store, resulting in a total of 5000 apps (on August 21, 2014). We plan to expand to almost all free apps in the Store.

2.2 Discussion of Study Outcomes

As shown in Table 1, for the top 100 apps, 81 out of 5000 (1.62%) use at least one sensor. This constitutes only a small increase over the previous estiate of 0.5% reported in 2012.

Evident from the statistics reported above is the fallacy of predictions made in the academic literature and in industry that mobile devices will become hubs in a sensor-rich world. In theory, many applications could benefit from leveraging all sensors on devices and in a user’s current environment (e.g., sensors from wearable devices, homes, and cars). In practice, this is not the case, and it remains unclear why today’s apps do not actively leverage even the potential utility of the sensors present on their own devices.

vSensor’s approach in reaction to these results is to try to make it easier for apps to use on-device and

nearby sensors, (i) by addressing the fragmentation issue we have found for the sensors and devices used in our work (see Section 5.1), and (ii) by dealing with the dynamic nature and with sensor protection for the sensors currently accessible from a mobile device.

3. APPLICATIONS WITH VSENSOR–CURRENT AND FUTURE USE CASES

This section describes sensor-rich applications written with vSensor, able to easily leverage surrounding sensors to collect notable events and employ nearby resources to process or/and react to them.

3.1 Prototype Applications with vSensor

3.1.1 Unmodified Store Apps using Local and Remote Sensors

A goal of vSensor is to make it easier for existing applications to transparently and seamlessly scale in terms of numbers of sensors with their associated data processing activities, without the need to modify said applications whenever additional or new sensors are used. We demonstrate this functionality with a simple Play Store applications, termed ‘sensor readout’, which reads the outputs produced by all sensors embedded in an Android device. With vSensor, an unmodified app like this can interact with both on- and off-device accessible sensors.

3.1.2 Composed Sensors – Combining Sensors via Services

vSensor can abstract from individual physical sensors to form new sensor types, but a common application need is to combine and make use of multiple sensors to realize some desired app-level functionality. We consider such functionality another case of sensor virtualization and extension. vSensor can group many sensors of the same type, e.g., all the light sensors in a home. A concrete example evaluated in our work is an application permitting end users to check the current time (on their smartphone), but without turning on the smartphone’s battery-consuming screen. This ‘don’t turn on the screen’ application uses a set of virtual sensors offered by vSensor with access to a home camera and a phone’s speaker, and it runs its some complex software services on the PCloud. The service implements finger-gesture recognition using the camera: if the user approaches the smartphone with two fingers, this is interpreted as a desire to check time rather than grasp the phone; the response is the phone’s speaker stating the current time, without unlocking and activating the screen, thus conserving phone power.

3.2 Future Use Cases

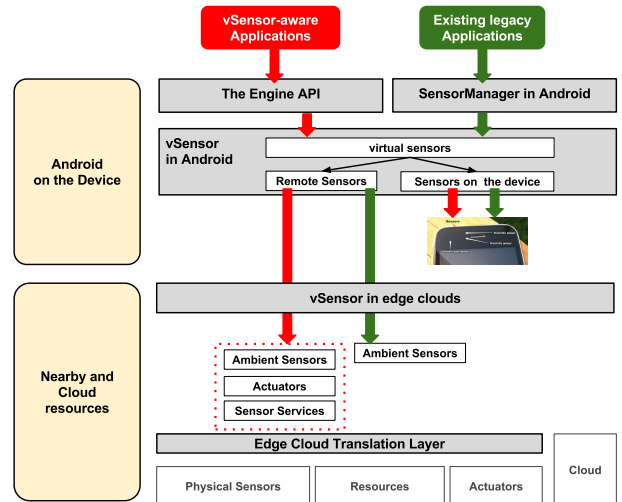


Figure 2: vSensor Design

An additional application currently under development realizes the aforementioned ‘health virtual sensor’. This application will allow mobile devices to become hubs for health-related information about the device owner. The application periodically measures data like blood sugar levels, data acquired from wearable devices (e.g., *health bracelets*), and even data acquired from say, an exercise bike used by the owner in a health club. Analytic services part of the object immediately raise alarms if unusual readings are detected from the object’s virtual sensors, and in addition, they interact with a cloud-resident service to compute long term health statistics, implement a dashboard, etc.

4. VSENSOR DESIGN

The key design goal of vSensor is to enable apps running on smartphones and similarly resource-constrained end user devices to more easily and efficiently interact with the many sensors present on the devices themselves and in their surroundings. Concerning ease of use, vSensor seeks to overcome difficulties related to using many diverse sensors and to handle their potential incompatibilities and custom APIs. Concerning efficiency, vSensor makes possible the use of ambient nearby resources and of the remote cloud, for complex processing and transformation of sensor data. Figure 2 depicts an overview of vSensor’s design, discussed in more detail next.

4.1 Design Principles

The vSensor abstraction enables apps to create points of access and use for sensors and associated sensor services and actuators, for single or sets of sensors and regardless of where these sensors are physically located, i.e., on the app’s local platform or accessible remotely. **Logical decoupling – sensor virtualization.** vSensor virtualizes physical sensors, the goal being to cre-

ate meaningful sensor abstractions with uniform APIs vs. the diverse, custom APIs of available physical sensors. Virtualization also provides substantial flexibility for crafting exactly the abstractions desired by applications, an example being a ‘health’ sensor implemented by virtual sensors comprised of a scale, a heartbeat sensor on a bracelet, and cameras watching the exercise equipment being used.

Flexible service execution – sensor services. The location of a virtual sensor should not determine where its service codes manipulating sensor data are run: on the sensor itself, on the platform hosting it, or on other available processing resources. There is ample previous work demonstrating the need for such flexibility, ranging from early work on adaptive sensor processing [11] or adaptive media manipulation [13] to recent work on edge cloud functionality [6] and cloud offloading for sensors [10]. In accordance with such insights, virtual sensors are able to run on any underlying platform made available by edge cloud infrastructures [9, 7, 8, 4] with which vSensor can interact (see the ‘translation layer’ in Figure 2). Currently, this is the PCloud edge cloud [9], but it is straightforward to write such a layer for other infrastructures.

Backward compatibility – sensor illusions. The virtual sensors via vSensor abstractions should remain available to existing applications that use them. We achieve this backward compatibility by having vSensor export ‘illusions’ of those sensors to legacy apps, while at the same time, providing the advanced functionality of use to new applications via a new set of our APIs.

4.2 Basic Operations

Virtual sensors are created and managed with operations that include *group*, *glance*, *aggregate*, *trim*, *bundle*, and *reconfigure* to meet the design principles articulated in Section 4.1. Apps invoke such operations via the ‘Engine API’ on the Android platform also shown in Figure 2 and as with object services, these operations can again be run on any of the resources made available by the underlying edge cloud – on the device making the call, on some nearby resource, or in the remote cloud. We next describe some of these API calls and their functionality.

5. SELECT IMPLEMENTATION DETAIL

This section presents the implementation details needed for understanding the performance evaluation in Section 6.

5.1 Android Realization

Our Android implementation must be (i) backwards compatible, (ii) transparent to their physical nature, and (iii) portable permitting virtual sensors to run on any of the variety of edge cloud infrastructures currently

Name	Hardware/Role(s)
Sensor Nodes	Intel Galileo 1st Gen
Actuator Nodes	Intel Galileo 1st Gen
Camera Nodes	Exynos 5420 and AMD E450
EC2	m3.large (Cloud resource)
Local 1	Intel i5 (Local resource 1)
Device	Galaxy S4 (User’s mobile device)

Table 2: vSensor Testbed Setting

under development [7, 8, 9]. We obtain these properties as follows. First, to existing sensor-based apps, virtual sensors wrapping existing sensors provide to these apps the aforementioned ‘illusions’ of the physical sensors these apps currently employ.

Second, with sensor virtualization and the ability to construct entirely new sensors from both physical sensors and events (e.g., calendar events) associated with logical artifacts, transparency becomes a pervasive property of vSensor’s Android implementation. Third, we obtain portability by providing a translation layer for running vSensor on the PCloud edge cloud.

5.2 The Engine API

In Android, apps using vSensor access virtual sensors via the ‘Engine API’, which implements methods that correspond to the operations described in Section 4.2. The API provides access to both local and remote sensors, via virtual sensors wrapped by vSensor abstractions.

6. EXPERIMENTAL EVALUATION

We evaluate use cases in Section 3.1 to demonstrate the utility and performance of the vSensor abstraction.

6.1 Evaluation Setup

Table 2 describes a set of resources in our experiments. Intel Galileo boards(1st Generation) are used for remote sensor and actuator nodes. The remote sensor nodes measure humidity and temperature at their locations while the actuator nodes send notifications to users via embedded LED bulbs. The camera nodes keep track of the figure’s gesture of users.

6.2 Initial Results

6.2.1 Unmodified Store Apps using Local and Remote Sensors

We verify the backward compatible capability of the vSensor abstractions by checking if the existing apps in the Google Play Store can transparently access the virtual sensors resulting from what vSensor virtualizes local and remote sensors. Figure 3 is a screenshot showing that such an app, Sensor Readout, transparently interacts with local (on-device) and remote (off-device)

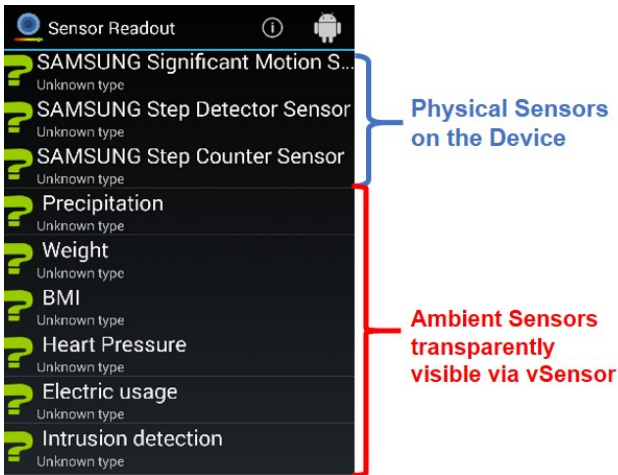



Figure 3: A Screenshot of the Sensor Readout App

sensors. First three sensors in Figure 3 are those on our test device, Samsung Galaxy S4, while the rest are virtual sensors that are local illusions of remote sensors created by the vSensor abstraction.


6.2.2 Composed Sensors – Combining Sensors via Services

The Engine API enables mobile applications to combine different virtual sensors to a new kind of a single virtual sensor that interacts with actuators and services. This ‘don’t turn on the screen’ application combines a set of cameras at home into a single virtual sensor object, which is ‘user’s finger gesture detection sensor’. The application uses this virtual sensor and runs its finger-gesture recognition service on either the nearby resources or the remote cloud (Amazon EC2) based on a scheduling decision made by the PCloud. The application, then, makes the speaker on the phone stating the current time. Results in Figure 4 summarize our experiments for the app. We check the current time ten times during ten seconds for each by turning on the screen (labeled as ‘Screen’) and by the app (labeled as ‘Speaker’). Figure 4 (a) shows the elapsed time between the finger detection and stating the current according to locations where the recognition service runs. If the service runs on local resources (labeled as ‘Pcl-WLAN’), the response time is around 97.4ms while the cloud (labeled as ‘EC2-WLAN’) 294.5ms. We also compare the power and energy consumption of the app with those without the app. Figure 4 (b) clearly shows that the app can dramatically reduce the energy consumption by 46%.

These initial results for two use cases shows that (i) the existing apps can interact with not only sensors on a device but those remotely available without having to change their source code, and (ii) vSensor can contribute to open up a new opportunity for mobile applications via composing sensors.



(a) Response time to state the current time. Error bars show 95% confidence intervals, respectively.



(b) Power consumption

Figure 4: Results of the ‘Don’t turn on the screen’ application. For the total energy consumption, the speaker and the screen are 22.14 and 40.94 mWH, respectively. Each circle shows the moment that a user acknowledges the current time.

7. CONCLUSIONS

With the vSensor approach and abstractions, we address several issues with today’s sensors and sensor processing ecosystem. First, with vSensor’s uniform virtual sensors, it becomes easier for applications to interact with and leverage available custom and specialized sensors. Second, by leveraging edge and remote clouds, vSensor makes feasible complex sensor processing and integration activities that are not limited by a smartphone’s computing, storage, and battery constraints.

8. REFERENCES

- [1] Dex2jar tool. <https://code.google.com/p/dex2jar/>.
- [2] Press release. <http://goo.gl/CziKJa>.
- [3] Sensor-based apps offer lots of potential but are hindered by fragmentation. <http://goo.gl/aUhi8N>.
- [4] Smartthings. <https://www.smartthings.com>.
- [5] BHARDWAJ, K., SREEPATHY, S., GAVRILOVSKA, A., AND SCHWAN, K. Ecc: Edge cloud composites. In *Mobile Cloud Computing, Services, and Engineering (MobileCloud), 2014 2nd IEEE International Conference on* (April 2014), pp. 38–47.
- [6] BUEVICH, M., WRIGHT, A., SARGENT, R., AND ROWE, A. Respawn: A distributed multi-resolution time-series datastore. In *Real-Time Systems Symposium (RTSS), 2013 IEEE 34th* (Dec 2013), pp. 288–297.
- [7] DIXON, C., MAHAJAN, R., AGARWAL, S., BRUSH, A. J., LEE, B., SAROIU, S., AND BAHL, P. An operating system for the home. In *Proceedings of the*

- 9th USENIX Conference on Networked Systems Design and Implementation* (Berkeley, CA, USA, 2012), NSDI'12, USENIX Association, pp. 25–25.
- [8] HA, K., CHEN, Z., HU, W., RICHTER, W., PILLAI, P., AND SATYANARAYANAN, M. Towards wearable cognitive assistance. In *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services* (New York, NY, USA, 2014), MobiSys '14, ACM, pp. 68–81.
- [9] JANG, M., SCHWAN, K., BHARDWAJ, K., GAVRILOVSKA, A., AND AVASTHI, A. Personal clouds: Sharing and integrating networked resources to enhance end user experiences. In *INFOCOM, 2014 Proceedings IEEE* (April 2014), pp. 2220–2228.
- [10] LIU, B., JIANG, Y., SHA, F., AND GOVINDAN, R. Cloud-enabled privacy-preserving collaborative learning for mobile sensing. In *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems* (New York, NY, USA, 2012), SenSys '12, ACM, pp. 57–70.
- [11] ROSU, D., SCHWAN, K., YALAMANCHILI, S., AND JHA, R. On adaptive resource allocation for complex real-time applications. In *Real-Time Systems Symposium, 1997. Proceedings., The 18th IEEE* (Dec 1997), pp. 320–329.
- [12] SILVA, L. C. D., MORIKAWA, C., AND PETRA, I. M. State of the art of smart homes. *Engineering Applications of Artificial Intelligence* 25, 7 (2012), 1313 – 1321. Advanced issues in Artificial Intelligence and Pattern Recognition for Intelligent Surveillance System in Smart Home Environment.
- [13] STEERE, D. C., GOEL, A., GRUENBERG, J., MCNAMEE, D., PU, C., AND WALPOLE, J. A feedback-driven proportion allocator for real-rate scheduling. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation* (Berkeley, CA, USA, 1999), OSDI '99, USENIX Association, pp. 145–158.
- [14] THRUN, S. Toward robotic cars. *Commun. ACM* 53, 4 (Apr. 2010), 99–106.