Unveiling the Data Shadow: A Scalable Software Architecture for Public Health and Electronically Assessed Data (PHEAD)

Tor-Morten Grønli[1], Andreas Biørn-Hansen[1], Siri Fagernes[1], Ragnhild Eg[2], Kjeld Hansen[1], and Per Morten Fredriksen[2]

[1]Mobile Technology Lab, Department of Technology, Kristiania University College, Norway
[2]Kristiania University College, School of Health Sciences, Norway

Abstract

Health expenditures in Norway amounted to 10,2 percent of GDP in 2018, and the budget was dedicated predominantly to treating disorders and maintaining support functions. Only 3 percent of the budget went to preventive health, despite the fact that preventive measures hold the greatest promise. At the same time, computer-based technology enables measuring and gathering large amounts of health-related data, like a person's activity level, heart rate, sleep patterns, eating habits, exercise habits and even emotional state. The popularity of personal electronic recording equipment, along with the extensive use of social media, result in large collections of data that are continuously expanding. Currently this wealth of information is not accessible to healthcare professions, instead it is contained within a data shadow. Even with access, the majority of health personnel would be unable to break the data down to sensible and usable information.

This research project is founded in the perspective of preventive health, aiming to synthesize available personal health information by utilizing commodity mobile and wearable hardware. It sets out to harvest the potential inherent in the constellation of shadow data, returning the insight to the individual as personalised health advice. Through a persuasive technologies experiment, the project will explore efficient means to motivate healthy lifestyle choices. Furthermore, the project presents a prototype architecture for the collection and processing of unfiltered data that will serve as a foundation for health workers and medical doctors to base diagnoses, treatments and council upon. Our major contribution is a proof of concept implementation and leveraging state of the art cloud based function as a service approach to build a scalable software architecture for a ubiquitous and heterogeneous environment harvesting the data shadow through activity tracking devices.

1 Introduction

Countries worldwide are experiencing an increase in life-style related diseases of epic proportions. The global burden of disease is mainly caused by overweight, obesity and

This paper was presented at the NIK-2019 conference; see http://www.nik.no/.

inactivity. World-wide numbers count more than 650 billion overweight or obese adults and close to 400 million children and youths [1]. Combined with an ageing western population, these numbers bring with them disorders such as diabetes, muscle skeletal conditions and cardiovascular diseases [2]. One consequence is a financial burden that will increase dramatically. Health expenditure in Norway amounted to 360 billion in 2018, corresponding to 10,2 percent of the national GDP, a 0,6 increase per inhabitant since the previous year [3]. 97 percent of this budget is spent on treating disorders and on support functions. A meager 3 percent are allocated preventive health measures, despite the amassing research body on their beneficial effects [3]. Considering that the challenge with many of the most common illnesses today is their relation to lifestyle and individual choices, it seems the economically most feasible and effective solution is the prevention of illness. The catch is that prevention requires information. To provide the appropriate preventive care and advice, health workers, medical doctors and other professionals would need all relevant information related to the current condition of a patient. However, significant amounts of information related to a person's lifestyle are largely concealed from these professionals. Consequently, it is difficult (if not impossible) to make the correct diagnosis, prescribe the correct treatment, supply meaningful advice or give life counseling. Up until now, information regarding a person's health has only been available if: measured by a healthcare professional, provided through medical history, and knowledgeably shared by the patient.

Meanwhile, the world drifts towards digitisation and digitalisation in close to every field and area, so also in terms of health related information. There are vast amounts of information available and health related information is now used to improve a number of aspects of personal life. In a matter of short time, changes in health diagnosis, treatment and research are forthcoming. Large international players such as Apple, Google and Amazon are investing heavily and general opinion awareness is growing together with governmental and national enquiries and campaigns. Internet of Things, mobile devices and wearable computing have led to a growing ecosystem for data rich services and a diverse universe for the end users. Systems and application development have typically been conducted targeting specific user needs. We now see a shift towards more heterogeneous environments, which are *task oriented* instead, hence recognizing the users' flow of action over multiple devices to complete a task [4].

Rapid technological innovations and continued digitalisation ensure that almost any and every routine and activity can be recorded and analysed, and many choose to do so. Over time, conscious choices can become habitual; furthermore, some systems come with with interfaces that are designed to incentivate continued activity. Something as simple as a milestone reached or an acknowledging remark can serve as rewards, and rewards increase motivation [5]. With continued use, a system can analyse input data and adapt its rewards to the user's behaviour. In social media and various entertainment applications, the adaptation can go beyond simple reward mechanisms, for instance by prioritising content to pique interest further [6]. The same type of adaptation could serve as a motivational mechanisms in health-promoting applications, using adaptive feedback to establish and reinforce a reward cycle for healthy lifestyle choices.

2 Related Work

In this section we will elaborate on related work to our research in accordance with the three major pillars in the project, personal health information, software architecture and persuasive technologies.

Personal Health Information

Diseases move along a continuum from healthy to sick, yet current treatment strategies prioritise patients at the sicker end of the spectrum [7]. This implies that the vast majority of health measures is directed at a minority of the population, while those at lower risk are left to themselves. Assuming an eventual progress along the disease continuum, the majority of the population can be considered in a stage of sick, albeit without clinical diagnoses or symptoms. This creates a paradox; although the number of sick or high risk patients are reduced, the number who will eventually require medical help increases.

The challenge remains to identify potential risks for future disease, to enable early action for the larger portion of the population. One fruitful solution is to create an analytical model that can assess grade of disease along the continuum and to identify early stages of action [7]. This redefines the current disease-reducing strategy to a risk-reducing strategy, encouraging a wider focus on population health and a larger budget allocation to preventive health work. Population-based research will be crucial for future public health.

Information is key to preventive health care; not only knowledge about the relation between lifestyle and illness, but insight about the individual and their behaviour. At the moment, information on patients' health and lifestyle is confined both in access and applicability:

- **Information measured by healthcare professionals.** Administrative data recorded during routine health care include attendance, procedures and diagnoses. These are entered digitally into administrative systems, and are then pooled at local, regional and national levels [8].

- **Information provided in patient journals.** Patient data include medical history, journals and previous treatments, hospital visits, physiological tests and laboratory tests.

- **Information shared by patients.** Health assessment is no longer restricted to healthcare professionals, anyone can register and monitor their own symptoms, physiological functions and various activities using any number of devices and applications. For example, activity levels, heartrate, sleep rhythms, eating habits, exercise habits, emotional states, and sentiments shared in social media, to name a few [8].

Daily interactions with technology have extended to many different aspects of life, from mundane routines and work chores to entertainment and physical activity. Moreover, social media applications, such as Facebook, Instagram, Snapchat and Twitter, gather indirect data on interests, relations and emotional states through consumed content and communication patterns. Most of these behaviours are voluntary, yet the expanse of data left behind online is difficult for the casual user to comprehend. The more we do online, the larger the shadow of our online actions grows. The data shadow encompasses all unintentional, or less intentional, data that an individual generates online, including clinical data. Together, activities on different technologies and arenas make up arrays of data that share the potential to yield unique insights about an individual's life-style related behaviour.

The inherent potential in these massive data collections is left largely untapped, for several reasons. Firstly, although a lot of this information is publicly available online, little or nothing is openly accessible to healthcare professionals. Secondly, health personnel underestimate the value of self-recorded and self-generated information. And thirdly, there is a general lack of knowledge and skill in using this type of data [8]. Nonetheless, advances in data analytics demonstrate the potential in applying statistical algorithms on patient data, for instance in calculating the risk of outbreaks of infectious diseases [9]. With the power of algorithmic analyses, a system could monitor all measures and activities and offer customised health advice, which in turn could be adjusted in accordance with the resulting effect. In other words, everyday actions can deliver the insight needed to provide the still-healthy population with customised life-extending advice.

The puzzle pieces are all there, the only thing needed to complete the puzzle is to assemble the pieces. Public Health and Electronically Assessed Data (PHEAD) aims to do just that, by developing a system that collects and analyses data from mobile and wearable devices. With the potential that follows large-scale data analysis, PHEAD will enable early detection of disease and other consequences related to life-style. Fine-meshed data limits will ensure that the system can divide the continuum from completely healthy to sick in the same way that many diseases are now divided into stages. This will socially, economically and not least individually be very important for future treatment and prevention of illness and disease.

Architectural Paradigms

Through years of evolution different software architectures and paradigms have developed. From the early days of client-server to modern day serverless architectures and with many stages in between. We will further highlight some of these paradigms with particular importance.

Service Oriented Architecture (SOA) was coined as a term in 1996 and it has since become the state-of-the-art in many software architecture. All large software vendors today offer various frameworks and implementations of SOA [10] and Microsoft, Amazon, Heroku to name a few host easy configurable cloud based environments for such deployments. SOA is a framework for designing flexible and loosely integrated services, in distributed environments with the main goal to address the challenge of supporting architectural innovation and flexibility in an ever changing software development environment. SOA as an architectural principle should aim to support architectural flexibility with the ability to include both new and legacy systems into the software platform. Further, its distributed nature should support efficiency, reuse, security, adaptability and maintainability [11].

Following the work on Service Oriented Architectures other perspectives evolved as well with Microservice Software Architectures as the most predominantly one [12, 13]. Miroservices is debated and not uniformly defined. However, usually most descriptions will recognise it by design principles. The design principles [14] often include the possibility to rapidly change and adapt, that the individual units (services) are small, coherent functions of work solving one specific task and by living in a cloud based environment they out of the box fully support a highly scalable approach.

In recent years, the serverless architectural paradigm has enjoyed the attention of practitioners and researchers alike – and interest is growing [15]. From the perspective of a software developer, the serverless paradigm's main purpose is to provide an abstraction over traditional operations related to hosting and orchestration of servers, and hardware-based scaling (e.g., vertical and horizontal system scaling) [16]. In Figure 1, five common abstractions in the X as a Service umbrella are displayed, ordered from closest to the metal (left) to the furthest away (right). The serverless paradigm aligns closer to the right, typically seen as an umbrella term covering Back-end as a Service (BaaS) while incorporating the Functions as a Service (Faas) paradigm as a core enabling component for executing business logic. The choice of runtime environment for FaaS is directly tied to cost and warm and cold startup performance of the FaaS functions, challenges investigated in depth by Jackson and Clynch [17]. The BaaS paradigm is composed of a set of Software as a Service (SaaS) components, typically user management, databases, file storage and similar. For the PHEAD project, the FaaS paradigm is in focus – however, we also make use of other BaaS and SaaS components including CosmosDB for persistent data storage, and Azure Event Grid enabling our event-driven architecture. Also (federated) identity management is likely a component we will integrate with in future work.

Figure 1: Layers of abstraction for X as a Service

Persuasive Technologies

Many behavioural responses are triggered by external stimuli, be it the reply to a question or the orientation towards a sound. Traditional learning theories attribute behavioural learning to the association that follows repeated exposure to a new stimulus and a subsequent response, mediated by an established stimulus [18]. This approach to learning has been adopted by a large number of technology providers, who incorporate basic stimulus-response mechanisms in their product designs. Names may vary, but playful, motivational and persuasive designs all aim to motivate the continued use of a product [5]. For instance, gamification mechanisms such as reward schemes are implemented in many social media applications [19]. These mechanisms are typically presented as rewarding responses that provide value to a voluntary behaviour, thereby reinforcing the relation between the behaviour and the response [6]. The response can be as simple as a high-score, a badge, a running streak, or a number of views, likes or comments to a piece of content shared by a user. However, the mechanisms are not as simple as they appear. A user's behaviour and responses provide data that can be used to accommodate the user further, such as filtering out relevant content or adjusting the frequency of notifications [6]. This type of nudging can certainly get someone hooked on an application [20], but it can also be used to encourage lifestyle changes.

Fogg [21] took a conceptual view on persuasive technologies and proposed the idea that there are five primary types of social cues for people to recognize computers as social actors: Physical, psychological, language, social dynamics and social roles. Fogg [22] developed a behavioral model for persuasive design named the Fogg Behavioral Model (FBM). The model consists of three elements: Motivation, ability, and triggers.

Persuasive design of health applications incorporates a number of features to encourage users both to use the application, and to change their behaviour. The features can be grouped into four categories, social support, system credibility support, dialogue support and primary task support. These categories contain features with different mechanisms, some appeal to social needs, some use call-for-actions to trigger behaviour, and some are aimed at the individual's preferences [23]. However, none of these features consider the context or the timing of the nudge. Previous work has suggested that increased focus on the user's current context of use is of key importance for improving the user experience of the wearable. In particular, lack of context-awareness in fitness trackers has been listed as a significant drawback, for instance when the user is not able to act on a nudge motivating to move, when the user is sitting on an airplane [24].

3 Scenario

Imagine the following scenario for our user, *Laura*: After getting out of bed this morning the user opens the smartphone to read the latest news, check Facebook, e-mail, and complete the morning ritual in the online world. Whilst doing this the activity tracker on Laura's arm are synchronizing the sleep pattern to the vendor cloud via the smart phone - she gets a thumbs up for having completed a good nights sleep. On the way out the apartment, the system picks up on this and calculates the likeness for that she now will choose the bus, rather than walk, to work. A gentle nudge with a motivating suggestion are shown encouraging Laura to walk today since the exercise pattern from the previous two days are below her expected standards. Well at work she has a busy day with meetings and office time combined. The system is aware of her schedule and gives updates on the progress towards the daily activity goal in between meetings and work session to be non-intrusive. After work hours, when she has finished her round of errands, the systems has a suggestion ready for an evening activity in line with the system profile for a user of her category - she completes her power-yoga session successfully. When she is at bed for the night, again the smartphone synchronizes today's activities and information to the backend cloud. An analytics-engine is crunching the data from the day from all the registered users, detecting anomalies, identifying patterns and organizing the data storage. Where consent has been given, data is shared with health care services for individual health profile updates, sharing the newly gained knowledge. In preparation for the next day, machine learning algorithms are gaining new insight from the last 24 hours of data, and accordingly updates the individual users profile. Whilst all this has happened, Laura has almost completed a good nights sleep. She is now one hour away from waking up to a new day with personalised, individually tailored health profile activities.

4 Prototype Architecture

Instead of manufacturing our own wearables, PHEAD is using existing off-the-shelf hardware widely available to the general public. The challenge is to find a common ground to make these wearables accessible for use. Establishing a joint software platform and an Application Programming Interface (API) to collect data from wearables of different manufactures is therefore necessary. All manufactures have their own mechanisms of collecting data based on their own hardware and algorithms, hence; we will access and pull this information from the dedicated manufacturers upon given user consent.

When examining the architecture components in Figure 2, which shows the software architectural implementation with emphasis on the major components, we can identify three major components: End user component with smartphone and wearable; wearable device vendor cloud; Azure Cloud Environment. All three components are vital parts and we will further elaborate on them.

The wearable device vendor clouds represent the remote service for the individual producers such as Garmin, Polar, FitBit etc. This is the standard cloud based services

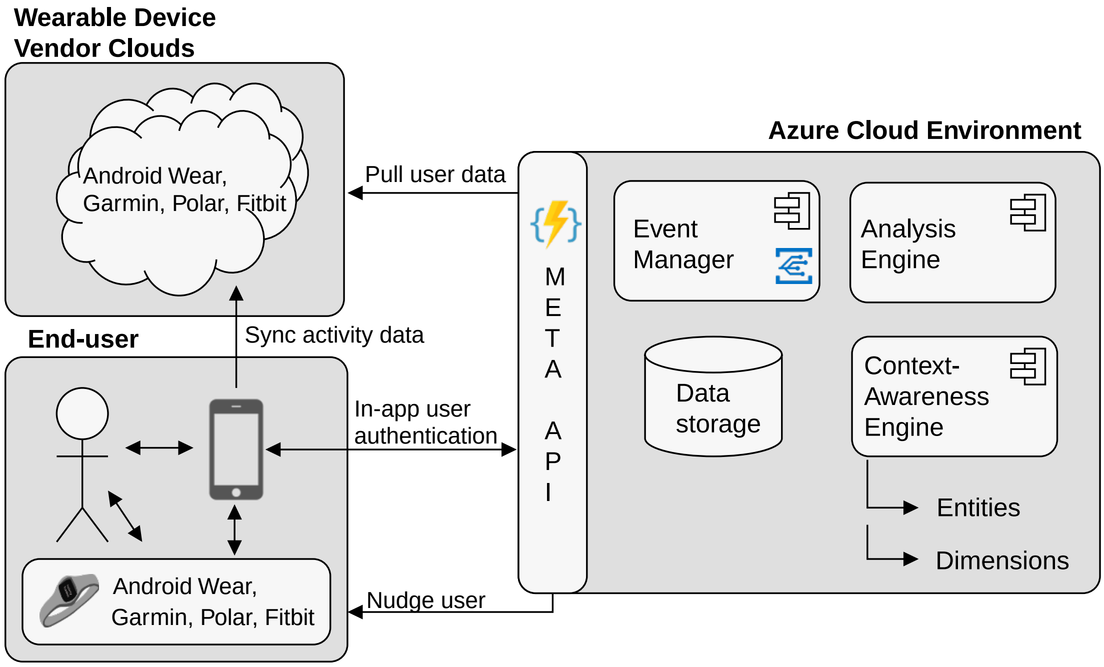

Figure 2: Conceptual PHEAD architecture

for each vendor where users typically register an account upon purchase of a new device. Vendors enable access to user data through proprietary APIs, given that a user token is acquired.

On the end-user side, PHEAD is build to expand across and to use several different platforms i.e. the central cloud based software platform, a variety of wearable devices and native applications on Android and iOS platforms. Taking a persuasive technology approach, the design of this applications incorporates a number of features to encourage users both to use the application, improve their health and to change their behaviour. The user perspective contains two major components, where the first one is the wearable device i.e. typically a watch (Garmin, Polar, Apple etc) or an activity tracker (FitBit, JawBone etc). Secondly, the users smartphone device is important as it has two main features: Firstly to act as the synchronization hub towards the vendor cloud and secondly providing the local application which is used for token generation for the backend system so that data can be fetched from the vendor cloud on the users behalf.

5 Result and Discussion

To elaborate on our findings, we will describe and discuss them in comparison with related work to our research. By doing so we will revisit it from the perspective of the three pillars of: personal health information, software architecture and persuasive technologies.

Personal Health Information challenges

Health personnel typically faces three challenges when consulting a patient; inadequate update of healthcare workers in the medical field due to limited time resources, lack of knowledge of the behaviour of the patients and an inability to change patients habits for improving health. Firstly, a GP typically sees 30 patients a day, and there is little room for self-study. As there are close to 2 million medical papers published each year, keeping updated in the field of medicine is an impossible task for a GP. Secondly, patients, intentionally or unintentionally, does not reveal all information the GP needs to perform qualified medical treatment. The information perhaps does not occur to be important to the GP nor the patient, the patients does not remember and sometimes the patients are unwilling to share important information. Thirdly, making people change to a more healthy lifestyle in order to prevent diseases has proven very difficult, and many GP's feel helpless as their advice are not embraced by the patients.

All of these three challenges may be improved by using information gathered by wearables, social media and algorithms detecting important health-related behaviour, as well as connecting the personalised information with the wast amount of medical information available on internet and conveying it to health personnel as the GP. Perhaps more importantly, health personnel may encourage health enhancing behaviour and dissuade behaviour related to poor future health. The latter by using premeditated and situation related nudging.

The challenge is to easily obtain these data, decipher the output and make appropriate use of the information to help patients. Unfortunately, the amount of data is so large it is impossible to process this in a normal manner. This, together with an increasing specialization of the medical field, the amount of information will cause problems for healthcare professionals. That is, physicians are not expected to possess the knowledge available about the individual diseases and the best treatment. The physicians and other healthcare professionals must therefore make use of the knowledge available on the internet. This information is very complex and fragmented. Using algorithms to search the internet for the latest information directly related to the patients symptoms and diseases will allow health personnel to give the latest treatment and advises to the patients with significantly more precision.

Architectural Challenges

The main architectural flow is showcased through one vendor in Figure 3 by exemplifying with FitBit. Initial dataflow is started by having the local user activity synced to the vendor cloud through a smartphone device. Given that new data is uploaded, the FitBit vendor cloud will POST to the webhook endpoint that new data is available, accordingly triggering a FaaS function to publish this event on the Azure Event Grid. Following this, another FaaS function collects data from the vendor cloud and transforms it to a local entity format for further processing. We build and run a data validation and reliability tests for the different wearables to ensure correct flow of data from user, through vendor systems, to our software system in the Azure cloud. Wearables that are not accessible or which fails to produce valid and reliable data are omitted. The approach with FaaS architecture, support continued development and addition of new functions for data processing as needed following established principles from previous work [14, 16]. The last step of data processing is the computing of the activity context for a given user. By doing this we achieve through a context-aware taxonomy to decide whether the user should be nudged using a persuasive approach to reach h/s goal, or to terminate the current operation.

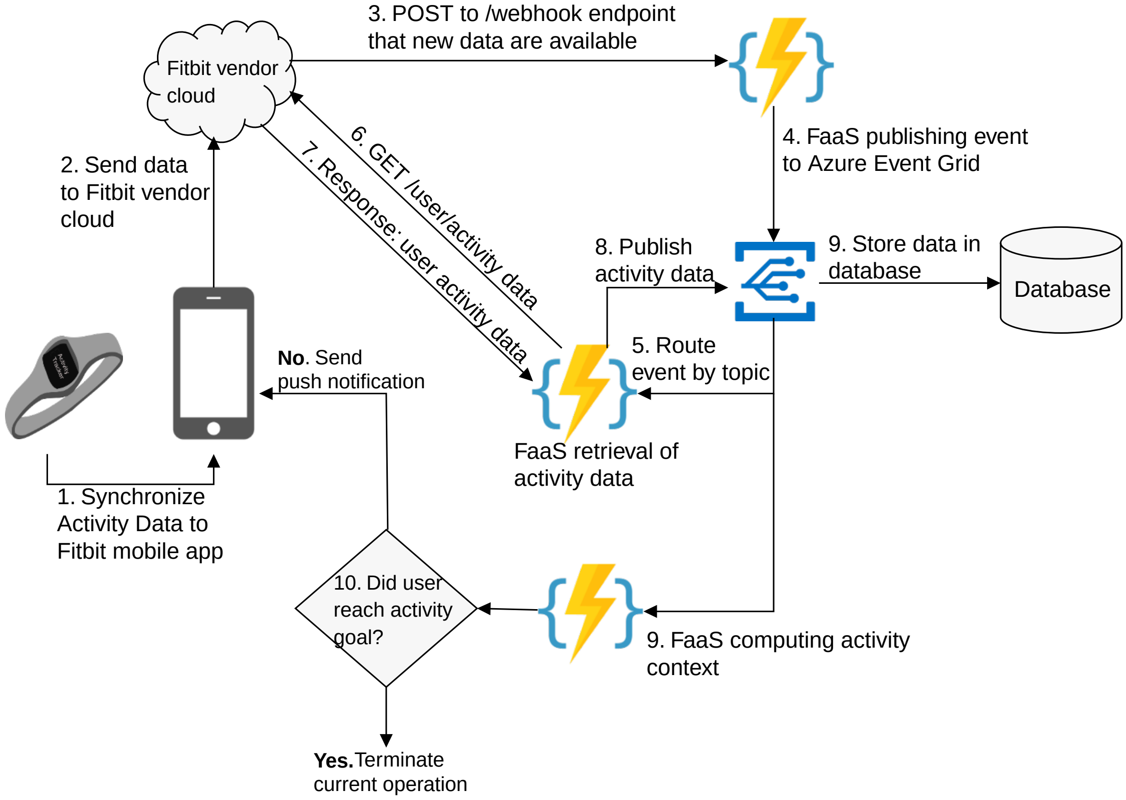

Figure 3: Fitbit PHEAD architecture

Persuasive Technology Challenges

This application function as a hub, collecting pre-planned data from the chosen applications by applying playful, motivational and persuasive designs all aim to motivate the continued use [5]. These mechanisms are typically presented as rewarding responses that provide value to a voluntary behaviour, thereby reinforcing the relation between the behaviour and the response [6]. As the main goal is to help people improve their life-style and prolong peoples life, an applicable feedback system must be developed. Telling people to change life-style on a general and not specific level has limited effect. However, human computer interaction studies have shown that so called "nudging" do have effect [23]. To encourage people to follow guidelines given based on the personalised data collection, nudging needs to be individualised and appropriate for the situation the individual is in. Small comments on a smartwatch telling you to move every hour is not very helpful, especially if you are on a plane for five hours, nor has it shown any large effects. The nudging have to be based on all the information given by the PHEAD-setup and programmed to nudge people in a setting where the nudging has an effect and people are able to follow the instructions.

Technical Challenges

The process of developing the FaaS-based PHEAD architecture is illustrated in Figure 4. One particular aspect of the process, which complicates local development and debugging of this architecture, is the absence of tooling. At the time of writing, Microsoft Azure does not provide an official emulator or local environment for their Event Grid product, a central component of the architecture, enabling routing of- and temporary storage for events. As the loosely coupled Azure Functions depend on being triggered by events published by the Event Grid, the lack of a local emulator renders development and debugging cumbersome. The issue of debugging is also raised in a whitepaper by Fox et al. [25] reporting on the status of FaaS and serverless computing both in industry and research. Based on our experience developing this system using the Microsoft Azure platform, we note that years after Fox et al. published their report, the state of debugging and local development is still work-in-progress.

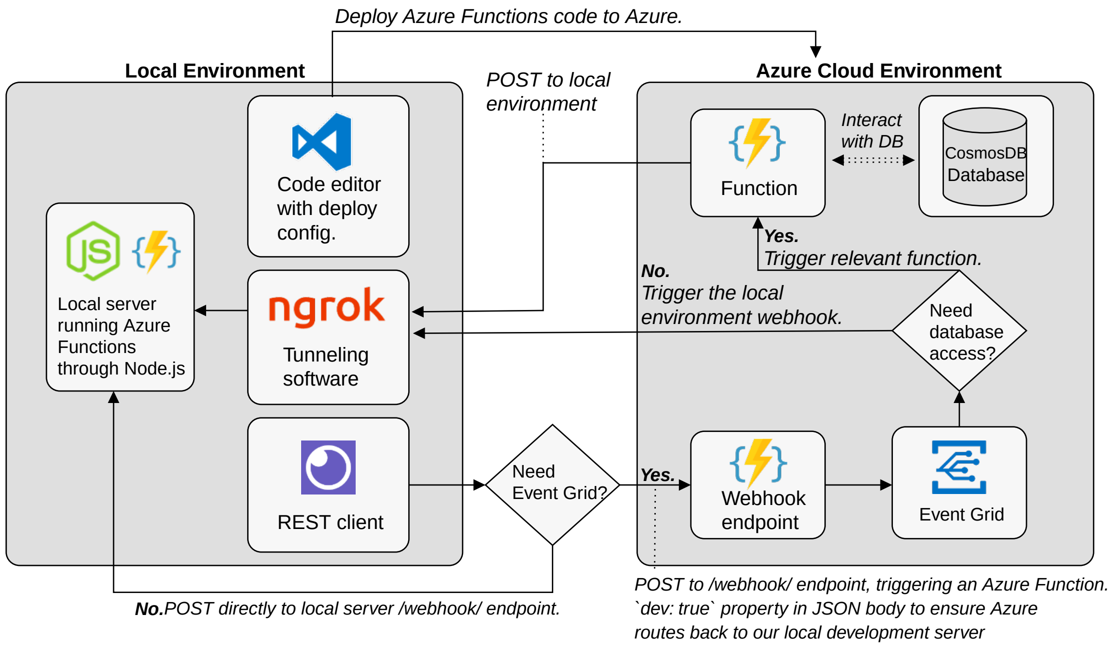

Figure 4: The process of developing PHEAD FaaS semi-locally

To route events from the Azure Event Grid back to the local development environment, the *ngrok* tunneling software is typically used to expose externally a server running on localhost, allowing the localhost server to be reached from outside the local network [26]. A RESTful API endpoint from the localhost server is then provided to the Azure Event Grid, formatted as follows after being exposed by the ngrok tunnel: https://<unique_app_id>.ngrok.io/api/endpoint. While this method of semi-local debugging is recommended by Microsoft [26], expanding the development team from one developer to two (or more) will require the Azure Event Grid to also handle event routing not only back to the local environment, but more specifically to the correct developer's environment. Setting up and configuring an online development environment, such as the upcoming Microsoft cloud-hosted development environments [27], to which multiple developers could connect to could be one solution to avoid riddling the Azure Event Grid with configuration and patterns for handling routing to multiple and distributed local development environments.

Limitations

We acknowledge that since this initial result and evaluation does not involve actual users with personal information, we do only initial verification of the suitability of the data and related security mechanisms. For further studies a robust security platform and data protection regime must be incorporated and tested.

6 Conclusion

This research project anchored in the cross-disciplinary fields of preventive health and software architecture showcased how to synthesizing personal health information from wearable devices to gain insight into the health data shadow. Through a developed prototype architecture and data flow test with the FitBit device, we show how individuals can gain a fact-based awareness of own health and make informed choices. Our proof of concept implementation further leverage the state of the art cloud based function as a service approach to build a scalable software architecture for a ubiquitous and heterogeneous environment harvesting the data shadow through activity tracking devices. Further to this, we highlight the benefit for such a solution for health workers and medical doctors that can be provided with a comprehensive, unfiltered data foundation to base diagnoses, treatment and council upon. This work on establishing a solution to unveil the data shadow show the huge potential for demystifying this huge amount of data, utilising state of the art serverless architectures and further highlight openings for data analysis, information tailoring and context-aware health information as highly relevant areas worthy of further pursue.

7 Acknowledgement

This work was supported by interdisciplinary funds from Kristiania University College. Regional Ethic Committee reference number: 2019/648/REK south-east C.

References

[1] World Health Organization. Obesity and overweight. `https://www.who.int/news-room/fact-sheets/detail/obesity-and-overweight`. (Accessed on 05/09/2019).

[2] World Health Organization. Double burden of malnutrition. `http://www.who.int/nutrition/double-burden-malnutrition/`, 2019. (Accessed on 03/17/2019).

[3] Statistisk Sentralbyra. 68 000 per innbygger til helse. `https://www.ssb.no/nasjonalregnskap-og-konjunkturer/artikler-og-publikasjoner/68-000-per-innbygger-til-helse`, 2019. (Accessed on 04/09/2019).

[4] Tor-Morten Grønli, Andreas Biørn-Hansen, and Tim A Majchrzak. Software development for mobile computing, the internet of things and wearable devices: Inspecting the past to understand the future. In *Proceedings of the 52nd Hawaii International Conference on System Sciences*, 2019.

[5] Sebastian Deterding. Eudaimonic design, or: Six invitations to rethink gamification. 2014.

[6] Ryan Tan Rui Yang and Vivian Hsueh Hua Chen. Gamification: Influencing value-perception of target behaviors. In *Proceedings of GamiFIN Conference*, 2017.

[7] Geoffrey Rose, Kay-Tee Khaw, and Michael Marmot. *Rose's strategy of preventive medicine*. Oxford University Press, 2009.

[8] Sarah R Deeny and Adam Steventon. Making sense of the shadows: priorities for creating a learning healthcare system based on routinely collected data. *BMJ Qual Saf*, 24(8):505–515, 2015.

[9] Jeremy Ginsberg, Matthew H Mohebbi, Rajan S Patel, Lynnette Brammer, Mark S Smolinski, and Larry Brilliant. Detecting influenza epidemics using search engine query data. *Nature*, 457(7232):1012, 2009.

[10] Michael Rosen, Boris Lublinsky, Kevin T Smith, and Marc J Balcer. *Applied SOA: service-oriented architecture and design strategies*. John Wiley & Sons, 2012.

[11] Haluk Demirkan, Robert J Kauffman, Jamshid A Vayghan, Hans-Georg Fill, Dimitris Karagiannis, and Paul P Maglio. Service-oriented technology and management: Perspectives on research and practice for the coming decade. *Electronic commerce research and applications*, 7(4):356–376, 2008.

[12] Cesare Pautasso, Olaf Zimmermann, Mike Amundsen, James Lewis, and Nicolai Josuttis. Microservices in practice, part 1: Reality check and service design. *IEEE Software*, 34(1):91–98, 2017.

[13] Cesare Pautasso, Olaf Zimmermann, Mike Amundsen, James Lewis, and Nicolai Josuttis. Microservices in practice, part 2: Service integration and sustainability. *IEEE Software*, (2):97–104, 2017.

[14] Martin Fowler and James Lewis. Microservices a definition of this new architectural term. *URL: http://martinfowler. com/articles/microservices. html*, page 22, 2014.

[15] Cloudability. State of the cloud 2018. Technical report, 2018.

[16] Neil Savage. Going serverless. *Communications of the ACM*, 61(2):15–16, January 2018.

[17] D Jackson and G Clynch. An investigation of the impact of language runtime on the performance and cost of serverless functions. In *2018 IEEE/ACM International Conference on Utility and Cloud Computing Companion (UCC Companion)*, pages 154–160. IEEE, December 2018.

[18] Burrhus F Skinner. Operant behavior. *American psychologist*, 18(8):503, 1963.

[19] Natasha Singer. Can't put down your device? that's by design. *The New York Times*, 2015.

[20] Nir Eyal. *Hooked: How to build habit-forming products*. Penguin UK, 2014.

[21] B Fogg. Computers as persuasive social actors, chapter 5, 2003.

[22] Brian J Fogg. A behavior model for persuasive design. In *Proceedings of the 4th international Conference on Persuasive Technology*, page 40. ACM, 2009.

[23] John Matthews, Khin Than Win, Harri Oinas-Kukkonen, and Mark Freeman. Persuasive technology in mobile applications promoting physical activity: a systematic review. *Journal of medical systems*, 40(3):72, 2016.

[24] Vivian Genaro Motti and Kelly Caine. Smart wearables or dumb wearables?: Understanding how context impacts the ux in wrist worn interaction. In *Proceedings of the 34th ACM International Conference on the Design of Communication*, page 10. ACM, 2016.

[25] Geoffrey C Fox, Vatche Ishakian, Vinod Muthusamy, and Aleksander Slominski. Status of serverless computing and Function-as-a-Service(FaaS) in industry and research. Technical report, IEEE, August 2017.

[26] Craig Shoemaker, Theano Petersen, and Nick Schonning. Azure functions event grid local debugging. `https://docs.microsoft.com/en-us/azure/azure-functions/functions-debug-event-grid-trigger-local`, October 2018. Accessed: 2019-9-6.

[27] Jonathan Carter. Intelligent productivity and collaboration, from anywhere. `https://devblogs.microsoft.com/visualstudio/intelligent-productivity-and-collaboration-from-anywhere/`, May 2019. Accessed: 2019-9-6.