A Fog-Computing architecture for Preventive Healthcare and Assisted Living in Smart Ambients

Luca Cerina *IEEE Member*, Sara Notargiacomo *IEEE Member*,
Matteo Greco *IEEE Student*, Luca Paccani *IEEE Student*,
Marco Domenico Santambrogio *IEEE Senior Member*

Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Italy,
{*luca.cerina, sara.notargiacomo, marco.santambrogio*}@polimi.it;
{*matteo3.greco, luca.paccani*}@mail.polimi.it

***Abstract*— In the last years, the emergence of pervasive connected devices and the development of the cloud computing paradigms brought a revolution in health-care and industrial applications. Cloud and Internet of Things (IoT) exploit large scale service providers to vastly reduce costs and gather Big Data. However, cloud-based services still face various issues related to: high bandwidth requirements, unpredictable delays, and security and safety concerns. These issues are critical to health-care and Active and Assisted Living (AAL) where a correct and timely reaction can result in saving a life or drastically reducing a disability (e.g. after a stroke).**

In this scenario, we present a flexible multi-level architecture using the fog approach, a computing paradigm in which heterogeneous devices at the *edge* of the network collect data, compute a task with minimal latency, and produce physical actions meaningful for the user, leveraging upon context and location awareness. In this paper, we envision also an edge node built upon Field-Programmable Gate Array (FPGA) technology. The hardware of a FPGA node can be reconfigured to produce maximum performance in tasks, to guarantee a minimal delay, or the capacity to scale on the number of devices connected, with a minimal power consumption. We present two case studies for assistive smart ambients and health applications designed on our Fog architecture.

I. Introduction

Cloud and IoT technologies represent some of the most powerful breakthroughs in computing of the last years; the aggregation of resources (e.g. storage, networking, and computation) inside large datacenters scattered around the globe, provided a new way for developers to create all different applications for the users, leveraging the amount of knowledge generated from the users' data. Today, two exabytes of data are generated each day and this reflects the requirements in terms of connectivity, computing and storage infrastructure to handle an equivalent mass of consumer services created upon IoT technologies. While IoT tech covers different industries like manufacturing, utilities, and transportation, with an expected spending of $1 trillion by 2025 [1], a Compound Annual growth rate (CAGR) of 38% and a coverage of 40% of the market in a 2020 forecast [2] [3], medical care and health-care services represent one of the most attractive fields for the development of IoT. Novel healthcare services are expected to reduce costs, and increase the quality of life of users. Main applications regard remote health monitoring, tailored well-being and treatment of chronic diseases, and most of the data is generated from biometric sensors, wearable or not, and from Electronic Health Records (EHR) services such as Microsoft *HealthVault* or Philips *HereIsMyData*.

Nevertheless, the impact of IoT in healthcare is yet to take off due to various issues. A crucial point is the sensors ability to provide medical-grade measurements for health status assessments. A fast response time is another key feature to improve therapeutic outcomes and boost service levels, keeping at the same time a high level of safety and privacy on data, plus quality services built upon pervasive sensing devices and users' active learning.

However, current cloud models do not seem to be the best solution to handle IoT challenges, since high bandwidth constraints, network infrastructure dependency and unpredictable response time from the cloud make it unsuitable for critical applications. Given these requirements, a good alternative resides in the Fog computing paradigm. This approach is complementary to the Cloud, and it is designed to harness all the computing power from the source of data to the cloud servers. The distribution of computing is beneficial, as it offloads different tasks and processes (e.g. security policies) on all the components of the network, with lower requirements on bandwidth and storage capabilities.

By processing data at the network edge, near to the source, it can yield shorter response time and more efficient processing while keeping sensitive data inside the network, enhancing security of the users. Given that reliability is mandatory in healthcare, Fog services can be also configured to gracefully degrade their functionality in absence of cloud, or even work in a totally independent manner, while Cloud services could suffer outages due to network failures or denials-of-service.

In this paper we present a multi-level Fog architecture composed by heterogeneous devices, specifically designed for the deployment of preventive healthcare applications, using an ambient quality control system and an emergency AAL application to demonstrate its capabilities.

Additionally, we envision the usage of FPGA-based Fog nodes. Reconfigurable Systems-on-Chip (SoC) represent a novelty inside the Fog approach, useful to enhance computing performance, provide a more scalable and fault tolerant

network, and achieve high power efficiency on a broad variety of tasks.
The main objectives of this paper are:

- to introduce the benefits of fog architectures in healthcare, and active aging
- to show the qualities of FPGA-based devices as efficient Fog nodes
- to show the role of humans as active components inside a smart ambient

The rest of the paper is organized as follows: Section II presents a bird's eye of view on the current status of Fog technologies, present challenges and fog-enhanced health applications. Section III describes in detail the proposed architecture, with a focus on the capabilities of the *accelerated* Fog node. Section IV presents a set of applications that leverage the proposed architecture to maximize outcomes. Lastly, Section V outlines the achievements obtained and the future developments.

II. Related works

This section describes technological aspects, issues and challenges present in literature regarding the usage of Fog computing.

The review of [4] operates a useful distinction between edge computing, which facilitates the computation at the nearest network hop from the user device, and fog computing, which harness computing power across the entire data path. It also provides exhaustive definitions of the nodes structure and computing models. Conversely, [5] describes Fog technologies in terms of expected capabilities of the system, application requirements, and runtime capabilities, highlighting current issue in terms of security, Fog reliability and management.
The challenges of Fog computing are thoroughly described in literature. The work in [6] focuses on scheduling issues, using custom policies from Fog nodes to the cloud to minimize latency and optimize network usage. Similarly, [7] shows the transition from *rudimentary* IoT to Fog in terms of scalability and reliability, proposing an orchestration model to handle component deployment and resource planning at runtime. Lastly, [8] covers security issues and possible solutions such as trust model authentication, detection of rogue nodes and certificate revocation schemes. While the preceding works hypothesize various healthcare applications enhanced by Fog techniques, it was difficult to find studies that propose health applications as the main topic (see the review on smart cities and healthcare of [9] for supplementary details). In [10] is presented a smart gateway to extract features from an Electrocardiographic (ECG) trace based on an OMAP4 platform acting also as local storage database. The improvements in latency from 3.5 to 48.5% depending on the network data rate conditions. An extensive usage of Fog is presented in [11] to prevent Zika virus outbreaks: a combination of user-generated data and the breeding density of mosquitoes are collected from local nodes with known location, then aggregated to the cloud to determine risk prone areas.
As the definition of fog-node is generally broad, [12] implements a two-step fall detector with a lightweight classifier on a smartphone (the node) and a more accurate classifier on Cloud. The system ensures high sensitivity combined with a very low response time. An interesting approach is described in [13], aimed at the assistance of visually impaired subjects during grocery shopping: the system uses Fog devices to locate the user inside the aisles and to offload the video processing necessary to assist the subjects. Moreover, it employs state-of-the-art techniques like FPGA-accelerated Cloud, convolutional neural networks and neuromorphic architectures for maximum power efficiency.
In the context of emergency management, where timeliness is a fundamental factor, fog computing can play a pivotal role. The study [14] presents Emergency Help Alert Mobile Cloud (E-HAMC), a fog infrastructure smartphone based service designed for optimizing the process of emergency notification. Comparing fog and cloud communication between two different testing scenarios, it results that fog infrastructure leads to a positive impact in terms of efficiency. In non-Fog scenarios, where data has to be communicated from end-nodes directly to the Cloud, delay in communication is up to 6 times higher than in the fog-scenario.
A crucial issue in IoT healthcare is the development of a personal data capturing system able to acquire data in an unobtrusive, low-cost and low power manner. At a framework level, [15] embraces different technical features focused on proper data interpretation, parameters recognition, and behavioural changes used to define interventions. The sensor network is composed by inertial wearable devices and different ambient meters. Bluetooth Low Energy (BLE) and GPS technology are sued respectively for indoor and outdoor location monitoring. Collected data are aggregated locally, and then sent to a central repository where they are analyzed applying Big Data methods. In [16] is addressed the problem of daily monitoring in Telehealth, answering the needs of two specific disease patient groups providing innovative sensor measures, tailored management processes and risk stratification strategies. Finally, [17] presents a remote monitoring system, which relies on both wearable devices and sensor embedded in the ambient, aimed to detect anomalies in daily activities and support active aging.
While FPGAs are already massively employed in telecommunications infrastructures to provide high reliability and low upgrade costs, this is the first work in literature to our knowledge that provides a FPGA-based, application specific approach inside the Fog landscape.

III. Proposed solution

This section presents our approach for the design of the fog architecture and how specific challenges for a reliable and scalable system.

In the edge computing approach, we can see how taking advantage of the location awareness of Fog nodes bring various benefits such as:

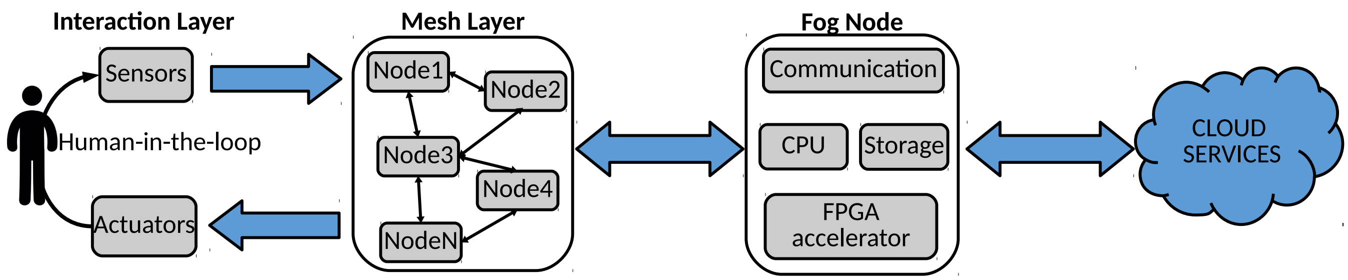

Fig. 1: Representation of the proposed architecture. The Interaction Layer collect data and receive commands. The Mesh Layer aggregate data, transfer commands and handle the security of lower nodes. The Fog Node process data to act on the ambient and possibly transfer information to the Cloud.

- A lower usage of bandwidth, prominent in all the situations in which connection could be limited (elderly houses, rural areas, and emergency vehicles)
- Minimized latency, higher context awareness and easier manageability
- Increased reliability and security, with flexible policies for data transfer and encryption
- The separation of Big Data analytics (useful for research on population and policy makers) and tailored *action items* that directly empowers the final user

The proposed system is shown in Figure 1. The lower layer is built upon three elements: *Sensors* that produces data, *Actuators* that consume it in form of commands to generate a physical outcome, and the *Human-in-the-loop* that can produce and consume data, and act upon it. The data is exchanged using proven BLE standards: sensor nodes exploit advertisement for low-security data (e.g. ambient temperature) or direct ad-hoc connections for sensitive data. The communication is handled from the Mesh Layer, which implement also a first security level, authenticating the nodes below. In our current implementation, the mesh nodes create point-to-point connections at runtime with the nearby nodes to increase security and optimize routing. Future implementations could rely on the location knowledge of nodes to achieve minimal latency transmission with optimal routing. Since this layer adds small, yet unwanted latencies to the transmission, its dimension is dependent on the specific application and it can be potentially reduced to zero for time-critical applications.

The higher layer is represented by the proper Fog node. From a networking point of view it sees the Mesh layer as a peer, collecting data and sending commands to actuators, and at the same time it acts as an access point to the internet and the cloud services. In our specific approach it consists in a Xilinx Pynq-Z1 development board, designed with IoT in mind, which combines an ARM Cortex A9 CPU and a Zynq XC7Z020 FPGA. The utilization of FPGA based systems as Fog nodes brings many advantages including:

- Reconfiguration of the hardware tailored on the specific application
- High performance in data manipulation and signal processing
- Low power, comparable with other board computers like Raspberry PI or BeagleBoard

The combination of these factors means that the node can be designed to handle some tasks (like data filtering or pattern recognition) in the most efficient way employing hardware accelerated algorithms. In the Fog paradigm it should lead to more responsive or more scalable systems, since it can handle an higher amount of information in lower time compared to a classical CPU. The details on how the FPGA improves the processing task is dependent on the application and it will be explained in the following section.

The data from the fog nodes is aggregated, according to the necessities of the application, and then transmitted to a central server. In the two case studies presented the system can operate independently from the server, but in many other applications the presence of a high-scale computing system is necessary to store and process global knowledge from the underlying fog structure. For instance, the Cloud can be employed to store information about patients and generate long-term analytics or predict risk stratification; after that, actuators at the lower level can adapt the ambient to increase the quality of life of patients.

With regard to the location awareness, the Fog architecture can be used to build models enhanced by this information, or to produce specific outcomes (like the location of injured subjects or the movement of crowds inside an ambient). Additionally, the knowledge about the position of elements can be used to modulate the distribution of resources (intended as sensors or computing nodes) optimally. In our implementation, the location of fixed objects can be defined at design time and the information can be shared between the fog nodes and the cloud. Objects deployed after that can be added to the location database.

IV. Case Studies and Results

This section introduces two case studies built upon our architecture. The first shows how a smart ambient application with a small sensor-actuator-human loop could be used to improve the air quality in working environments, with beneficial results on health and productivity. The other case study proposes an application for the recognition of emergency situations in elderly (fall and/or heart failures) and use FPGA-based node to process biological signals and generate an alarm with minimal latency.

Light & CO_2

Given the amount of time that humans spend indoor, the built environment plays a critical role in our overall well-being and can influence our health. In particular every day insides offices, houses, classrooms, people are routinely exposed to high levels of carbon dioxide (CO_2) that has direct and negative impact on human cognition and decision-making processes [18].

CO_2 indoors is mostly produced by human respiration, which in closed and lowly ventilated spaces, leads to an increase in CO_2. Therefore, CO_2 is a good indicator of proper building ventilation and indoor air exchange rates, and it can be measured to determine if the indoor air is adequate for humans to occupy the building.
In terms of worker safety, Occupational Safety and Health Administration (OSHA) has set a permissible exposure limit (PEL) for CO_2 of 5,000 parts per million (ppm) over an 8-hour work day. However, according to [19], human cognitive functions decline by 15% for a moderate CO_2 daily amount ($\sim$945 ppm) and by 50% with CO_2 concentrations of $\sim$1,400 ppm. For these reasons a sufficient amount of fresh air of appropriate quality must be periodically provided in crowded zones, to reach a sufficient comfort.

To investigate these phenomena we installed a CO_2 sensor in the laboratory meeting room (which is equipped with a ventilation system but void of windows) to study the CO_2 concentration variation during the day, also taking into account the activities that took place at that time (e.g. meetings). We used a Cozir CO_2 sensor with a sensitivity of 100 ppm and a range of 0 to 20% in ppm. Additionally, we installed a WiFi connected light bulb (LIFX Color 1000) to provide visual feedback to the people inside the room, and to point out to what concentrations of CO_2 they were exposed at that particular time. The system is programmed in such a way that the light displays a different color for each CO_2 range:

- 400 to 800 ppm $\rightarrow$ Green: CO_2 is perfectly bearable by a human being;
- 800 to 1700 ppm $\rightarrow$ Orange: it is recommended to open the room's door or the near window to allow the air exchange;
- 1700 to 2600 ppm $\rightarrow$ Red: it is strongly recommended to open the room's door and the near window to reduce the CO_2 concentration.

As it can be seen in the Figure 2, we obtained a great improvement in terms of air quality after the light bulb installation. In fact, it allowed people in the room to become aware of the CO_2 level in the air, and motivate their actions appropriately. This improve people active learning, and it helps to avoid the same high and distressful CO_2 concentrations that were measured before the light bulb installation, as it can be seen in Table I.

Given that the environmental conditions in every room, intended as the combination of the CO2 levels and the state

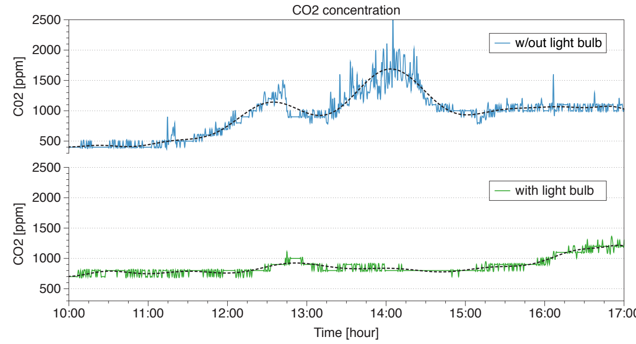

Fig. 2: Trends of CO_2 concentration before (top) and after (bottom) the installation of the light bulb inside the laboratory room. In the *before* graph are visible the curves caused by meetings during the day.

TABLE I: Statistical indicators before and after the light bulb installation. SD is standard deviation.

CO_2 [ppm]	**Median**	**SD**	**Max**
w/out light	1000.0	380.1	2600.0
w/ light	900.0	155.4	2100.0

of all lights, must be checked, the ideal solution seems to be the introduction of an intermediate layer in the area. The fog node can use a minimal amount of resources to analyze the data from sensors and switch the lights accordingly, with limited bandwidth consumption, no computation on cloud servers (which can then re-elaborate the data transmitted) and no delay between detection and execution. The FPGAs provides a flexible, standard based solution that combines software programmability, real-time processing, hardware optimization and any-to-any connectivity with the security and safety (detection and execution performed locally) needed by Industrial IoT systems.

Fall and cardiovascular shock assisted living

Here is presented a proof-of-concept of health monitoring system, which includes wearable and living environment embedded sensors, able to detect accidents or sudden irregularities in heart rate and automatically request assistance, specifying user's position. To prevent cardiac issues with an automatic system is relevant because many people couldn't recognize stroke events within an hour of symptoms onset, and fewer than a half of emergency callers are able to recognize the symptoms during the stroke event [20].

The system is built using two types of sensors: a commercial wearable wristband worn by the user and equipped with inertial sensors and a heart rate sensor, plus a set of Bluetooth beacons deployed in the user living place. Additionally, in the current version, the user needs to use a smartphone to retrieve the beacon signals and communicate its position.
The system is trained with common daily activities, including rest sessions, meal times and sleep, to generate an Activity Index (AI). The sampled data are separated and analysed

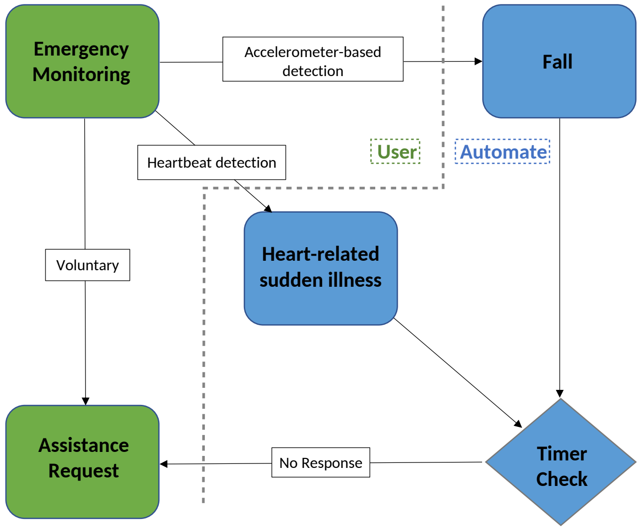

Fig. 3: Representation of unconsciousness detection model. In Green the processes that are embedded in the wearable system and depend from the user interaction, in Blue the logic that runs on the smartphone (timer check) or the Fog node

to create a baseline model. The singular observations, including the current AI, allow the system to build a Hidden Markov Model (HMM) on the HR. HMM is a state machine composed by two main components: the observable states (the measure of HR), and the hidden states (subject eating, working, unconscious, asleep) connected by transition probabilities. The user state variation is expected to be relatively slow (eg. from working to eating and vice-versa), hence a sudden transition could mean that an anomaly has happened. The tolerance of the system to false detections is ensured by multiple checkpoints in which the user has to interact with the system.

User position is estimated by BLE beacons and a mobile app which packs beacons' signal strength before sending data to a Fog node. In case of emergency, an assistance request is sent together with injured person position. This system could also work in remote areas where internet access is limited, because computational nodes are on site. The onboard FPGA is currently configured with soft-processors (low-power CPUs that are replicated on the FPGA fabric): two that perform HR and activity analysis, and another one that integrate the data in the HMM machine. For example, heartbeat anomalies or stress development could be detected not only in elderly subjects, but also in factory workers or long distance drivers, suggesting them to stop. Currently, the usage of a commercial wearable does not allow us to estimate important parameters, such as beat-to-beat Pulse Rate Variability (PRV), which could be useful to better estimate the cardiovascular stress level. (See [21] for HRV/PRV parameters of interest in job-related cardiovascular distress).

Since our system exploits very sensitive data (indoor location, biometric measurements), the system has to provide high fidelity in terms of privacy. That is guaranteed by the Fog architecture, using nodes with software *privacy mediators* that denature and enforce privacy-policy on the sensor streams [22]. Fog also permits to easily adapt the system according to the amount of users. If the number of subjects increase, the system can be upgraded with numerous nodes in the Mesh layer, but this increments also network's delay. However, the introduction of a new Fog node will reduce this delay to an acceptable level. The use of FPGA as Fog node increments both computational and power efficiency so that it has marginal impact on the structure power consumption where the system is installed. As FPGA is highly reconfigurable, it could be also reprogrammed in case of new biological signals, better wearable devices, or novel and improved algorithms to process the signals.

V. Conclusions and Future Works

In this paper, a novel architecture based on fog paradigm and enhanced by the usage of reconfigurable high performance nodes based on FPGA technology. The main objectives of this paper were to acknowledge how the Fog and reconfigurable systems can provide various benefits in the field of healthcare applications. With respect to the state-of-the-art, this is the first work to employ FPGAs in the Fog at both infrastructure and application level, highlighting how the most recent FPGA programming paradigms could be exploited to provide rich Fog applications with maximum power efficiency. In the first case we showed how the architecture can be employed to enforce positive behaviours and raise awareness on the quality of air in a working environment, using a co2 sensor and a visual feedback. The results shown that the system was able to function reliably during all the experimentation and that it was able to reduce the CO_2 levels in the environment without adding mechanical elements or affecting the energy impact of the building, and positively affecting the awareness of users inside the test environment.

The second system is aimed to monitor subjects at the risk of fall and cardiovascular issues (e.g. elderly or construction workers), and to issue an emergency call in case of necessity, leveraging on FPGA performance to provide a scalable and reliable service. The current system still requires a smartphone, which we aim to substitute with a general *thing* or the wearable device worn by the user, and more work must be done to find a wearable device with the necessary medical-grade requirements. Upon that, new algorithms that exploit Heart rate Variability features to detect critical situations will be embedded on the FPGA node.

Yet, the field of Fog computing is still research-intensive and various issues are to be tackled properly, in particular on the design of fog infrastructures. One point regards the ability to model the workload of the Mesh layer, which plays a critical role in the efficiency and reaction speed of the system. Simulation methodologies and tools will be developed to optimize its components at design-time. Another crucial point is the distribution of roles among the Cloud and the Fog nodes, which subsequently determine the amount of computing resources that must be deployed across the infrastructure.

We envision that FPGA-based nodes could represent a major driver for the implementation of Fog technologies in the future, with major benefits in terms of energy consumption, cost reduction, and systems' reactivity.

Acknowledgements

This paper has been funded by *Fondazione Cariplo* (www.fondazionecariplo.it/en) under the project "Progetto Cariplo MORPHONE 2016-1010: A Challenges Driven Design for Effective and efficient Autonomic Mobile Computing Architectures".

REFERENCES

[1] J. Manyika, M. Chui, P. Bisson, J. Woetzel, R. Dobbs, J. Bughin, and D. Aharon, "Unlocking the potential of the internet of things," *McKinsey Global Institute*, 2015.

[2] marketsandmarkets.com, "IoT Healthcare Market by Components (Medical Device, System and Software, Service, and Connectivity Technology), Application (Telemedicine, Workflow Management, Connected Imaging, Medication Management), End-User - Global Forecast to 2020," Tech. Rep., 2015. [Online]. Available: http://www.marketsandmarkets.com/Market-Reports/iot-healthcare-market-160082804.html

[3] H. Bauer, M. Patel, and J. Veira. (2016) The internet of things: Sizing up the opportunity. [Online]. Available: http://www.mckinsey.com/industries/high-tech/our-insights/the-internet-of-things-sizing-up-the-opportunity

[4] B. Varghese, N. Wang, D. S. Nikolopoulos, and R. Buyya, "Feasibility of fog computing," *arXiv preprint arXiv:1701.05451*, 2017.

[5] P. Varshney and Y. Simmhan, "Demystifying fog computing: Characterizing architectures, applications and abstractions," *arXiv preprint arXiv:1702.06331*, 2017.

[6] Z. Hao, E. Novak, S. Yi, and Q. Li, "Challenges and software architecture for fog computing," *IEEE Internet Computing*, vol. 21, no. 2, pp. 44–53, 2017.

[7] Z. Wen, R. Yang, P. Garraghan, T. Lin, J. Xu, and M. Rovatsos, "Fog orchestration for internet of things services," *IEEE Internet Computing*, vol. 21, no. 2, pp. 16–24, 2017.

[8] A. Alrawais, A. Alhothaily, C. Hu, and X. Cheng, "Fog computing for the internet of things: Security and privacy issues," *IEEE Internet Computing*, vol. 21, no. 2, pp. 34–42, 2017.

[9] A. Solanas, C. Patsakis, M. Conti, I. S. Vlachos, V. Ramos, F. Falcone, O. Postolache, P. A. Pérez-Martínez, R. Di Pietro, D. N. Perrea *et al.*, "Smart health: a context-aware health paradigm within smart cities," *IEEE Communications Magazine*, vol. 52, no. 8, pp. 74–81, 2014.

[10] T. N. Gia, M. Jiang, A.-M. Rahmani, T. Westerlund, P. Liljeberg, and H. Tenhunen, "Fog computing in healthcare internet of things: A case study on ecg feature extraction," in *Computer and Information Technology; Ubiquitous Computing and Communications; Dependable, Autonomic and Secure Computing; Pervasive Intelligence and Computing (CIT/IUCC/DASC/PICOM), 2015 IEEE International Conference on*. IEEE, 2015, pp. 356–363.

[11] S. Sareen, S. K. Gupta, and S. K. Sood, "An intelligent and secure system for predicting and preventing zika virus outbreak using fog computing," *Enterprise Information Systems*, vol. 0, no. 0, pp. 1–21, 2017.

[12] Y. Cao, P. Hou, D. Brown, J. Wang, and S. Chen, "Distributed analytics and edge intelligence: Pervasive health monitoring at the era of fog computing," in *Proceedings of the 2015 Workshop on Mobile Big Data*. ACM, 2015, pp. 43–48.

[13] P. Zientara, S. Advani, N. Shukla, I. Okafor, K. Irick, J. Sampson, S. Datta, and V. Narayanan, "A multitask grocery assistance system for the visually impaired smart glasses, gloves, and shopping carts provide auditory and tactile feedback," *IEEE CONSUMER ELECTRONICS MAGAZINE*, vol. 6, no. 1, pp. 73–81, 2017.

[14] M. Aazam and E.-N. Huh, "E-hamc: Leveraging fog computing for emergency alert service," in *Pervasive Computing and Communication Workshops (PerCom Workshops), 2015 IEEE International Conference on*. IEEE, 2015, pp. 518–523.

[15] P. Paolini, N. Di Blas, S. Copelli, and F. Mercalli, "City4age: Smart cities for health prevention," in *Smart Cities Conference (ISC2), 2016 IEEE International*. IEEE, 2016, pp. 1–4.

[16] H. Reiter, A. Tesanovic, and A. Martinez-Romero, "Heartcycle: From insights to clinically evaluated ict solutions for telehealth," in *Engineering in Medicine and Biology Society (EMBC), 2013 35th Annual International Conference of the IEEE*. IEEE, 2013, pp. 6992–6995.

[17] L. Pigini, G. Bovi, C. Panzarino, V. Gower, M. Ferratini, G. Andreoni, R. Sassi, M. W. Rivolta, and M. Ferrarin, "Pilot test of a new personal health system integrating environmental and wearable sensors for telemonitoring and care of elderly people at home (smarta project)," *Gerontology*, vol. 63, no. 3, pp. 281–286.

[18] U. Satish, M. J. Mendell, K. Shekhar, T. Hotchi, D. Sullivan, S. Streufert, and W. J. Fisk, "Is co2 an indoor pollutant? direct effects of low-to-moderate co2 concentrations on human decision-making performance," *Environmental health perspectives*, vol. 120, no. 12, p. 1671, 2012.

[19] J. G. Allen, P. MacNaughton, U. Satish, S. Santanam, J. Vallarino, and J. D. Spengler, "Associations of cognitive function scores with carbon dioxide, ventilation, and volatile organic compound exposures in office workers: a controlled exposure study of green and conventional office environments," *Environmental Health Perspectives (Online)*, vol. 124, no. 6, p. 805, 2016.

[20] E. C. Jauch, J. L. Saver, H. P. Adams, A. Bruno, J. B. Connors, B. M. Demaerschalk, P. Khatri, P. W. McMullan, A. I. Qureshi, K. Rosenfield, P. A. Scott, D. R. Summers, D. Z. Wang, M. Wintermark, and H. Yonas, "Guidelines for the early management of patients with acute ischemic stroke," *Stroke*, vol. 44, no. 3, pp. 870–947, 2013.

[21] T. Föhr, A. Tolvanen, T. Myllymäki, E. Järvelä-Reijonen, K. Peuhkuri, S. Rantala, M. Kolehmainen, R. Korpela, R. Lappalainen, M. Ermes, S. Puttonen, H. Rusko, and U. Kujala, "Physical activity, heart rate variability–based stress and recovery, and subjective stress during a 9-month study period," *Scandinavian Journal of Medicine and Science in Sports*, vol. 27, no. 6, pp. 612–621, 2017. [Online]. Available: http://dx.doi.org/10.1111/sms.12683

[22] M. Satyanarayanan, "The emergence of edge computing," *Computer*, vol. 50, no. 1, pp. 30–39, Jan 2017.