Engineering IoT Systems through Agent Abstractions - Smart Healthcare as a Case Study

Eloisa Vargiu[1], Franco Zambonelli[2]

Abstract The augmenting number of elderly people is part of a demographic change that will have an enormous impact on the society in the next few years. Thus, with the final goal to give remote support to elderly people at their home, intelligent solutions that rely on Internet of Things have been proposed in the literature. Those solutions are typically able to automatically send alarms in case of anomalies, putting in contact the end-user with her/his GP or alerting the emergency center or familiars, according to the specific needs. However, although the great deal of worldwide researches in the area of the Internet of Things and its applications to several fields, the technologies to apply it in the real world are far form being assessed. In this paper, we propose novel concepts that rely on synthesizing the common features of existing proposals and application scenarios, and on bringing in the lessons of agent-based computing and agent-oriented software engineering. Being focused in teleassistance, how those abstract concepts have to be implemented in this specific scenario is presented and discussed.

1 Introduction

By 2025, about one-third of Europe's population will be aged 60 years and over –with a particularly rapid increase in the number of people aged 80 years and older. This demographic change will have an enormous impact on the society. Thus, at European and International level, there are powerful arguments for investing in health and well-being as an objective in its own right, but also because the intrinsic relationship to economic growth and competitiveness.

[1] Eurecat Technology Center - eHelth Unit - Spain e-mail: eloisa.vargiu@eurecat.org
[2] Dipartimento di Scienze e Metodi dell'Ingegneria - Università di Modena e Reggio Emilia - Italy e-mail: franco.zambonelli@unimore.it

Elderly people aim to preserve their independence and autonomy at their own home as long as possible. However, as they get old the risks of disease and injuries increase making critical to assist and provide them the right care whenever needed. Unfortunately, neither relatives, private institutions nor public care services are viable long-term solutions due to the large amount of required time and cost. In that scenario, several technological solutions have been proposed and adopted: self-management systems for user's empowerment [21]; training systems for elderly people and their family for a better and participatory life [30]; social networks to avoid solitude [31]; and teleassistance to remote give assistance and support [8]. Focusing on the latter category, a great work has been done to move from pure reactive, unlinked, and not-integrates solutions, to proactive, linked, and integrated-in-the-hospitals solutions.

Early teleassistance efforts were structured mostly as pilot projects that were small in sample size and proof-of-concept in nature, yet they demonstrated that some treatment techniques and rehabilitation assessment could be delivered to end-users located in physically separate locations, thus overcoming obstacles of distance and lack of access to trained providers. In some of the first teleassistnce projects, clinicians used the telephone to provide client follow-up and caregiver support, and to administer end-user's self-assessment measures [29, 19]. By the late 1980's, this approach expanded to include the use of closed-circuit television and pre-recorded video material to provide visual interaction with end-users [32, 33]. Solutions belonging to this first generation of teleassistance systems were totally reactive because assistance was given only in response to a user's request. Moreover, the majority of them only linked to emergency numbers and not directly to GPs of users. Furthermore, they not allow updating and synchronizing the electronic medical records of the involved end-user, thus do not providing integration with the hospital system.

With the advent of low-cost technology, the second generation of teleassistance systems introduced sensors, wearables, and devices (as for instance the panic button) at end-user's homes [15]. In so doing, a direct contact between the end-user and her/his GP is established. Nevertheless, those solutions are still reactive and no integration with the hospital is provided.

In the recent years, novel pro-active solutions have been proposed. In particular, being the final goal to automatically keep update the electronic medical record in an integrated care perspective, intelligent solutions that rely on nets of sensors and devices have been proposed (third generation of teleassistance) [24]. Those solutions are able to automatically send alarms in case of anomalies, putting in contact the end-user with her/his GP or alerting the emergency center or familiars, according to the specific needs. In particular, there are several of efforts to utilize solutions based on Internet of Things (IoT) for monitoring elderly people, most of which target only certain aspects of elderly requirements from a limited viewpoint [14, 7, 3, 2, 11, 13].

However, despite the great deal of worldwide researches in the area of the IoT [17] and its applications to several fields [5] –such as teleassistance [26], independent living [12] and, more generally, healthcare [23]– the technologies to apply it in the real world are far form being assessed. In the direction of proposing general characteristics related to the engineering of complex IoT systems and applications

–focusing on real needs identified in the specific field of teleassistance– in this paper we propose novel concepts that rely on synthesizing the common features of existing proposals and application scenarios, and on bringing in the lessons of agent-based computing and agent-oriented software engineering [35]. The so analyzed common characteristics are then used to identify some key software engineering abstractions around which the process of developing IoT systems and applications could revolve. Such abstractions –due to the inherent presence in IoT systems and applications of autonomous and goal-oriented behaviors– can be used to define a guideline for IoT-oriented software engineering. The availability of that guideline is then assessed considering the teleassistance scenario with particular reference to how the full set of abstractions may be realized there.

2 IoT systems design and development

The definition of general software engineering principles requires identifying the general features and issues that characterize most current approaches to IoT systems design and development.

2.1 Things

The "things" in the IoT vision may encompass a large number of physical objects, and also include places and persons.

Physical objects and places can be made trackable and controllable by connecting them to low-cost wireless electronic devices. At the lower end of the spectrum, RFID tags or bluetooth beacons, based on low-cost and short-range communication protocols, can be attached to any kind of objects to enable tracking their positions and status, and possibly to associate some digital information with them. More advanced devices integrating environmental or motion sensors (i.e., accelerometers) can detect the present and the past activities associated with objects or with some place. In addition, one can make objects actuable –enabling the remote control of their configuration/status via proper digitally-controller actuators– and possibly autonomous –delegating them of autonomously direct their activities. In this perspective, autonomous robots and autonomous objects [1] are components that will increasingly populate the IoT universe.

Concerning persons, other than simply users of the technology, they can also be perceived at first-class entities of the overall IoT vision. Simply for the fact of having a mobile phone or a wearable device, they can be sensed in their activities and positions, and they can be asked to act in the environment or supply sensing.

2.2 Software Infrastructures

To make "things" usable and capable of serving purposes, there is need of software infrastructures (that is, of IoT middleware [25]) capable both of supporting the "gluing" of different things and of providing some means for stakeholders and users to access the IoT system and take advantage of its functionalities.

Concerning the "glue", there are *interoperability* issues, to enable a variety of very heterogeneous things to interact with each other, via a set of common name spaces, uniform communication protocols and data representation schemes; and *semantic* issues, because a common semantics for concepts must be defined to enable cooperation and integration of things. For both these issues, however, a large body of proposals (dating back to the early years of IoT research) exists. Thus, for our purposes, we assume the existence of proper technical solutions.

Rather, key open "gluing" issues of relevance for software engineering include *discovery, group formation, and coordination*. IoT systems functionalities derive from the orchestrated exploitation of a variety of things, possibly involving a variety of users and stakeholders. Thus, it implies to discover and establish relations among things, between things and humans, and coordinating their activities also accounting for their social relations [4]. Clearly, for the above coordination mechanisms to work, *context-awareness* and *self-adaptation* are required. In fact, the inherent ephemerality, unreliability, and mobility of system components makes it impossible to anticipate which things will be available and for how long during their exploitation. This requires mechanisms for discovery, group formation, and coordination are that are capable of dynamically self-adapting to the general context in which they act, or possibly even self-organize in a context-aware way [20, 36].

Concerning the "access" to the functionalities and capabilities of individual things by users, the scene is currently dominate by the so called "Web of Things" (WoT) vision [16]. The idea is to expose services and functionalities of individual things in terms of REST services, enabling the adoption of assessed web technologies as far as discovery of things and provisioning of coordinated group services are concerned. Concerning middleware infrastructures, a variety of proposal to support the provisioning of IoT services and applications have appeared [34, 6, 25]. Beside their specificities, most of these proposals rely on: some basic infrastructure to support the WoT approach (i.e., to expose things in terms of simple services); some means to support, in according to a specific coordination model, the discovery of things (and of their associated services), and the coordinated activities of groups of things; and some solutions to make services and applications capable of self-adapting and self-organizing in a context-aware and unsupervised way.

2.3 Services and Applications

With the term "IoT System" we generally refer to the overall set of IoT devices and to the associated middleware infrastructure devoted to manage their networking and

their context-aware interactions. Logically above an IoT system, specific software can be deployed to orchestrate the activities of the system so as to provide:

- A number of specific *services*. That is, means to enable stakeholders and users to access and exploit individual things and direct/activate their sensing/actuating capabilities, but also coordinated services that access groups of things and coordinate their sensing/actuating capabilities.
- A number of more general-purpose *applications* or *suites*, intended as more comprehensive software systems intended to both regulate the overall functioning of an IoT system (or of some of its parts), so as to ensure specific overall behavior of the system, as well as to provide an harmonized set of services to access the system and (possibly) its configuration.

Clearly, depending on the specific scenario, one can think at IoT systems in which services may exist only confined within the context of some general application, but also at scenarios in which there are services that can be deployed as stand-alone software.

3 Software Engineering Abstractions

Based on the above overview of IoT issues, we now try to synthesize the central concepts and abstractions around which the development of IoT systems (spanning analysis, design, and implementation) should be centered, and discuss how these directly relate to concepts and abstractions developed in the context of agent-based computing [18, 35].

3.1 Actors

The first activity in the analysis of a system-to-be concern identifying the stakeholders and users of the system, aka the "actors". That is, those persons/organizations who will own, manage, and/or use the system and its functionalities, and from which requirements should be elicited.

In the case of IoT systems, the distinction between IoT services and applications, and the presence of an IoT middleware to support them and to manage individual things, naturally leads to the identification of three main abstract classes of "actors":

- *Global Managers*: These are the owners of an overall IoT system and infrastructure, or delegates empowered to exert control and establishing policies over the configuration, structure, and overall functioning of its applications and services.
- *Local Managers*: These are owners/delegates (whether permanently or on a temporary basis) of a limited portion of the IoT system, empowered to enforce local control and policies for that portion of the system.

- *Users*: These are persons or groups that have limited access to the overall configuration of the IoT applications and services, i.e., cannot impose policies on them, but are nevertheless entitled to exploit its services.

The three identified classes of actors are of a very general nature. For example, in a scenario of energy management in a smart city, they could correspond to, respectively: city managers, house/shop owners, private citizens and tourists. In the area of urban mobility, they could correspond to, respectively: mobility managers, parking owners or car sharing companies, private drivers.

3.2 Functionalities

Once the key actors are identified, the analysis preceding design and implementation cannot –for IoT systems and applications– simply reduce to elicit from them the functionalities (i.e., the specific services) that things or group of things has to provide, but has to be taken into account for a more comprehensive approach. In fact:

- Beside things provided with basic sensing/actuating functionalities, one should consider the presence of smarter things that can be activated to perform in autonomy some long-term activities associated with their nature and with their role in the socio/physical environment in which they situates. These can range from simply cleaning robots to more sophisticated autonomous personal assistants [1].
- IoT applications are not simply concerned with providing a suite of coordinated functionalities, but they should also globally regulate the activities of the IoT systems on a continuous basis, according to the policies established by its stakeholders and to their objectives.

As a consequence, other than analyzing the specific functionalities to deliver, one also has to identify the *policies* and *goals* to be associated with services and applications, i.e., the desirable "state of the affairs" to strive for in the context of the socio-cyber-physical system where IoT applications and services operate.

In this perspective, the general classes of functionalities to be identified for the development of IoT applications and services include:

- *Policies* express desirable permanent configurations or states of functioning of an overall IoT system (global policies) or portions of it (local policies), and have the aims of regulating the overall underlying IoT system. Policies are meant to be always active and actively enforced. Although, from the software engineering viewpoint, the focus is mostly on application-level policies, policies can also account for the proper configuration of the underlying hardware and network infrastructures. The definition of global and local policies is generally in charge of the global managers, although local managers can be also entitled to enforce temporary local policies on local portions of the system (provided they do not contrast with the ones imposed by the global managers).

- *Goals* express desirable situations or state of the affairs that, in specific cases, can/should be achieved. The activation of a goal may rely on specific pre-conditions (i.e., the occurrence of specific events or the recognition of some specific configurations in the IoT system) or may also be specifically activated upon user action (e.g., the activation of a goal is invokable "as a service"). The typical post-condition (deactivating the pursuing of a goal) is the achievement of the goal itself. As it was the case for policies, the definition of global and local goals is generally in charge of global, and sometimes of local, managers, whereas users can be sometimes entitled to activate simple local goals (or goals associated to individual things) "as a service".
- *Functions* define the sensing/computing/actuating capabilities of individual things or of group of things, or the specific resources that are to be made available to managers and users in the context of specific IoT applications and services. Functions are typically made accessible in the form of services, and can sometime involve the coordinated access to the functions of a multitude of individual things. Functions and the associated services are typically defined by global and possibly local managers, but are exploited also by the everyday users of the IoT systems.

Clearly, the concepts of goals and policies are central in the research area of agent systems and multiagent systems, and will require, to be realized, components with autonomous and social behaviour, capable of working together towards the achievement of goals and the enforcement of policies.

3.3 Software Components and Their Coordination

Moving from analysis to the design of an actual system and of its components, one should consider that the "things" to be involved in the implementation of the identified functionalities can correspond to a variety of different objects and devices, other than to places and humans, each relying on a pletora of different technologies and capabilities. Accordingly, from both the gluing software infrastructure and the software engineering viewpoints, it is necessary to define higher-level abstractions to practically and conceptually handle the design and development of application and services, and to harmonically exploit all the components of the IoT system.

Most of the proposal for programming models and middleware acknowledge this need, by virtualizing individual things in some sort of software abstraction [16]. The WoT perspective abstracts things and their functionalities in terms of generic resources, to be accessed via RESTful calls, possibly associating external software HTTP "gateways" to individual things if they cannot directly support HTTP interfacing. Other approaches suggest adopting a more standard SOA or object-oriented approach. Surprisingly, only a few proposals consider associating autonomous software agents to individual things [27], despite the fact goals to be pursued in autonomy may be associated to things, a feature that service-oriented approaches can hardly accommodate.

In addition, as already stated, some "things" make no sense as individual entities as far as the provisioning of specific services and applications is concerned, and are to be considered part of a group and be capable of providing their services as a co-ordinated group. This applies both to the cases in which a multitude of equivalent devices must be collectively exploited abstracting from the presence of the individuals [6], and to the cases in which the functionalities of the group complement with each other and needs to be orchestrated [27]. However, due to the dynamic and contextual nature of IoT scenario, traditional service-oriented orchestration methods, although necessary, are not enough to

With these considerations in mind, in an effort of synthesizing from a variety of different proposals and of bringing in as needed agent-oriented concepts, we suggest the unifying abstractions of *avatars* and *coalitions* (see Figure 1).
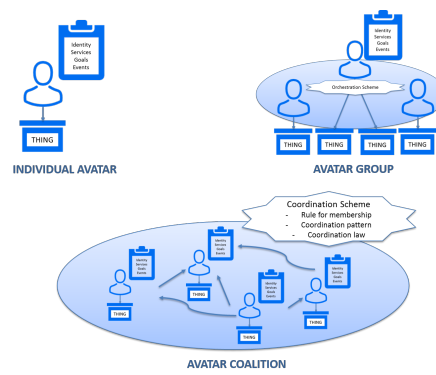

Fig. 1 Avatars, groups, and coalitions.

Avatars. Borrowing the term from [22] (to distinguish from software agents but nevertheless borrowing several features from them) we define an avatar as the general abstraction for individual things and also for group of things (and possibly other avatars) that contribute to define a unique functionality/service. Avatars abstract away form the specific physical/social/technological characteristics of the things their represent, and are defined by means of:

- *Identity*. An avatar has a unique identity and is addressable. An avatar representing a group does not necessarily hides the identities of inner avatars, but it has its own identity.
- *Services*. These represent access point for exploiting the peculiar capabilities of avatars. That is, depending on the kinds of things and functionalities a service abstracts: triggering and directing the sensing/computing/actuating capabilities, or accessing some managed resources.
- *Goals*. Goals, in the sense of desired state of the affairs, can be associated to avatars. A goal may have a pre-condition for autonomous activation, or may be explicitly activated by a user or by another avatar.

- *Events*. Events represent specific state of the affairs that can be detected by an avatar, and that may be of interests to other avatars or to users. Other avatars or users can subscribe to events of interest.

Clearly, for group of avatars, an internal *orchestration scheme* must be defined for coordinating the activities/functionalities of the things (or of the other avatars) it includes. In general terms –and in accordance with assessed service-oriented approaches– an orchestration scheme defines the internal workflow of activities among the composing things and avatars, and the constrains/conditions they are subjected to. Orchestration scheme may also account for contextual information, to make the activities of the group of context-aware. The need of defining orchestrations schemes and constraints to rules the access and usage of (group of) things is generally attributed –with specific characteristics and terminologies– in most middleware and programming approaches for IoT [34, 6].

The avatar abstraction is in line, and account for all the typical characteristics, of most existing IoT approaches. However, the stateful concepts of goals and events make avatars go beyond RESTful approaches. Indeed, these concepts make an avatar more than simply a service provider, turning them into autonomous entities capable of goal-oriented and situated behaviour. Although most existing approaches recognize the need to somehow incorporate similar concepts within RESTful architectures [16], a few of them explicitly refer to agent-based computing, where such concepts belong to.

Coalitions. In this case, and without fear of borrowing the term from the area of multiagent systems [10], we define a coalition as a group of avatars that coordinate each other's activities in order to reach specific goals, or enact specific policies. Accordingly, coalitions may be of a temporary or permanent nature. Unlike avatar groups, coalitions do not necessarily have an identity, and do not necessarily provide services.

To define and bring a coalition in action, the abstraction of coalition must be defined (at least) in terms of a *coordination scheme* that should include:

- *Rules for membership*, to specify the conditions upon which an avatar should/could enter a coalitions. From the viewpoint of individual avatars, the act of entering a coalition can be represented by the activation of a specific goal based on pre-conditions that correspond to the rules for membership [9].
- *Coordination pattern*, to define the pattern (interaction protocol and shared strategy) by which the members of the coalition have to interact. The coordination pattern may include an explicit representation of the goal by which the coalition has been activated. However, such goal can also be implicit in the definition of the protocol and of the strategy.
- *Coordination law*, to express constraints that must be enforced in the way the avatars involved in the coalition should act and interact.

In addition, one can consider the possibility to subscribe to events occurring within the coalition.

The view of avatar coalitions can be of use to realize policies, or to aggregate groups of avatar based on similarity, so as to make them work collectively in a

mission-oriented way without forcing them to specific identity-centered orchestration scheme. This is coherent with the idea of multiagent societies and, in general, of distributed dynamic coordination [18]. Also, this is in line with nature-inspired approaches [36], and approaches to aggregate programming.

The main software engineering concepts discussed in this Section are graphically frames in a logical stack in Figure 2.
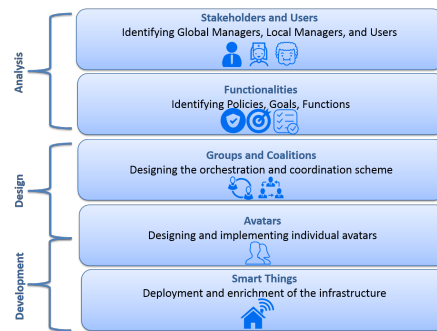

Fig. 2 Key concepts and abstractions for IoT engineering.

4 From Abstraction to Real World: The Teleassistance Scenario

In third generation teleassistance, IoT enriches houses to support smart health monitoring and care. Thus, houses are densely augmented with connected sensors and actuators: light and heat controllers, gas and smoke detectors, presence and motion sensors, door (main doors, internal doors, fridge, kitchen furniture) sensors, electric consume sensors, shutter/curtain controller, as well as sensorized everyday objects (e.g, cup, fork, cane). Suitable "things" may become smart to accomplish this goal:

- RFID tags may be attached to everyday objects in houses such as glasses to detect the quantity of ingested water;
- suitable controllers (e.g., Arduino-based) may be integrated in order to enable turning on/off the light in a specific room remotely (e.g., via a mobile phone);
- autonomous systems that may open and close the shutter/curtain depending on the performed activities, the context (the hour, the day), and/or user's habits;
- robots for home assistance are gaining momentum (e.g., the Giraff plus [28]).

Moreover, also medical devices (e.g., pulse-oximetry, smart scale) may be provided to elderly people in order to automatically send health status information and measures.

In such a scenario, different actors (from medical doctors to elderly people and their family members) can contribute to set up a variety of IoT services to support

both medical doctors in the monitoring and care activities of individuals, and to help individuals and their family members in their everyday self-managed healthcare activities.

Functionalities have to be performed fusing together different things and, possibly, involving different users and stakeholders. For instance, it is desirable to automatically configure a given room (e.g., bedroom) for a given context (e.g., time to go to sleep). This implies a discovery service able to detect the right devices (e.g., the bedroom light actuator) among those belonging to the overall lightening and shutter system. Moreover, coordination is needed to ensure that all the involved devices act in accordance (e.g., if is time to go to sleep, the bedroom light must be turned off if the end-user is on the bed through the coordination with the mattress sensor). Context-awareness is also required to actuate when some pre-conditions are triggered (e.g., the end-user is on the bed and are 10:00 PM) taking also into account that the context may change according to recommendations by caregivers and clinicians (e.g., caregiver may suggest to not to get to sleep a specific day for medical reasons) guaranteeing self-adaptation.

Concerning the general classes of functionalities listed above, in teleassistance we may identify:

- Global and local policies may be defined. For instance, a policy could be introduced to specify the maximum sleeping hours at day or the maximum allowed time for sedentary activities. Policy compliance may be verified relying on the system, for instance, not-intrusive sensors may monitor activities and invite elderly people to be more active or, on the contrary, to go for resting whenever needed.
- Goals to be automatically achieved to give support and assistance, especially in emergency situation. One example could be that of activating an evacuation procedure upon detection of fire by a smoke sensor (pre-conditions), whose goal (and post-condition) is to achieve a quick evacuation of the patient from her/his home. To this end, the activation of a goal can trigger the activities of digital signages and controllable doors in order to rationally guide people towards the exits. Another example could be the case of a fall has been detected. An audio sensor automatically recognizes the help request by the patient (pre-conditions), whose goal is to immediately send assistance at home (e.g., an ambulance) and to communicate with the familiars to make a visit and support the patient (post-condition). To this end, the activation of a goal can trigger the activities of contacting caregivers and familiars.
- Both individual and complex functions should be required: a door sensor in a fridge, e.g., to control opening/closing (individual) to control food in the fridge to verify if the shopping list updating it with needed food (complex).

Finally, let us consider stakeholders and users in the teleassistance scenario:

- Global managers: system managers devoted to control the overall IoT system of the smart houses set according to the directives of the medical doctors, e.g., for deciding heating levels or for surveillance strategies;

- Local managers: house owners, empowered to control the IoT system in their houses and rooms, and tune the local parameters and exploit its services according to own specific needs;
- Users: elderly people with limited abilities, authorized to access specific services (e.g., regulating specific appliances), but not entitled to modify the overall configuration of their houses (in charge of medical doctors and partly of their responsible family members.

Figure 3 shows the different roles of IoT actors in defining and exploiting the above framed functionalities in the teleassistance scenario.
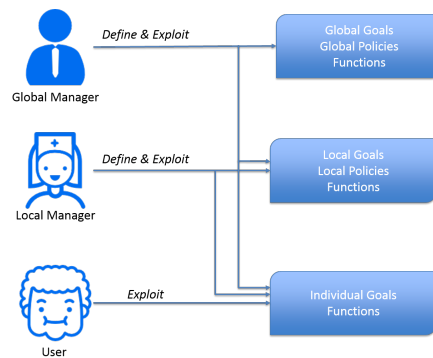

Fig. 3 IoT actors and the functionalities of IoT systems.

5 Conclusions and Future Directions

Third generation teleassistance aims to provide remote support and online help to people that need assistance, as elderly people. Several solutions aimed at monitoring elderly people and their homes have been presented in the literature, most of them based on the IoT technology. In this scenario, this paper proposed and framed some key conceptual abstractions revolving about the IoT universe and showed how they may refer to the teleassistance scenario. The abstractions presented in the paper together with their implementation in this specific scenario represent a first small step towards a general discipline for engineering IoT systems and applications.

As future directions, we are currently defining the overall infrastructure required to put in practice the concepts illustrated and discussed in the paper in order to start testing its real application in teleassistance. Once infrastructure and the basic concepts will be implemented, their effectiveness will be evaluated through a living-lab approach.

Acknowledgments

References

[1] Harshit Agrawal, Sang-won Leigh, and Pattie Maes. L'evolved: Autonomous and ubiquitous utilities as smart agents. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 487–491, New York, NY, USA, 2015. ACM.

[2] Urs Anliker, Jamie A Ward, Paul Lukowicz, Gerhard Tröster, Francois Dolveck, Michel Baer, Fatou Keita, Eran B Schenker, Fabrizio Catarsi, Luca Coluccini, et al. Amon: a wearable multiparameter medical monitoring and alert system. *Information Technology in Biomedicine, IEEE Transactions on*, 8(4):415–427, 2004.

[3] Louis Atallah, Benny Lo, Raza Ali, Rachel King, and Guang-Zhong Yang. Real-time activity classification using ambient and wearable sensors. *Information Technology in Biomedicine, IEEE Transactions on*, 13(6):1031–1039, 2009.

[4] Luigi Atzori, Davide Carboni, and Antonio Iera. Smart things in the social loop: Paradigms, technologies, and potentials. *Ad Hoc Networks*, 18:121–132, 2014.

[5] Debasis Bandyopadhyay and Jaydip Sen. Internet of things: Applications and challenges in technology and standardization. *Wireless Personal Communications*, 58(1):49–69, 2011.

[6] Jacob Beal, Danilo Pianini, and Mirko Viroli. Aggregate programming for the internet of things. *IEEE Computer*, 48(9):22–30, 2015.

[7] NP Bidargaddi, A Sarela, et al. Activity and heart rate-based measures for outpatient cardiac rehabilitation. *Methods of information in medicine*, 47(3):208–216, 2008.

[8] Peter Bower, Martin Cartwright, Shashivadan P Hirani, James Barlow, Jane Hendy, Martin Knapp, Catherine Henderson, Anne Rogers, Caroline Sanders, Martin Bardsley, et al. A comprehensive evaluation of the impact of telemonitoring in patients with long-term conditions and social care needs: protocol for the whole systems demonstrator cluster randomised trial. *BMC health services research*, 11(1):184, 2011.

[9] T. Bures, F. Plasil, M. Kit, P. Tuma, and N. Hoch. Software abstractions for component interaction in the internet of things. *Computer*, 49(12):50–59, Dec 2016.

[10] Yongcan Cao, Wenwu Yu, Wei Ren, and Guanrong Chen. An overview of recent progress in the study of distributed multi-agent coordination. *Industrial Informatics, IEEE Transactions on*, 9(1):427–438, 2013.

[11] Marie Chan, Eric Campo, and Daniel Estève. Assessment of activity of elderly people using a home monitoring system. *International Journal of Rehabilitation Research*, 28(1):69–76, 2005.

[12] Angelika Dohr, R Modre-Opsrian, Mario Drobics, Dieter Hayn, and Günter Schreier. The internet of things for ambient assisted living. In *2010 Seventh International Conference on Information Technology*, pages 804–809. Ieee, 2010.

[13] Céline Franco, Jacques Demongeot, Christophe Villemazet, and Nicolas Vuillerme. Behavioral telemonitoring of the elderly at home: Detection of nycthemeral rhythms drifts from location data. In *Advanced Information Networking and Applications Workshops (WAINA), 2010 IEEE 24th International Conference on*, pages 759–766. IEEE, 2010.

[14] Hulya Gokalp and Malcolm Clarke. Monitoring activities of daily living of the elderly and the potential for its use in telecare and telehealth: a review. *TELEMEDICINE and e-HEALTH*, 19(12):910–923, 2013.

[15] G Sen Gupta, SC Mukhopadhyay, Michael Sutherland, and Serge Demidenko. Wireless sensor network for selective activity monitoring in a home for the elderly. In *Instrumentation and Measurement Technology Conference Proceedings, 2007. IMTC 2007. IEEE*, pages 1–6. IEEE, 2007.

[16] J. Heuer, J. Hund, and O. Pfaff. Toward the web of things: Applying web technologies to the physical world. *Computer*, 48(5):34–42, May 2015.

[17] Marco Iansiti and Karin Lakhani. Digital ubiquity: How connections, sensors, and data, are revolutionizing business. *Harvard Business Review*, 2014.

[18] Nicholas R. Jennings. An agent-based approach for building complex software systems. *Commun. ACM*, 44(4):35–41, 2001.

[19] Nicol Korner-Bitensky and Sharon Wood-Dauphinee. Barthel index information elicited over the telephone: is it reliable? *American journal of physical medicine & rehabilitation*, 74(1):9–18, 1995.

[20] Alexander Kott, Ananthram Swami, and Bruce West. The internet of battle things. *Computer*, 49(12):70–75, 2016.

[21] Kate R Lorig and Halsted R Holman. Self-management education: history, definition, outcomes, and mechanisms. *Annals of behavioral medicine*, 26(1):1–7, 2003.

[22] M. Mrissa, L. Medini, J.-P. Jamont, N. Le Sommer, and J. Laplace. An avatar architecture for the web of things. *Internet Computing, IEEE*, 19(2):30–38, Mar 2015.

[23] Xavier Rafael-Palou, Alexander Steblin, and Eloisa Vargiu. Remotely supporting patients with obstructive sleep apnea at home. In *International Conference on IoT Technologies for HealthCare*, pages 119–124. Springer, 2016.

[24] Xavier Rafael-Palou, Carme Zambrana, Stefan Dauwalder, Enrique de la Vega, Eloisa Vargiu, and Felip Miralles. Third generation teleassistance: Intelligent monitoring makes the difference. In *2nd Workshop on Artificial Intelligence*

and Internet of Things (AI-IoT) @ECAI 2016, The Hague, August 30, 2016. CEUR Workshop Proceedings, Vol. 1724, pages 1–6, 2016.

[25] M. A. Razzaque, M. Milojevic-Jevric, A. Palade, and S. Clarke. Middleware for internet of things: A survey. *IEEE Internet of Things Journal*, 3(1):70–95, Feb 2016.

[26] Jesica Rivero-Espinosa, Andrés Iglesias-Pérez, Jose Antonio Gutiérrez-Dueñas, and Xavier Rafael-Palou. Saapho: an aal architecture to provide accessible and usable active aging services for the elderly. *ACM SIGACCESS Accessibility and Computing*, (107):17–24, 2013.

[27] N. Spanoudakis and P. Moraitis. Engineering ambient intelligence systems using agent technology. *Intelligent Systems, IEEE*, 30(3):60–67, May 2015.

[28] Jonas Ullberg, Amy Loutfi, and Federico Pecora. A customizable approach for monitoring activities of elderly users in their homes. In *International Workshop on Activity Monitoring by Multiple Distributed Sensing*, pages 13–25. Springer, 2014.

[29] GR Vaughn. Tel-communicology: health-care delivery system for persons with communicative disorders. *Asha*, 18(1):13–17, 1976.

[30] Asra Warsi, Philip S Wang, Michael P LaValley, Jerry Avorn, and Daniel H Solomon. Self-management education programs in chronic disease: a systematic review and methodological critique of the literature. *Archives of Internal Medicine*, 164(15):1641–1649, 2004.

[31] G Clare Wenger. Social networks and the prediction of elderly people at risk. *Aging & Mental Health*, 1(4):311–320, 1997.

[32] Robert T Wertz, Nina F Dronkers, Ellen Bernstein-Ellis, Yvonne Shubitowski, Roberta Elman, and Gregory K Shenaut. Appraisal and diagnosis of neurogenic communication disorders in remote settings. *Clinical Aphasiology*, 17:117–123, 1987.

[33] Robert T Wertz, Nina F Dronkers, Ellen Bernstein-ellis, Lisa K Sterling, Yvonne Shubitowski, Roberta Elman, Gregory K Shenaut, Robert T Knight, and Jon L Deal. Potential of telephonic and television technology for appraising and diagnosing neurogenic communication disorders in remote settings. *Aphasiology*, 6(2):195–202, 1992.

[34] Lina Yao, Q.Z. Sheng, and S. Dustdar. Web-based management of the internet of things. *Internet Computing, IEEE*, 19(4):60–67, July 2015.

[35] Franco Zambonelli and Andrea Omicini. Challenges and research directions in agent-oriented software engineering. *Autonomous Agents and Multi-Agent Systems*, 9(3):253–283, 2004.

[36] Franco Zambonelli, Andrea Omicini, Bernhard Anzengruber, Gabriella Castelli, Francesco L. De Angelis, Giovanna Di Marzo Serugendo, Simon Dobson, Jose Luis Fernandez-Marquez, Alois Ferscha, Marco Mamei, Stefano Mariani, Ambra Molesini, Sara Montagna, Jussi Nieminen, Danilo Pianini, Matteo Risoldi, Alberto Rosi, Graeme Stevenson, Mirko Viroli, and Juan Ye. Developing pervasive multi-agent systems with nature-inspired coordination. *Pervasive and Mobile Computing*, 17, Part B:236–252, 2015.