

The Ensafe Home System

Lead Partner: UniPR
Authors: Niccolò Mora, Guido Matrella, Paolo Ciampolini
Contributors: UniPR
Date: 31/05/2018
Revision: V1.1
Dissemination Level: Public



Project Acronym: ENSAFE
Project full title: Elderly-oriented, Network-based Services Aimed at independent liFE
AAL project number: AAL 2014-1-112
With support of:



Content

1	Abstract.....	3
2	System architecture	3
2.1	ZigBee Wireless Sensor Network	3
2.2	Home gateway	5
2.3	Cloud backend.....	5
2.4	Sensor statistics.....	7
	Examples of analytics services	8
2.5	Basic statistics	8
2.6	Trend and outlier detection	9
2.7	Sensor profiles	10
2.8	Further developments	11
3	Conclusions	12

1 Abstract

This document presents the final architecture of the Ensafe Home Kit, which is present in Ensafe levels 3, 4. The system is composed of a Zigbee Wireless Sensor Network (WSN), a Cloud backend, and an analytics part. All pilots share the same architecture, described in Section 2, even though they may vary slightly (in terms of deployed sensors), based on users' needs and constraints. In this sense, the system allows for an installation personalization, tailored to each specific scenario. The cloud infrastructure is flexible and automatically scalable, allowing to deal with multiple pilots. Analytics services, presented in section 0, automatically determine which sensors are present in each pilot and deliver the related results to the users. Some examples of services are presented and discussed.

2 System architecture

The distributed architecture of the ENSAFE can be divided into a local (i.e., home) and a cloud level. At the local level, wireless sensors produce data, which are then sent to a gateway. Relying on the low-power ZigBee protocol. The gateway is in charge of both the management of physical local network and of bidirectional data exchange with the cloud. It actually supports heterogeneous wireless networking standards: in particular, Bluetooth devices can be connected as well, including a weight scale, a blood pressure monitor and a glucose meter, this encompassing the integration of telemedicine features.

Higher-level management tasks are demanded instead to cloud modules, which allows to fine tune and deploy different setups for each pilot scenario. Sensor configuration and diagnostics are carried out at the cloud level, with a rule-based supervisor constantly watching over proper functioning of all devices. Besides supervision, all analytics tasks are performed within the cloud, exploiting unrestricted computing power on demand, and thus freeing the gateway from heavy computational tasks.

In the following subsections, all the components of the system are described in more details.

2.1 ZigBee Wireless Sensor Network

The ENSAFE sensor kit includes the following sensors:

	<p>bed occupancy sensor, to trace sleeping patterns</p>
	<p>chair occupancy sensor, to gather information on how much time and when a user sits on a chair/armchair/sofa</p>

	<p>toilet presence sensor, specifically developed to trace toilet use</p>
	<p>passive InfraRed (PIR) sensors for motion detection, suitable for tracing room occupancy</p>
	<p>magnetic contact sensors, useful for monitoring open/close states of different objects, including, for example, doors, drawers, medical cabinet</p>
	<p>power meter, to monitor home appliances use (e.g., TV, oven, etc.)</p>

Not every sensor is necessarily installed in every home: based on the monitoring needs of each user, a suitable, personalized subset can be defined. Besides sensors, ZigBee networking gear completes the ENSAFE device fleet:

- A coordinator node, consisting of a USB dongle, to be plugged into the home gateway and responsible for the physical ZigBee network management. Through the coordinator, the gateway connects to and gathers information from the sensor network.
- Additional router nodes, for extending (when needed) the ZigBee signal reach all over the home environment.

The above picture includes both commercial, off-the-shelf components, and custom-made ones: in particular, whenever specific functions not available on the market were needed, new devices were designed from scratch, including environmental devices (for instance, bed, chair, fridge, toilet), coordinator and router nodes. To this aim, a platform approach was followed, in which a few general-purpose modules were designed, to be configured and assembled on specific purpose. More specifically, two main modules were designed:

- Radio module. This unit, built around a Texas Instruments CC2531 SoC, takes care of implementing the ZigBee protocol, thanks to the dedicated stack. The CC2531 is also in charge of dealing with basic tasks such as sensor reading and power management. The module exploits an omnidirectional chip antenna to make sensor installation easier.

- Sensor carrier module, that acts as an adaptor and physical interface to the actual sensing element (e.g. bed/chair pressure pad or toilet proximity sensor). Different sensors require different unit assembly, but the base PCB (Printed Circuit Board) is the same

Both radio module and sensor carrier boards have been successfully tested for electromagnetic compatibility issues, qualifying for CE certification.

2.2 Home gateway

The gateway, at the heart of each ENSAFE kit, provides local intelligence and management features. The gateway is implemented using a low-cost embedded (thin client) Ubuntu machine, running the *Ubuntu Server 16.04 LTS OS* for high stability and security features. The gateway is designed to be tolerant to power losses, with all services automatically restarting at each reboot, with no need of human interaction. All programs, ranging from functional network management to cloud connectivity, were coded in *C++* language, exploiting standard libraries (to allow future porting). This allows maximal code efficiency, and make the gateway core suitable for being ported to low-cost, resource-constrained devices (such as *Raspberry PI* platforms).

Each gateway is connected to the USB ZigBee coordinator mentioned above, to enable sensors communication. It is the gateway's responsibility to manage the functional operation of the wireless sensor network, whereas physical operation is demanded to the coordinator and to the ZigBee protocol stack. Once a new sensor requests to join the network, the gateway automatically identifies the device and assigns it a unique name, which will be used as identifier also by the ENSAFE analytics pipeline, and accounts for automatic setting of sensors' parameters settings (via radio commands).

Besides streaming the home sensors data to the cloud, the gateway similarly gathers and forwards data from ENSAFE Bluetooth medical devices. A dedicated service running on the gateway manages pairing, connection and communication with such devices. In particular, data coming from the weight scale, blood pressure meter and glucose meter are decoded and re-formatted to be compatible with the ENSAFE cloud backend, exploiting the same channel used for environmental sensor data.

2.3 Cloud backend

The ENSAFE project heavily relies on cloud technologies. Two high-level tasks can be identified for the cloud backend:

- performing data analytics to get insights from user-generated data
- managing the operating aspects of the pilots.

In particular, the former one can be thought as a *cold path*, meaning that it does not operate in real-time on streaming data, whereas the latter can be called a *hot path*, focusing mostly on the real-time generated data. In the following, we will refer to the hot path as *streaming analytics pipeline*, whereas the cold path will be referred to as the *analytics pipeline*.

The streaming analytics pipeline must meet several requirements, given our application scenario:

- high-bandwidth handling of incoming messages, originating from multiple IoT devices;




Fig. 1: high-level view of the ENSAFE cloud architecture

- ability to perform operations on streaming data, possibly being automatically triggered at the arrival of specific messages;
- Scalability, allowing future massive deployment of ENSAFE pilots.

For the sake of easier deployment and flexibility, such tight requirements are dealt with by resorting to mainstream managed services, such as those offered by *Microsoft Azure*. This also supports automatic scaling of the services, adjusting the computing capability to meet the actual user demand by applying predictive load modelling and simple rule-based actions.

Similar consideration in terms of scalability and flexibility also apply to the analytics pipeline, which exploits powerful statistical and machine learning techniques to mine and extract information from raw data. With this respect too, the Microsoft Azure suite allows for easy deployment of customized machine-learning models: RESTful web services can be automatically deployed, suitable for serving users' requests.

A high-level representation of the ENSAFE cloud architecture is shown in Fig. 1. Functional blocks are described in the following. In particular, elements arranged in the upper row pertains to the hot Stream Analytics path, whereas the lower row represents the cold analytics pipeline.

Data ingestion: High-bandwidth data flow, coming from home gateways, is ingested by exploiting *Azure IoT Hub* technology. It is a fully managed service that securely connects, monitors, and manages multiple IoT devices. Data coming from sensors are stored in the Hub for a fixed (configurable) amount of time, allowing worker nodes/functions to efficiently fetch them.

Stream Analytics: The online, real-time processing capability of the ENSAFE system is achieved via *Azure Stream Analytics*. This is, again, a fully managed service (e.g. avoiding the need of configuring clusters and allowing for easy scaling capabilities), which allows to efficiently query real-time

streaming data gathered by the IoT Hub and perform simple processing on them. The results can then be fed, for example, to a similar message-buffering service or to a permanent storage. Indeed, all incoming sensor data are committed to an *Azure SQL database*, to enable subsequent, off-line processing.

Analytics: Exploiting permanently stored data, mining and modelling is performed relying on the *Azure ML Studio* framework. Within such framework, *Python* functions have been implemented to create customized models, and made available to user’s request through simple web services.

Output unit: output buffering of the Stream Analytics and of the Analytics pipeline is performed via the *Azure Event Hub*, a service similar to IoT Hubs, better suited for machine-to-machine message exchange. In particular, this channel is used for feeding end-users interfaces.

End user interface: a panel for data presentation is delivered to the end-users. It can also be used as a tool for recording users’ inputs, such as subjective wellbeing state (on a scale from 1 to 4), or any contact request.

2.4 Sensor statistics

A typical Ensafe pilot produces the following median number of triggering events per day:

Sensor Type	Daily Median triggering amount
Pir	90
Toilet	36
Bed	26
Plug	16
Chair	18
Door	6

Such detected sensor events do not represent distinct human activities, but rather variations detected by the sensing element. Post-processing and analytics reduce those activations just to meaningful behavioural events.

Overall, summing contribution from all pilots, the following distinct triggering events were detected and logged in the DataBase:

Sensor Type	Total trigger events	Percentage
Pir	129044	34.7%
Toilet	115821	31.2%
Bed	58721	15.8%
Plug	26462	7.1%
Chair	26331	7.1%
Door	15162	4.1%
TOTAL	371541	100.0%

Examples of analytics services

In section 2.3, we discussed the implementation details of the cloud backend, introducing two separate paths for sensor data processing. This section focuses on the so-called cold path, i.e. all the analytics carried out on stored data (i.e., not in real-time), aiming at extracting information and insights on users' habits.

The ENSAFE production environment offers three types of data analysis for each pilot:

- basic statistics about daily activation of sensors (e.g. counts, active time, etc.);
- interpretable Generalized Linear Models for assessing the presence of linear trends, also depending on other (periodic) factors, such as weekend days. This is augmented with outlier detection;
- expected probability/amount of sensor activation, also known as sensor profiles, along with statistical significance analysis for difference between periods.

It is worth remarking that, given the nature of the experimentation, which does not consider active tagging and annotation by end users, all the analyses are run in an unsupervised way.

2.5 Basic statistics

Basic statistics may be derived for visualization purpose. For example, a dashboard may report basic statistics about a sensor activity (or detected events) during the last 24-48 hours.

A similar tool may be exploited to monitor correct functioning of each sensor. An example is reported in Fig. 2, which displays the daily amount of sensor detected events, along with summary histograms.




Fig. 2: visualization of daily detected events counts for each sensor (left), and related summary histograms (right).




Fig. 3: example of regression on the bed usage hours during night.

2.6 Trend and outlier detection

The ENSAFE analytics engine provides automated regression analyses, exploiting Generalized Linear Models (GLM). Such analyses are useful to detect statistical trends, as well as to identify outliers in data. Depending on the quantity being analysed, two regression frameworks are used:

- Real-valued regressions, particularly suited to analyse, for instance, the amount of time a person spends sitting, lying in bed, watching TV and so on.
- Discrete-valued regressions, suitable for analysing the number of toilet visits during night or the whole day, the number of awakenings from bed and many others.

It is worth noting that a given sensor is not bound to a real or discrete-valued framework, but rather it may contain relevant information for both cases. For example, a bed sensor can be used to monitor both the amount of hours spent in bed (continuous value), as well as the number of awakenings during the night (discrete value).

Regression models are applied on a given time-window: for instance, looking at last 30 days allows for spotting meaningful trends while getting rid of seasonal variability of behaviour. Both regressions yield interpretable results, and allow assessing the impact that different features have on the predicted quantity. Common features used in those regression frameworks include, for instance:

- *bias term*: the average event count or measure;
- *linear trend term*: a linearly varying feature trying to detect longitudinal increments (or decreases);
- *weekend day*: a binary feature, which is false during weekdays and true for weekends. This allows to capture differences between those two categories.

The values of such feature help in interpreting possible underlying conditions. For example, an increasing trend in daily bed usage may be interesting to spot. Furthermore, regressions are exploited to detect and identify anomalous, possibly important data.




Fig. 4: sensor profiles for bed and chair sensors. Shaded area represent the uncertainty (confidence intervals) in the probability estimate.

For instance, Fig. 3 illustrates the application of regression model to bed occupancy time: figure refers to a real case, monitored in the framework of ENSAFE pilots. It is shown that an average of about 11.7 hours is spent in bed in the observed period, and a significant increase is shown along time ($\approx +2.9$ h, over the observed period). It is important to note that, given the explanatory purpose of the model, such a trend has not to be interpreted as a prediction for subsequent periods, but rather as a significant attribute of underlying data. The model also detects the presence of outliers (grey crosses in the scatter plot). Such outcomes are relevant per se, allowing (possibly remote) caregivers to discover anomalies which would be hardly noticed otherwise.

2.7 Sensor profiles

From the gathered data, it is also possible to extract the so-called sensor profiles, which attempt to model the expected probability of having a sensor active during a given time interval.

Fig. 4 shows the outcome of such process: activity profiles for bed (solid line) and chair (dotted line) sensors are shown, along with their confidence intervals (shaded areas around each curve). This provides an informational picture of customary living patterns throughout the day. By suitable statistics, shifts or changes in user's habits can be automatically detected, again providing the caregiver with an easy-to-interpret insight, possibly relevant to health or wellness assessment. To this purpose, at each time bin, two populations are compared (e.g. a reference 20-day period against another 20-day one) using the framework of hypothesis testing.

Comparison of profiles, reported in Fig.5, allows to detect a change of behaviour between 10:30 a.m. and 12:30 p.m., which is made evident by non-overlapping confidence intervals. In particular, an increased activity in the second period (dotted line) is noticeable, with respect to the preceding one (solid line). Such increase is consistent with the regression plot in Fig. 3, with the profile analysis providing the caregiver with more detailed and expressive description.




Fig. 5: comparison of bed sensor profiles for two consecutive periods (20 days each). Non-overlapping confidence intervals indicate automatically-detectable, meaningful behavioral changes.

2.8 Further developments

Many more complex analyses can be derived, offline, capable of providing further insights in users patterns. Such methods are not present in the *Ensafe production pipeline*, nonetheless, they yield interesting conclusion.

An example is offered by Sensor profile clustering, which allows to automatically discover the presence of different sensor profiles, for the same person. In fact, one can behave normally in many different modes. For instance, there may be many daily bed usage patterns, in which sometimes the users take a nap after lunch, sometimes they do not. Fig. 6 shows a similar situation.

It is important to underline that such patterns can be automatically extracted from the data, by leveraging clustering techniques. Such patterns may then be used as normal behavioural references, to detect days in which such patterns are not matched.




Fig. 6: extraction of different clusters of patterns regarding daily bed usage. Two prototypes are identified: one having the user rest on bed after lunch (left), and one in which the bed is just used in night hours (right).

3 Conclusions

This document presented the final Ensafe system configuration for level 3-4, featuring a personal home sensor kit. The system is presented in all its components, i.e. Home Wireless sensors, Cloud backend, and Analytics services. Overall, the system is flexible in terms of installation (different sensors may be deployed, tailored to each user's needs), and management (easy scaling of cloud services). Some examples of analytics results were presented, to highlight the capabilities of the system; furthermore, some future improvements are envisioned, to provide more powerful analytics results.