A Self-Organizing Approach to Activity Recognition with Wireless Sensors *

Clemens Holzmann and Michael Haslgrübler

Institute for Pervasive Computing
Johannes Kepler University Linz, Austria

Abstract. In this paper, we describe an approach to activity recognition, which is based on a self-organizing, ad hoc network of body-worn sensors. It makes best use of the available sensors, and autonomously adapts to dynamically varying sensor setups in terms of changing sensor availabilities, characteristics and on-body locations. For a widespread use of activity recognition systems, such an opportunistic approach is better suited than a fixed and application-specific deployment of sensor systems, as it unburdens the user from placing specific sensors at predefined locations on his body. The main contribution of this paper is the presentation of an interaction model for the self-organization of sensor nodes, which enables a cooperative recognition of activities according to the demands of a user's mobile device. We implemented it with an embedded system platform, and conducted an evaluation showing the feasibility and performance of our approach.

1 Introduction

The recognition of user activities is an important aspect in context-aware systems and environments, as it enables their adaptation to the user's current situation and hence allows for providing services with reduced human intervention. The recent availability of *body sensor networks* made it possible to recognize activities with wireless sensors which are mounted on different body parts, like for example embedded in wrist bands, belts and clothes, and are able to communicate to each other as well as to a mobile device [1]. A common approach to activity recognition is the use of accelerometers and the classification of acceleration data into a set of output classes using supervised machine learning techniques [2]. However, usually a precise, application-specific deployment of sensors is used, which does not take into account different numbers, displacements and failures of sensors, and therefore limits the widespread use of such context-aware systems. For example, a person which is running in the woods and equipped with a body sensor network for monitoring his activities, may lose sensors, carry more or less sensors, or their on-body locations may vary due to his movement. In the recently started European research project OPPORTUNITY [3], an alternative

* This work is supported by the FP7 ICT Future Enabling Technologies programme of the European Commission under grant agreement No 225938 (OPPORTUNITY).

approach is proposed, which is based on an *opportunistic recognition of activities* with sensors that are currently available.

A key issue in this regard is the *self-organization* of the wireless sensors [4] into goal-oriented, cooperative sensing ensembles, in order to recognize the activities which are relevant for a user's mobile device – or, in general, for devices in his environment – from a dynamically varying and a priori unknown sensor configuration. This means, that from a set of available sensors, just those which are capable of providing the relevant information cooperate to achieve the goal. In [3], goal-oriented sensing is achieved based on (i) the formulation of a *recognition goal* which represents the activities to be recognized and (ii) its transformation into a coordinated *sensing mission* which is communicated to the sensor network. The scope of the present paper is the *interaction between body-worn sensors*, which is necessary for their ad-hoc formation according to a received sensing mission. First, in Section 2, we describe our general approach for the opportunistic recognition of activities with a body sensor network. In Section 3, the *self-organization* of sensors to achieve a common recognition goal as well as their *self-adaptation* to changing sensor availabilities, characteristics and on-body locations are explained. Finally, Section 4 presents first results of a performance evaluation using a state-of-the-art hardware platform.

2 An Approach for Opportunistic Activity Recognition

Our approach is explained best with an initial application example. Consider a person with a mobile phone in his pocket, and whose clothes are equipped with wireless acceleration sensors; a respective image with our prototype system is shown in Fig. 1(a). These sensors can be used for recognizing his locomotion activity, in order to (i) change the state of the mobile phone accordingly (e.g. to accept or reject phone calls depending on the user's activity) and (ii) notify the caller about the current activity if he is in the user's contact list for example. To become aware of certain activities, the mobile phone *broadcasts* a respective request, which causes the body sensor network to self-organize and provide the needed activity information to the phone. In particular, the phone first formulates a *recognition goal* that basically represents a class of required activities (e.g. locomotion activities such as sitting, standing, walking or running), and automatically translates it into a *sensing mission* describing how the sensors have to cooperate to provide the requested activity. This translation process is conducted by the system and beyond the scope of the present paper.

We define the sensing mission with a *tree data structure* containing those parts of the human body from which sensor data are required to achieve the recognition goal (see Fig. 1(b)). It represents a pre-defined *containment hierarchy* that is given by the human anatomy, with the human as root node and his body parts as child nodes, and by defining which functionality is needed at which node of the tree. As shown in the sensing mission of Fig. 1(c) for example, the sensor mounted on the left thigh has to provide certain features (F), and his parent sensor – which represents the containing left leg – has to do a classification (C)

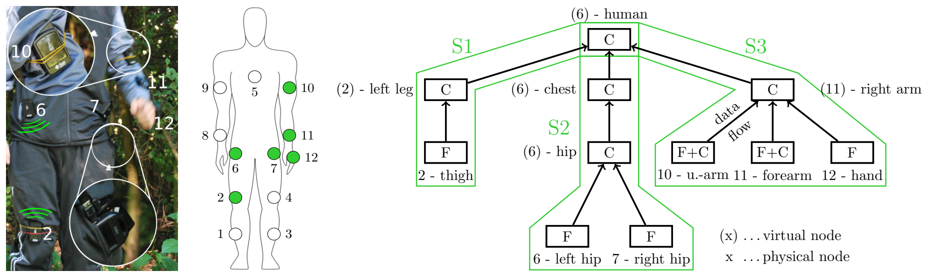

(a) deployment (b) scheme (c) sensing mission

Fig. 1. Exemplary sensor deployment and sensing mission.

of these features. Which feature extractions and classifications are required at which level of the hierarchy, and how the provided results are linked together, is specified in the tree nodes of the sensing mission. A related approach is described in [5], where so-called "processing layers" are used to structure the information flow in a similar way. We assume that every sensor platform has equal capabilities concerning the extraction of features and their classification, and that a sensor (physical node) can undertake the tasks of multiple tree nodes which do not have a physical representation (virtual nodes); more about physical and virtual nodes is explained in the next section. Note that different configurations of sensors may be able to fulfill a given sensing mission; for example, if acceleration information from the left leg is needed, both a wireless sensor located on the left thigh and on the left shank may be applicable.

Furthermore, it is possible to orchestrate simple activities into complex activities by classifying features and the results of lower-level classifiers to high-level activities. This has also the advantage that the amount of information which has to be delivered bottom-up in the tree decreases drastically, which preserves energy due to the reduced wireless communication, and makes such a system feasible for accurate activity recognition (cf. [6]). Since different activity classes (e.g. locomotion and household activities) may require the same low-level activities, which causes overlapping branches in the trees, they can be reused in the fulfillment of *multiple sensing missions* without requiring additional computations. To the best of our knowledge, this *distributed* recognition of activities described above is a novelty of our approach.

For each sensing mission that is broadcast, the receiving body sensors negotiate their participation in the *cooperative* recognition of the respective activities, including especially the extraction of features from the sensor data as well as their classification to activities; a detailed explanation is given in the next section. We have developed a software framework for this approach, which consist of different services. Among others, it comprises a *communication service* for the abstraction of the underlying communication protocol, as well as an *interaction services* for achieving the self-organizing behavior (see Fig. 2).

3 Self-Organization and -Adaptation of Wireless Sensors

Every sensor in the network has to know its own body position, which can be pre-defined or detected as proposed in [7], and is thus able to infer a *local tree* representing its body position from this knowledge. The local tree is based on the containment hierarchy of the sensing mission tree, with the same root and the sensor itself as a *leaf node* (e.g. S1 for sensor 2 in Fig. 1(c)). This makes it possible to compare those trees, which is fundamental for our self-organization approach. When a sensor receives a sensing mission from the mobile device, it participates by providing features that are required for the classification, if and only if its local tree covers a leaf node of the broadcast sensing mission. As every participating sensor is a leaf node in the sensing mission, the leaves are called *physical nodes*; all the ancestor nodes are called *virtual nodes*. The virtual nodes, which contain classifications of features belonging to the corresponding body parts, are also executed by one of their physical child nodes. In the tree S3 of Fig. 1(c) for example, the nodes 10 to 12 extract features from their sensors, and some also perform classification functions. The data of all three nodes are input for the classifier located on the virtual node *right arm*. In this example, sensor 11 additionally undertakes the tasks of this virtual node; this means, that the nodes 10 and 12 send the respective data to the physical node 11, which in turn delivers it – together with its own data – to the virtual node *right arm*.

In order to *self-organize*, each sensor which receives a sensing mission compares it with its local tree, and determines the number of nodes in which they differ. Based on this comparison, we distinguish four different ways to proceed, which are also visualized in the sequence diagrams of Fig. 2:

- *No overlap of the local tree with a leaf node of the sensing mission*: The sensor does not have valuable information for the respective sensing mission and simply ignores the request (e.g. sensor 1 in Fig. 1(c)).
- *Partial overlap of the trees (difference $>$ 1)*: The sensor is part of the solution of the sensing mission, and will therefore try to get elected as a *master* to provide the results to the sensor which has broadcast the sensing mission. For the election, each sensor broadcasts the difference number (e.g. difference 7 for the sensors 6 and 7 in Fig. 1(c), as they are both able to provide the functionalities of the nodes *hip*, *chest* and *human*) together with a random number. A sensor elects itself as a master if it has the smallest difference number, or – in the case of equal differences – the highest random number. Instead of the difference and random number, other metrics such as the battery status of the sensor could also be taken into account, but have not been considered yet. The elected master broadcasts new sensing missions for those parts of the tree which it cannot fulfill (e.g. S1, S2 without *left hip*, and S3 for the master sensor 6 in Fig. 1(c)), which are again compared by receiving sensors with their local trees. These sub-requests may lead to the election of sub-masters, like for example sensor 11 which serves as a master for S3. The master responds with the MAC addresses of the sensors assigned to the respective nodes of the received sensing mission tree.

– *Full overlap of the trees (difference = 0)*: The sensor can satisfy the whole sensing mission by itself (e.g. sensor 2 in Fig. 1(c) would satisfy a sensing mission consisting of S1 only), and thus responds with the tree of the sensing mission in which its own MAC address is assigned to all nodes of the tree.
– *Almost full overlap (difference = 1)*: The sensor needs the information of just *one* other node, and directly requests that information by formulating a sensing mission for the missing partial tree by itself (e.g. sensor 6 in Fig. 1(c) would satisfy a sensing mission consisting of S2 only, and request the missing part from sensor 7). Upon receiving a response to this sub-request, the sensor responds with a tree containing its own MAC address and that of the other sensor. This special case has been introduced due to the fact that a master election process would be inefficient for just two nodes.

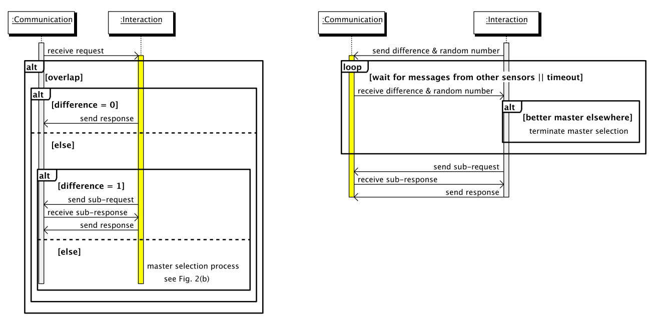

(a) interaction process (b) master election process

Fig. 2. Interaction between sensors depending on their difference to the sensing mission.

If multiple sensors respond to the same superior master or to the mobile phone, just the first answer will be used. After the self-organization is complete, the root master reports to the mobile device that the sensing mission can be satisfied. The mobile device acknowledges this by notifying the master that it can start, which in turn relays this notification to its children to start with their feature extraction and classification tasks. As the sensing mission contains the addresses of all participating sensors, the communication between them is done directly by establishing *reliable point-to-point connections*. With regard to related work on the self-organization of body sensor networks (e.g. [8, 9]), our approach is novel insofar as it (i) hierarchically organizes the sensors according to the anatomy of the human body, whereby the sensors belonging to a certain body part cooperate in the classification of the respective sub-activities and therewith reduce the overall communication load, and (ii) additionally imposes a data-flow that is optimized for changing sensor configurations and activities.

The tree structure is not only used to self-organize the body sensor network for information processing, it is also essential for its *self-adaptation*. If a node fails, the network will try to recover, whereas three scenarios are possible depending on whether (i) a leaf node, (ii) a sub-master or (iii) the root master is concerned. For the cases (i) and (ii), the corresponding parent node will detect a time-out, whereupon it has several handling strategies. First, it will try to notify the failed nodes again and continue with sending the data requested by the sensing mission (by reusing previously obtained data). If still no data is received, e.g. due to a loss or failure of the sensor, the corresponding parent will formulate a new sensing mission for the missing part of the tree, and the network will organize itself again (e.g. by switching to a backup sensor with the same capabilities than the lost one); for example, if sensor 2 in Fig. 1(c) fails, sensor 6 (i.e. the root master) would request the tree S1 again. If this also fails, the root node is informed that the activity recognition cannot continue, whereupon it notifies all child nodes participating in the sensing mission to (i) stop working and (ii) also notify their children accordingly. If the children of a failed node receive a transmission error, they also stop working and notify their sub-nodes accordingly. For the case (iii), the application that requested the sensing mission in the first place can also notify the nodes or send out a sensing mission again, or it can request an *alternative sensing mission* for the recognition goal.

4 Evaluation

We have conducted a first evaluation of our approach by implementing the framework for the Sun Microsystems *Small Programmable Object Technology (SPOT)* platform [10] and the Openmoko *Neo FreeRunner* mobile phone [11] as the user's mobile device (cf. Fig. 1(a)), and measuring the time which is needed for (i) the self-organization of the sensors for different sensing missions consisting of the trees S1, S2 and S3 of Fig. 1(c) as well as compositions of them, and (ii) for the self-adaptation to failures of a leaf node, a sub-master or the root master, according to the strategies re-notification (strategy 1) and sensing mission re-broadcast (strategy 2) at a time; the results are shown in Fig. 3. For the self-adaptation, the complete sensing mission shown in the Fig. 1(c) has been used. It should be noted that there is a timeout of 300ms if the master does not receive the difference and random numbers of the other sensors as well as a retry after 1500ms if it does not receive a response to a broadcast sub-request (cf. Fig. 2(b)), which may lead – due to the unreliability of broadcasts that causes messages to be lost – to high delays for sensing missions with many nodes.

5 Conclusions

The presented work is a first step towards building a distributed activity recognition system with opportunistic sensor configurations. We discussed a novel interaction model for the self-organization of a body sensor network according to a given recognition goal and its self-adaptation to changes in the sensor network, and implemented it as a software framework for a wireless sensor platform. The evaluation showed the feasibility of our approach and the time it requires for the self-organization and -adaptation in different real-world scenarios. As for future work, we will first implement feature extraction methods and classification

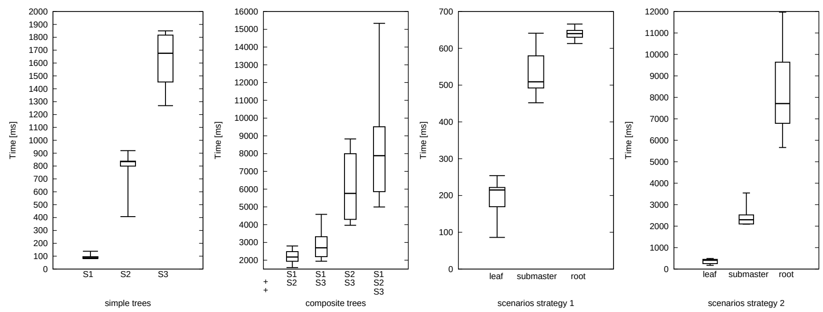

Fig. 3. Time required for the self-organization and -adaptation of the network.

algorithms for recognizing locomotion activities, and evaluate them with respect to their accuracy and the required processing and memory resources; a particular focus will be on the distributed recognition with multiple hierarchically linked classifiers. Another issue of future work is the extension and evaluation of the presented interaction model for its use with multi-hop networks, which would allow for extending it to wireless sensors that are spread in the environment.

References

1. Farella, E., Pieracci, A., Benini, L., Rocchi, L., Acquaviva, A.: Interfacing human and computer with wireless body area sensor networks: the WiMoCA solution. Multimedia Tools and Applications **38**(3) (2008)
2. Ravi, N., Dandekar, N., Mysore, P., Littman, M.L.: Activity recognition from accelerometer data. In: Proc. of IAAI'05
3. Roggen, D., Förster, K., Calatroni, A., Holleczek, T., Fang, Y., Tröster, G., Lukowicz, P., Pirkl, G., Bannach, D., Kunze, K., Ferscha, A., Holzmann, C., Riener, A., Chavarriaga, R., del R. Millán, J.: OPPORTUNITY: Towards opportunistic activity and context recognition systems. In: Proc. of WoWMoM Workshop AOC'09
4. Mills, K.L.: A brief survey of self-organization in wireless sensor networks. Wireless Communications and Mobile Computing **7**(7) (2007)
5. Harms, H., Amft, O., Tröster, G., Roggen, D.: Smash: A distributed sensing and processing garment for the classification of upper body postures. In: Proc. of BodyNets'08
6. Lombriser, C., Bharatula, N.B., Roggen, D., Tröster, G.: On-body activity recognition in a dynamic sensor network. In: Proc. of BodyNets'07
7. Kunze, K.S., Lukowicz, P., Junker, H., Tröster, G.: Where am I: Recognizing on-body positions of wearable sensors. In: Proc. of LoCA 2005
8. Watteyne, T., Augé-Blum, I., Dohler, M., Barthel, D.: Anybody: a self-organization protocol for body area networks. In: Proc. of BodyNets'07
9. Osmani, V., Balasubramaniam, S., Botvich, D.: Self-organising object networks using context zones for distributed activity recognition. In: Proc. of BodyNets'07
10. Sun Microsystems Inc.: Sun SPOT World. http://www.sunspotworld.com (2009)
11. Openmoko Inc.: Neo FreeRunner. http://www.openmoko.com (2009)