

Modeling Service-oriented Context Processing in Dynamic Body Area Networks

Clemens Lombriser, *Student Member, IEEE*, Raluca Marin-Perianu,
Daniel Roggen, *Member, IEEE*, Paul Havinga,
Gerhard Tröster, *Senior Member, IEEE*

Abstract

Context processing in Body Area Networks (BANs) faces unique challenges due to the user and node mobility, the need of real-time adaptation to the dynamic topological and contextual changes, and heterogeneous processing capabilities and energy constraints present on the available devices. This paper proposes a service-oriented framework for the execution of context recognition algorithms. We describe and theoretically analyze the performance of the main framework components, including the sensor network organization, service discovery, service graph construction, service distribution and mapping. The theoretical results are followed by the simulation of the proposed framework as a whole, showing the overall cost of using dynamically distributed applications on the network.

Index Terms

Simulation, analytical modeling and measurements of BAN, Integration of BAN with heterogeneous networks, Platforms for the support of BAN applications

C. Lombriser, D. Roggen, and G. Tröster are with the Wearable Computing Lab, ETH Zurich, Switzerland

R. Marin-Perianu and P. Havinga are with the Pervasive Systems Group, University of Twente, The Netherlands

I. INTRODUCTION

Body Area Networks (BANs) are composed of wireless nodes, ranging from hand-held devices such as mobile-phones, over smart objects in the environment to miniaturized sensor nodes integrated into garments. These devices provide a heterogeneous collection with varying capabilities in terms of sensors, actuators, processing power, memory, and available energy. The number and type of devices forming a BAN change over time, as a result of the interaction with other BANs, e.g. people exchanging objects, or between the BAN and the environment, e.g. clothes or objects taken from chairs.

A user-interface to such a system must hide the complexities involved in handling all the devices involved in the process [1]. One way of achieving this is to make the system context-aware, such as to infer the user state from a set of body-worn sensors [2]. Another type of context is the activity of the user, for which the Context Recognition Network (CRN) toolbox [3] provides a modular composition of algorithms for fast prototyping. Previous work has shown that miniaturized and low-profile sensor nodes with limited processing power are capable to run complex classification tasks, such as the recognition of sound [4].

In this paper, we focus on the organization of activity recognition on dynamic and resource-constrained BANs. We propose a flexible service-oriented framework, which provides the core-components for real-time adaptation to network mobility and heterogeneity. Network mobility is handled through dynamic, context-aware clustering and service discovery. The activity recognition is structured as a service graph, where the various services represent sensors, actuators, and miscellaneous data processing functions. The service graph is dynamically mapped to appropriate nodes in the BAN to execute the recognition algorithm. We model our approach for cluster stability, service graph executability and execution cost.

The paper continues with an analysis of the special properties of BANs and presents the characteristics distinguishing BANs from traditional networks in section II. We present a framework for the distributed execution of applications using the service-oriented approach in section III.

1 To provide a stable processing environment, where nodes stay interconnected, the network is
2 clustered as described in section IV. A model of service distributions is introduced in section V
3 to get an estimation on how many services need to be available on sensor nodes such that
4 applications can be expected to be executable. Section VI introduces a cost model for applications
5 mapped onto the network. This cost is then evaluated using simulations in section VII, showing
6 the behavior and requirements on devices for successful deployment of applications in the BAN
7 environment. The findings in the whole chain from network organization to application execution
8 cost are discussed in VIII, whereupon the paper is concluded.

9 II. CHARACTERISTICS OF CONTEXT-AWARE BAN APPLICATIONS

10 We give two examples of wearable smart assistants, in order to pinpoint the recurring char-
11 acteristics of BANs and the challenges to distributed context processing. The first example is a
12 generic personal sports trainer, while the second example is an assistant for industrial workers.

13 A. *Personal sports trainer*

14 In many sports, the effectiveness of the athlete can be improved by optimizing the course of
15 movement. A BAN system can act as a personal trainer by monitoring the movements using
16 sensors attached to different limbs of the body. It can suggest improvements or training sets to
17 enhance agility or strength. Using actuators such as vibrators or sound, the system can also give
18 feedback, helping the athlete to improve his motions.

19 As an example, we consider running tracks through forests, which are interrupted by stops
20 with equipment, such as bars for pull-ups, or wooden dumbbells. The runners wear clothing with
21 integrated motion sensors, sensing their whole body motion. When the runners reach a stop, they
22 perform one of several activities available at this location, depending on their training targets:
23 strength, agility, or endurance.

24 The BAN is able to automatically detect the sensors attached to the clothing and the ones the
25 athlete interacts with. Depending on the activities to be performed at a certain stop, the BAN

1 downloads the corresponding service graph to monitor the runner's motion. The system must be
2 able to distinguish among multiple runners and the various activities they are performing. For
3 this application, a wide range of body-worn devices and accessories may be worn by athletes,
4 requiring service graphs to be adaptable to many device types and configurations.

5 *B. Worker training and support*

6 A second scenario is the recognition of user activities at work [5]; workers in assembly
7 manufacturing are supported with just-in-time context-aware information on the activities they are
8 doing. Sensors in their clothing and tools measure movements to determine the activity currently
9 performed. The BAN system can thus monitor whether all necessary steps were performed. If a
10 mistake was made, the system may alert the worker in a suitable way.

11 Consider the example of two workers mounting the engine of a car. When the engine is
12 suspended into the engine compartment, the BAN checks if one worker does not move the
13 engine in such a way as to trap the other worker's hand. When they pick up the automatic
14 screwdrivers, the tools are automatically checked for the correct settings, such as rotating speed
15 or torque for the screws. The suspension system holding the engine cannot be released before
16 both workers have completed screwing and stepped back from the car.

17 In this scenario, the BANs are dynamically adapted to incorporate the tools the workers are
18 using. The system needs to distinguish among the nodes attached to different workers by correctly
19 associating the tools with the corresponding BANs. The collaboration among workers becomes
20 also a collaboration among BANs, which may connect to the company database and register
21 completed process steps or issue emergency alarms.

22 *C. Characteristics of a context-processing framework for BANs*

23 From the two scenarios presented above, we derive the specific requirements of BANs, distin-
24 guishing them from typical Wireless Sensor Networks (WSN), and the resulting challenges for
25 executing algorithms in these environments:

- 1 • **Heterogeneous devices** – BANs consist of devices with highly varying sensing, processing,
2 and communication capabilities. This contrasts to the mostly application-specific and homo-
3 geneous networks generally considered for WSN research. As a consequence, a programming
4 abstraction needs to be found which hides the details of the implementation and enables a
5 programmer to write applications executable on a wide range of devices.
- 6 • **Dynamic topology** – As people move in the environment, BANs constantly connect and
7 disconnect to other BANs, static devices or wireless networks. A first challenge lies in
8 identifying clusters of nodes which retain communication for an extended amount of time.
9 For achieving a stable distributed processing, a selection of nodes that can be expected to
10 remain accessible with a high confidence is necessary.
- 11 • **In-network processing** – dynamic BANs require real-time in-network processing, delivering
12 just-in-time feedback to the user. Sensor data is processed where it is sensed, and only the
13 resulting events and alarms are transmitted for further context inference, therefore reducing
14 the network load and at the same time enabling quick response times.
- 15 • **Continuous operation** – The timescale of events that need to be sensed on the body and its
16 immediate surrounding is much shorter than that of the more common WSN deployed e.g.
17 for crop field monitoring or building automation. Therefore, higher sample rates are needed
18 (typically of the order of 10-100 Hz for human activity recognition from motion sensors).
19 As a consequence, only limited time is available for duty cycling or entering sleep modes
20 between data acquisition.
- 21 • **Energy resources** – Handheld devices and sensors integrated into clothes are in most
22 applications removed from the body at night, which gives the opportunity to recharge devices
23 during this time, e.g. using coat hangers in the wardrobe [6]. Energy resources must thus
24 sustain a runtime of a few days rather than months or years.

25 A context processing framework designed for BANs must take into consideration these character-
26 istics. Failure to do so is likely to result in poor or unpredictable performance in this environment.

1 We developed a service-based approach which takes into account the points highlighted above
2 for efficient context processing in BANs.

3 III. SERVICE-ORIENTED CONTEXT PROCESSING FRAMEWORK

4 In previous work, we have implemented the Titan framework for the execution of context
5 recognition algorithms in BANs [7]. Applications within Titan are described using services and
6 their dataflow interconnections. *Services* describe sensors, actuators, or data processing functions
7 offered by members of the BANs. They are considered black boxes offering input and output
8 ports, which can be interconnected to form a *service graph*, representing the application to be run.
9 A service graph description includes parameters passed to the services and attributes constraining
10 the assignment of services to network nodes.

11 Service graphs typically describe the dataflow from sensor readings to data processing result,
12 such as shown in figure 1 for the recognition of motion activity. The service graph implements
13 the algorithms used for the worker support example presented in [5]. The top row of services
14 consist of motion sensors located on the body of a worker and on the tools he uses. The sensors
15 produce input to the algorithm, while the next levels introduce the services selected for feature
16 extraction and first local classification of the activity, which is then fused in the network for an
17 overall result.

18 Each node in the network provides a set of services in a *service pool*, from which they can be
19 instantiated when needed. An unused service does not consume RAM nor CPU cycles, such that
20 the content of the service pool depends on program memory only, which is usually available in
21 larger quantity than RAM. Devices with more resources may provide a larger number of services,
22 and devices with different types of sensors may provide different sets of services in their service
23 pools.

24 The distributed processing is organized as depicted in figure 2. Each network cluster contains a
25 *Network Manager* responsible for the execution of applications, which are stored in their service

1 graph representation in a *Service Graph Database*. On a request for execution, the Network
2 Manager retrieves the corresponding service graph from the database and passes it to the *Network*
3 *Mapper*. The Network Mapper partitions the service graph in subsets and assigns them to network
4 nodes for processing. Therefore it uses the information in the *Service Directory*, which maintains
5 a database of node capabilities and service pools available on the cluster nodes. On the individual
6 nodes, *Node Managers* are responsible of instantiating and executing the services of the service
7 graph subset assigned to them.

8 Upon changes in the cluster, such as new nodes appearing or nodes going lost, the Network
9 Manager is notified by the service directory, and reevaluates the mapping of the application
10 service graph. Different error types are reported to the Network Manager as well and are
11 forwarded to the initiator of the application execution request.

12 We want evaluate the behavior of the Titan framework in BANs. A first important process is
13 the clustering algorithm, which organizes the network in which Titan is run.

14 IV. CLUSTERING ALGORITHM

15 The Network Mapper prefers instantiating services in the local cluster as it assumes those
16 nodes will remain available with higher confidence than others. Thus the cluster should include
17 all and only the nodes carried on the body of the user, which will be referred to as the *correct*
18 state. We measure the *stability* p_S by the percentage of time the cluster is in the correct state.

19 To provide a stable cluster, network nodes are dynamically clustered according to a mutually
20 shared contextual state using the Tandem algorithm [8]. Here, the shared contextual state is
21 whether nodes are located on the same person, which can only be determined with a certain
22 accuracy [9]. Due to this non-perfect accuracy, the perceived shared context may vary in time
23 and lead to cluster instability. In order to analyze the influence of this instability on the service
24 graph mapping, we evaluate the behavior of Tandem using the following mobility model for the
25 BAN.

1 The BANs of individual persons are modeled by groups of sensor nodes moving around
2 an area using a Random Point Group Mobility (RPGM) model [10]. When groups come into
3 communication range of each other, they choose with probability p_m to move together to the next
4 waypoint and to wait there for a random time. After the waiting period, they choose whether or
5 not to continue with each other. This mobility model follows considerations of [11] to produce
6 interconnectivity patterns resembling the social behavior of people.

7 We assume that the moments when the nodes start to cluster and when the clustering structure
8 is freed are triggered by a contextual change or a change in the environment. As an example,
9 we use the moment when a person starts and stops walking or running, which can be reliably
10 detected [12]. When the person is walking or running, Tandem permanently evaluates the shared
11 context among the nodes wirelessly connected and distributively clusters the person's devices.
12 Clusters may change depending on the perceived context (whether nodes are part of the same
13 BAN) or as a result of topological changes (people meet or separate). At each time step, every
14 node chooses itself or a neighboring node with which it shares a common context as *clusterhead*.
15 Tandem tries to achieve stable clusters by keeping the same clusterhead as long as possible and
16 by having each node make decisions based on the status of its neighbors with which it shares a
17 common context. When the person stops walking, Tandem ceases to run and the clusters remain
18 as they are.

19 As soon as a clusterhead is elected, it becomes the Network Manager and receives the service
20 descriptions from each node in its cluster and stores it into its *service directory*. The Network
21 Mapper can then query the service directory for services it needs to instantiate. In case services
22 are not found locally, the queries are forwarded to service directories in adjacent clusters. Using
23 service directories allows a quick assessment of the capabilities of a cluster and speeds up the
24 mapping process.

25 As the next step, we model the distribution of services in the network. This distribution
26 determines what services are available in each cluster and provide the solution space for the

1 Network Mapper.

2

V. SERVICE DISTRIBUTION MODEL

Titan relies on the services available in the service pools of the individual nodes in the cluster. This allows fast reconfiguration and quick adaptation, but raises the question of whether applications can be mapped to the network at all. Therefore we introduce a system model which allows deriving a probability of a service graph being executable in a cluster. Every sensor node has a subset S_a of all known services S available in its service pool. For the analysis we describe the probability of each service $s \in S$ to be available on an arbitrary node with probability $p(s) = \Pr[s \in S_a]$, with $p(s)$ being referred to as the service distribution. The availability, or the probability that a service is available on at least one node in a homogeneous network with n nodes, is:

$$a(s, n) = 1 - (1 - p(s))^n \quad (1)$$

Using equation 1, we can derive a general expression for whether a given service graph is executable in a network of with n nodes. A service graph $A = (T, I)$ is described by a set of services T and their interconnections $I = (t_i, t_j)$. Note that services $s \in S$ may be contained multiple times in T . We thus further introduce the set of *different* service $T_D \in T$ the service graph requires. The product of the service availabilities indicates whether all required services are available:

$$p_{exec}(n, T_D) = \prod_{s \in T_D} a(s, n) = \prod_{s \in T_D} (1 - (1 - p(s))^n) \quad (2)$$

3 The execution probability p_{exec} indicates whether a network with n nodes can provide all the
 4 different services T_D required to execute a service graph. It assumes that services can be
 5 instantiated multiple times at each sensor node and that the probability of a service being available
 6 is the same on every node. The execution probability can be adapted to heterogeneous networks
 7 as will be shown below.

The number of possible mapping solutions of service graphs to the network can be derived by calculating the expected number of nodes providing a service. As each of the nodes independently implements the service with probability $p(s)$, the expected number of nodes is following a binomial distribution having an expected value of $np(s)$. Titan allows to instantiate a service multiple times on every nodes, such that every service can be distributed independently on $np(s)$ devices. The number of possible mappings for a service graph is thus:

$$R(T, n) = \prod_{s \in T} np(s) = n^{|T|} \prod_{s \in T} p(s) \quad (3)$$

1 Services added to the service graph as well as the number of network nodes thus exponentially
2 enlarge the number mapping solutions.

3 In order to further analyze the executability of service graphs on BANs and to show how the
4 model can be adapted to a less abstract case, we analyze the special case where we assume no
5 knowledge about what kinds of application service graphs should be run. We therefore distribute
6 all services with uniform probability. However, we distinguish between powerful devices, such
7 as PDAs or mobile phones, and limited devices, which are integrated into clothing and due to
8 their minimal size can provide only limited processing power. Further we distinguish two classes
9 of services, a set of simple ones $S_e \subset S$, which are available with probability p_{el} on limited
10 devices, and with a probability $p_{ep} \geq p_{el}$ on powerful devices. The second set of complex services
11 $S_i \subset S$, $S_i \cup S_e = S$, can only be implemented on powerful devices with probability p_{ip} .

Using the availability equation 1 and having a fraction r of limited devices in the network, we can derive the expected total number of different services $S_{tot}(n) \leq |S_e| + |S_i|$ available in a network with n nodes:

$$S_{tot}(n) = |S_e|(1 - (1 - rp_{el} - (1 - r)p_{ep})^n) + |S_i|(1 - (1 - (1 - r)p_{ip})^n) \quad (4)$$

During development, S_{tot} may serve as an indication of how large a service directory should be to accommodate all services available in the cluster. S_{tot} also allows to give a general probability for the executability for a service graph requiring $|T_D|$ different services, by considering the

ratio of all possible combinations of service graphs from the available services to the number of service graphs combinations from all existing services:

$$p'_{exec}(n, |T_D|) = \frac{\binom{S_{tot}(n)}{|T_D|}}{\binom{|S_e|+|S_i|}{|T_D|}} \quad (5)$$

p'_{exec} indicates whether the network can be expected to provide all the services required for execution. Figure 3 shows how p'_{exec} changes with the number of nodes and the size of the service graph. The factorials in the binomial coefficient are approximated for the continuous S_{tot} using the gamma function. A low number of different services improves the probability considerably. Adapting the service distribution to service usage statistics of service graphs intended to be run in the network can thus considerably improve p'_{exec} for service graphs following the statistic. It will however decrease the range of different service graphs being executable.

The service distribution model delivers an executability measure for a service graph can serve as a guideline when designing the application service graph. However, during runtime more knowledge of the current state of the network is available. The actual cost in terms of resources required for execution can then be determined by the Network Mapper.

VI. MAPPING SERVICES TO NETWORK NODES

The Network Mapper assigns each service of a service graph to a node in the network and therefore needs to consider the availability of each service on the nodes as well as their resource constraints. Neither the processing nor the communication capabilities should be exceeded on any of the nodes. Additionally, it is favorable to find the implementation keeping resource usage as low as possible.

The task of the Network Mapper is formally described as to map a service graph $A = (T, I)$ onto a network graph $G = (V, E)$. The network graph is described by a set of nodes V and communication links $E = (v_i, v_j)$, $v_i, v_j \in V$. The Network Mapper's goal lies in finding a mapping $M : T \rightarrow V$, such that a given cost function $C(M)$ is minimized.

1 Various cost functions targeting different trade-offs have been proposed for this task, such
 2 as the minimization of transmission cost, total energy consumed, or the maximization network
 3 lifetime [13]. In this paper we use a metric targeting minimization of the total energy used in the
 4 network. The cost function makes use of a model of the sensor node using values stemming from
 5 benchmarking the Titan implementation on sensor nodes [7]. The metric used for the evaluation
 6 relies on three main cost functions:

- 7 • **Processing cost** $C_p(t, v)$ – the cost of processing all services assigned to a node. This cost
 8 results into a measure for whether enough CPU cycles are available to execute all services
 9 of the subset assigned to the given node.
- 10 • **Sensor cost** $C_s(t, v)$ – the cost of using sensors to collect data for the algorithm. As sensors
 11 can usually be turned off when not sampling, this cost value describes the additional energy
 12 dissipated on the node while sampling, and includes possible duty cycling.
- 13 • **Communication cost** $C_c(i, v, e)$ – the cost of communicating data from one service to an-
 14 other for the node v . The communication cost is zero for two services communicating within
 15 the same node. For external communication, it prioritizes intra-cluster communication and
 16 introduces penalties for cross-cluster communication. The cost includes energy dissipated
 17 at the sending and receiving part. It is defined by the message rate.

The mapping is constrained by the maximum processing power $C_{p,max}(v)$ and communication
 rate $C_{c,max}(v)$ a node can support. These limits ensure the executability of the tasks on the nodes
 and guarantee that the maximum transmission capacity is not exceeded. The constraints are given
 for the service graph subset $(T_v, I_{v,e})$ assigned to a node $v \in V$:

$$\sum_{t \in T_v} C_p(t, v) \leq C_{p,max}(v) \quad (6)$$

$$\sum_{i \in I_{v,e}} C_c(i, v, e) \leq C_{c,max}(v) \quad (7)$$

18 Each interconnection i is mapped to an edge e and added to two sets $(i, e) \in I_{v,e}$ as outgoing
 19 and incoming connections. Failure in meeting the constraints results in the service graph not

1 being implementable. In such a case the execution cost will be set to infinity.

The total execution cost of the network is achieved by summing up all costs incurring at nodes participating in the execution:

$$C_{total}(M(A, G)) = \sum_{T_v \in T} \sum_{t \in T_v} C_p(t, v) + C_s(t, v) + \sum_{I_{e,v} \in I} \sum_{i \in I_e} C_c(i, v, e) \quad (8)$$

2 The costs introduced above depend on the device type to which they apply. The parameters for
 3 the device model are sent to the service directory along with the node address. The Network
 4 Mapper further uses a model of the services to derive the service output data rate given a certain
 5 input data rate and the service parameters in the service graph description. When determining
 6 execution cost, the Network Mapper first derives an estimation of the data communicated from
 7 service to service by propagating the data rates generated from each service to each successor.
 8 The individual cost functions make use of the service models and device models to produce the
 9 total mapping cost.

10

11 In this work we are interested in finding the absolute minimum mapping cost of the system.
 12 We do not have an algorithm to find an optimal service graph mapping with minimal cost in the
 13 current implementation of Titan. Also, an exhaustive search is intractable for service graphs and
 14 networks of reasonable size, as the search space grows with $O(n^{|T|})$ (see equation 3). Therefore
 15 we use a Genetic Algorithm (GA) to optimize the mapping, as GAs are known to provide robust
 16 optimization tools for complex search spaces [14]. The GA parameters are selected in order to
 17 favor convergence to the global maximum by selecting a large population size, avoiding premature
 18 convergence, and by performing several runs. The resulting performance is the maximum of the
 19 performance obtained in each run.

20 The service graph is encoded for the GA as chromosome with $|T|$ genes. Every service
 21 is described by a gene containing all nodes in the network providing the service. Mutations
 22 are applied by moving services from one node to another. Crossovers arbitrarily select two

1 chromosomes, randomly pick a gene and swap the gene and all its successors between the
 2 two chromosomes, which are then added to the population. The fitness of the chromosomes is
 3 evaluated using the cost metric given above.

4 Once the implementation of the service graph with the lowest cost has been found, the service
 5 graph subsets can be sent to the individual nodes for execution. We stop our analysis at this
 6 level and do not further model the exact firing rules and scheduling problems on the individual
 7 nodes. In the next section, we simulate the whole chain of algorithms presented to this point and
 8 analyze their interactions.

9 VII. SIMULATION RESULTS

10 The interactions of the clustering algorithm, the service distribution, and the mapping are
 11 evaluated using the mobility model introduced in section IV. Simulations are run for 10'000
 12 virtual seconds with 2 to 10 groups of 10 nodes and the probabilities $p_m \in 0.3, 0.5, 0.7, 0.9$ of
 13 groups joining each other for the next waypoint. The simulation provides an area of 100x100
 14 units and a communication range of 5 units. The group speed averages at 3.62 units per second,
 15 while waiting times result in a mean of 10 virtual seconds. The accuracy of the shared-context
 16 recognition is modeled after experimental results that analyzed whether devices are located on
 17 the same person by correlating sensed motion patterns [9]. Figure 4 shows the stability p_S of
 18 the clusters depending on the number of groups n_S and the joint movement probability p_m . A
 19 higher value of p_m results in longer connections among clusters, introducing more opportunity
 20 for wrong shared context associations and hence lower stability. In a similar way, increasing the
 21 number of groups adds more confusion due to the higher number of possible clusters nodes can
 22 associate with.

23 Services are assigned to nodes in each simulation run following a uniform distribution as
 24 described in section V, with two device and service classes. The model parameters are set to
 25 $p_{el} \in [0.3, 0.4, 0.5]$, $p_{ep} = 0.7$, $p_{ip} = 0.3$, $S_e = 40$, $S_i = 20$, and $r = 0.8$. We repeat each

1 service distribution 100 times for each simulation run. This allows a statistical investigation of
2 the execution cost of a given service graph according to the effective service distribution on
3 individual nodes.

4 The execution costs are determined for a service graph resembling figure 1, with 10 instances
5 of 6 different services, and costs requiring at least 3 nodes. The GA finds an optimal solution
6 for each cluster configuration found in the simulation runs and each service distribution. The
7 GA parameters constitute a population size of 150, a rank selection of the 90 best individuals,
8 6.7% mutation rate, and 50% crossover rate.

9 Figure 5 shows the average execution cost C_{total} , with its variance as shading, and the number
10 of cluster nodes over time. Where the execution cost line is interrupted, no executable mapping
11 could be found for any service distribution, indicating that the sink node receiving the result of
12 the service graph execution is not in the cluster any more. Peaks in the variance result from
13 powerful nodes having left the cluster, such that multiple low-profile nodes need to pick up the
14 processing, leading to higher execution costs.

15 In figure 6 we represent the average execution cost of a service graph as function of p_{el} and
16 r . In this case, we simulate two clusters, one of 5 and one external cluster of 20 nodes. We
17 deliberately selected a cluster, whose size allows to illustrate the influence on the cost of p_{el} and
18 the consequences of recruiting additional nodes external to that cluster. The variance in the cost
19 is shown in figure 6b. With increasing p_{el} , the local cluster is able to support a larger part of
20 the service graph, thereby reducing the cost due to the higher availability of powerful devices.
21 A similar effect can be observed for decreasing r , which increases the number of powerful
22 nodes and reduces the execution cost. The variances show the dependency of the cost from the
23 actual service distribution. A high variance indicates that the cost can be substantially reduced
24 for certain service distributions. The mean difference between maximum and minimum of the
25 mean cost for different service distributions is 32%. This implies that there is a strong benefit
26 for using statistical knowledge about the service graphs that might be issued onto the network.

1 By assuring a higher availability for often used services, the execution cost of these selected
2 service graphs may be substantially reduced.

3 We validate our execution model by comparing the model-derived p'_{exec} to the experimental
4 results observed in our simulation. We illustrate this in figure 7 by comparing p'_{exec} derived from
5 the model and simulation for $r = 0.4$, $p_{ep} = 0.7$ and two values for p_{el} . Results do not match
6 for low sensor number values, which is a result of the service graph not being executable on
7 less than 3 nodes. For higher numbers of nodes ($n > 5$), however, the effect is reduced and the
8 model closely matches simulation results. Therefore we conclude that our model can indeed be
9 used as an indication of whether a service graph is executable on a network with uniform service
10 distribution.

11 VIII. DISCUSSION

12 Within this work we have identified the main characteristics for distributed context processing
13 in BANs and have evaluated a service-oriented approach for its suitability to this environment.
14 Several solutions for service-oriented processing in Wireless Sensor Networks (WSN) have been
15 presented before, such as DFuse [13], OASiS [15], ATaG [16], TinySOA [17], or Viptos [18].
16 Those frameworks have been designed with different objectives and do not completely fit the
17 requirements imposed by the applications we have presented.

18 The Titan framework is specifically designed for *in-network processing* of activity recognition
19 algorithms. The challenge of *heterogeneity* encountered in BANs addressed by a service-oriented
20 approach. Each service is defined by a standard functionality, expected input, and output pro-
21 duced. Hence, it can be implemented in native machine code for the individual devices and
22 can still allow seamless (re-)programming, even in heterogeneous environments. Our framework
23 allows choosing a subset of all services for implementation on participating devices, such that
24 the service pool can be adapted to available processing resources.

25 *Dynamic topologies* are organized by our clustering algorithm to allow for stable execution

1 despite unpredictable dynamical changes in network topology. Therefore we make use of shared-
2 context algorithms, which provide a more robust approach in comparison to simple monitoring
3 of connectivity patterns. The clusterheads maintain service directories for quick access to their
4 clusters' services from foreign clusters.

5 The minimization of resource usage for *continuous operation* is achieved by an appropriate
6 cost model, which is used by the Network Mapper to assign services to nodes for execution.
7 The ability to reprogram nodes allows saving energy at times when sensors and devices are not
8 needed.

9 The general advantages of the service-oriented approach are e.g. the ease of programming. A
10 programmer has just to interconnect services to create applications. No code has to be written
11 and new services may be debugged individually, thus removing potential error sources from the
12 development process. A compact service graph representation allows fast reprogramming of the
13 heterogeneous network while native machine code implementations of services ensure efficient
14 processing. Therefore the approach combines the advantages of virtual machines and code update
15 mechanisms to provide an optimal solution for stream processing in sensor networks [7].

16 The results of this paper may be used by a context recognition algorithm designer to assess the
17 requirements on service distributions and execution costs incurring in a BAN when developing
18 his algorithm. Additional techniques, such as service composition [19], might allow adaptation
19 of service graphs to the actual network configurations encountered at runtime. The system model
20 developed here might serve as an opportunistic selection criterion for candidate service graphs.
21 Therefore it needs to consider only a low number of network parameters, i.e. $p(s)$, S , and n ,
22 instead of requiring complete knowledge about available sensor nodes and services.

23 Titan and consequently also the model do not include the ability to migrate the execution
24 of services from one node to another, such as e.g. DFuse allows it. By migrating the service
25 state, i.e. its main variables, to a node providing the same service, nodes could locally improve
26 the service mapping and thus optimize an initial mapping in a decentralized fashion [13]. This

1 however would require some form of standardized service state import/export format supported
2 by all service implementations.

3 Another point not included in the model is the energy consumption over time. Techniques such
4 as clustering sensors according to required classification accuracy and taking into account the
5 energy consumed by the sensor nodes might prolong network lifetime by turning off unneeded
6 sensor nodes [20].

7 The results of the simulations involving the whole chain from mobility model to service graph
8 execution cost validate the model we have derived. The influence of the model parameters on the
9 execution cost are shown for the special case of uniformly distributed services. This corresponds
10 to the worst-case scenario for which there is no a-priori knowledge about the service graphs
11 intended to be mapped onto the network. The results might lead to the conclusion that with a
12 sufficient count of services on every node, service graphs are always executable and technological
13 advancement in processing power density would give any kind of node the ability to hold enough
14 services. With BANs however, it is more interesting to take advantage of this development
15 by further reducing the size of the devices and by further integrating peripheral components.
16 This would allow more unobtrusive ways of embedding devices into clothing, thus making the
17 technology disappear.

18 IX. CONCLUSION

19 The detection of the context of a user enables the development of proactive and unobtrusive
20 interfaces to complex networks composed of sensor and actuators integrated into clothing, hand-
21 held devices, and smart objects in the environment. We have evaluated the characteristics of
22 a service-oriented context processing framework for such dynamic Body Area Networks with
23 respect to the distinctive challenges encountered in this type of network. Its capability of adapting
24 to node mobility, topology changes, and heterogeneous processing resources provides a stable
25 processing environment for applications in form of service graphs. The theoretical model and the

1 simulation results allow to configure a BAN system in such a way as to optimize the probability
2 of application executability on a given service distribution.

3 The results allow designers of context recognition algorithms to estimate the requirements
4 imposed by their algorithms on network size and configuration early in the design process. The
5 model further provides a basis for heuristical optimization of service distributions for networks
6 as well as service graphs.

7 ACKNOWLEDGMENT

8 This paper describes work undertaken in the SENSEI project (www.sensei-project.eu) sup-
9 ported by the European 7th Framework Programme with the contract number 215923.

10 REFERENCES

- 11 [1] T. Starner, "The challenges of wearable computing: Parts 1 and 2," *IEEE Micro*, vol. 21, no. 4, pp. 44–67, 2001.
- 12 [2] A. Krause, A. Smailagic, and D. P. Siewiorek, "Context-aware mobile computing: Learning context-dependent personal
13 preferences from a wearable sensor array," *IEEE Transactions on Mobile Computing*, vol. 5, no. 2, pp. 113–127, 2006.
- 14 [3] D. Bannach, O. Amft, and P. Lukowicz, "Rapid prototyping of activity recognition applications," *IEEE Pervasive Computing*,
15 vol. 7, no. 2, p. in press, 2008.
- 16 [4] M. Stäger, P. Lukowicz, and G. Tröster, "Power and accuracy trade-offs in sound-based context recognition systems,"
17 *Pervasive and Mobile Computing*, vol. 3, no. 3, 2007.
- 18 [5] O. Amft, C. Lombriser, T. Stiefmeier, and G. Tröster, "Recognition of user activity sequences using distributed event
19 detection," in *Proc. 2nd European Conference on Smart Sensing and Context (EuroSSC)*, 2007, pp. 126–141.
- 20 [6] A. P. Toney, B. H. Thomas, and W. Marais, "Managing smart garments," in *Proc. 10th International Symposium on Wearable
21 Computers (ISWC)*, 2006, pp. 91–94.
- 22 [7] C. Lombriser, D. Roggen, M. Stäger, and G. Tröster, "Titan: A tiny task network for dynamically reconfigurable
23 heterogeneous sensor networks," in *Proc. 15. Fachtagung Kommunikation in Verteilten Systemen (KiVS)*, 2007.
- 24 [8] R. Marin-Perianu, C. Lombriser, P. Havinga, J. Scholten, and G. Tröster, "Tandem: A context-aware method for spontaneous
25 clustering of dynamic wireless sensor nodes," in *Accepted for: Internet of Things, International Conference for Industry
26 and Academia*, 2008.
- 27 [9] J. Lester, B. Hannaford, and G. Borriello, "Are You with Me?" - using accelerometers to determine if two devices are
28 carried by the same person." in *Proc. 2nd International Conference on Pervasive Computing (PERVASIVE)*, 2004, pp.
29 33–50.

- 1 [10] T. Camp, J. Boleng, and V. Davies, "A survey of mobility models for ad hoc network research," *Wireless Communication*
2 *& Mobile Computing*, vol. 2, no. 5, pp. 483–502, 2002.
- 3 [11] M. Musolesi and C. Mascolo, "A community based mobility model for ad hoc network research," in *Proc. 2nd International*
4 *Workshop on Multi-hop Ad Hoc Networks: from Theory to Reality*, 2006, pp. 31–38.
- 5 [12] G. H. Jin, S. B. Lee, and T. S. Lee, "Context awareness of human motion states using accelerometer," *Journal of Medical*
6 *Systems*, 2007, Online First, DOI=10.1007/s10916-007-9111-y.
- 7 [13] U. Ramachandran, R. Kumar, M. Wolentz, B. Cooper, B. Agarwalla, J. Shin, P. Hutto, and A. Paul, "Dynamic data fusion
8 for future sensor networks," *ACM Transactions on Sensor Networks*, vol. 2, no. 3, pp. 404–443, Aug. 2006.
- 9 [14] D. E. Goldberg, *Genetic Algorithms in Search Optimization & Machine Learning*. Reading, MA: Addison-Wesley, 1989.
- 10 [15] M. Kushwaha, I. Amundson, X. Koutsoukos, S. Neema, and J. Sztipanovits, "OASiS: A programming framework for service-
11 oriented sensor networks," in *Proc. 2nd International Conference on Communication Systems Software and Middleware*
12 *(COMSWARE)*, 2007, pp. 1–8.
- 13 [16] A. Bakshi, V. K. Prasanna, J. Reich, and D. Lerner, "The Abstract Task Graph: A methodology for architecture-independent
14 programming of networked sensor systems," in *Proc. Workshop on End-to-End, Sense-and-Respond Systems, Applications,*
15 *and Services (EESR)*, 2005, pp. 19–24.
- 16 [17] A. Rezgui and M. Eltoweissy, "Service-oriented sensor-actuator networks: Promises, challenges, and the road ahead,"
17 *Computer Communications*, vol. 30, pp. 2627–2648, 2007.
- 18 [18] E. Cheong, E. A. Lee, and Y. Zhao, "Viptos: a graphical development and simulation environment for TinyOS-based
19 wireless sensor networks," in *Proc. 3rd International Conference on Embedded Networked Sensor Systems (SenSys)*, 2005,
20 pp. 302–302.
- 21 [19] S. Kalasapur, M. Kumar, and B. A. Shirazi, "Dynamic service composition in pervasive computing," *IEEE Transactions*
22 *on Parallel and Distributed Systems*, vol. 18, pp. 907–918, 2007.
- 23 [20] P. Zappi, C. Lombriser, E. Farella, D. Roggen, L. Benini, and G. Tröster, "Activity recognition from on-body sensors:
24 accuracy-power trade-off by dynamic sensor selection," in *Proc. 5th European Conference on Wireless Sensor Networks*
25 *(EWSN)*, 2008, pp. 17–33.




Fig. 1. Typical service graph for an activity recognition algorithm. The data flows from sensor services at the top through different data processing services to classify the activity observed by the sensors on the body of the user and the tools he uses.

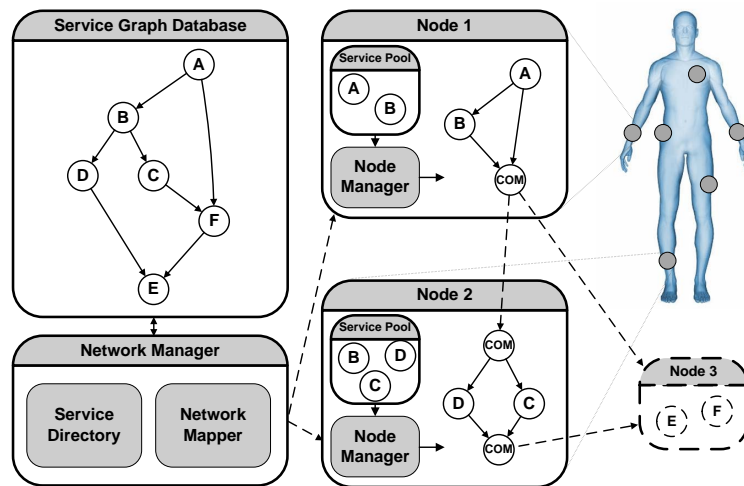


Fig. 2. Framework for the execution of service graphs. Each node can contribute to the distributed processing with the services in its service pool. The Network Manager assigns a part of the service graph to each node in the cluster to run the application.




Fig. 3. Execution probability p'_{exec} of number of nodes accessible n vs. number of service types $|T_D|$ in the service task graph (for $p_{ep} = 0.8$, $p_{ip} = 0.4$, $p_{el} = 0.4$, $S_i = 5$, and $S_e = 40$). The execution probability increases with the number of available nodes n and decreases with the number of required service types $|T_D|$




Fig. 4. Cluster stability p_S depending on the number of sensor groups n_S available in the simulation and the probability p_m of joint movement






Fig. 5. Timeline showing on top the number of nodes available in the cluster and below the execution cost C_{total} with a shading indicating its variance for 6 node groups and $p_m = 0.9$. The clustering algorithm mostly only loses single nodes and has a stability of $p_S = 0.5$.



(a) Mean cost for the execution of the example service graph



(b) Variance of the cost over different distributions

Fig. 6. Execution cost C_{total} for a service graph with $|T_D| = 10$ on a 5 node cluster connected to a 20 node cluster. With increasing r , more low-profile nodes are in the cluster, requiring higher cost. With increasing p_{el} more services are available on all nodes, allowing a cheaper mapping. High variance indicates large cost differences for different service distributions. Application-specific service distributions can improve cost especially in regions with high variance.




Fig. 7. Assessment of the model validity for the execution probability p'_{exec} for a range of nodes for a service graph using $|T_D| = 6$ service types, $r = 0.4$ and $pep = 0.7$. The model does not include execution costs, which result in the simulation values not reaching the model prediction for a low number of nodes, where costs invalidate most of the possible solutions.