

Using Auction based Group Formation for Collaborative Networking in Ubicomp

Christian Decker¹, Emilian Peev¹, Till Riedel¹, Martin Berchtold¹, Michael Beigl², Daniel Roehr², Monty Beuster²

¹Telecooperation Office (TecO), University of Karlsruhe

²Distributed and Ubiquitous Computing (DUS), University of Braunschweig
{cdecker,peev,riedel,berch}@teco.edu,
{beigl,roehr,beuster}@ibr.cs.tu-bs.de

Abstract. In many UbiComp scenarios tiny wireless embedded sensor devices are used, and devices often collaborate to accomplish a common goal. This paper presents a group formation method designed for collaboration of devices. The paper analysis requirements for typical application scenarios, and then presents the general concept of the group formation process. The proposed formation method is a variant of an auction-based algorithm. The algorithm works fully distributed and is based on individual, private beliefs of the devices in the auctioneered item - e.g. a sensor value. The paper introduces the algorithm and analyses its effects analytically and in simulation. We show, that fully distributed operation, high robustness in the case of network failures and extreme low resource (energy) consumption can be obtained. We introduce an application case and present results from a real-world implementation.

1 Introduction

The use of tiny battery powered wireless sensor devices is common to many Ubiquitous Computing (UbiComp) scenarios. Devices as Particles, BTNodes, MITEes or Berkeley motes enable UbiComp application to sense object state and surrounding context. Recently, RFID devices are equipped with sensors thus being able to perform similar functionality at a less complex level. Due to massive price drop of devices, their potential integration into many (everyday) objects becomes feasible. A higher density of such networked sensing devices brings in new opportunities for this novel class of computing devices: One example application is collaborative agreement on redundant sensor measurement which provides huge improvement in reliability of measured sensor values. Another example is collaborative group formation of sensor devices to detect critical situations not detectable by one single sensor system. Furthermore, "group behavior" is common to how human handle real world objects - think of paper files in your shelf! Thus, group behavior and formation may be seen as a necessary requirement for appropriate performance of applications in UbiComp.

The basic concept of collective group behavior, which is the central concept of this paper, is the **collaborative** achievement of a **common goal**. For instance,

several wireless sensor devices, which are attached to objects, work as a collaborative electronic seal [1]. They jointly monitor transport and storage conditions and reason on the common seal's state. Reliability and correct recognition of the seal state is acquired collectively by bringing together information from various types of detection sources from different kind of wireless sensor network devices. This paper presents a novel approach for such robust collective recognition using a specially designed auction algorithm.

The paper starts with a motivating example and requirement analysis followed by an overview over group formation and auction methods. The design of the auction method is presented formally, and its effects are described in theory and based on a performance evaluation. At the end, an application example shows an use case and analysis, how the proposed method helps in that context.

2 Motivation and Requirements Analysis

To informally identify the requirements and constraints for the proposed approach, we depict a motivating example. The example is based on experience collected for a feasibility study from one of our projects [1].

The scenario is taken from an analysis of a notary's office with a large archive containing important paper documents. The archive contains documents that are filed together according to instance or incident. Files are not allowed to be taken apart and documents must not be removed. So far there are only manual countermeasures possible, e.g. supervising each visitor personally while in the archive. The use of technology may be of help, but we learned that there are further restrictions: Not all of the files can be equipped with electronics due to cost reasons, only the important and precious once. Also, it must be possible to take out a archived file to a lawyers desk to allow work on the case, but the system must detect when a single document is removed from the file. No additional manual work should be required to maintain the technical system, such as check in or managing of documents to file order, as this would raise the handling costs. Technical requirements taken from the above example are:

- Autonomous start of the group formation process without user involvement
- Autonomous run of the the detection process
- Infrastructure-less operation: the absence of a document in a file must be detected independent of supporting infrastructure by the involved objects (documents and files). This requires a decentralized, but simple algorithm that could be carried out among the sensor nodes directly.
- For a maintenance free process, minimal energy consumption is required. Minimal energy consumption can be archived by minimizing network time. This impedes the use of existing communication methods and technologies. The exception is the use of superimposing signals as communication technology basis [2] for the algorithm. See discussion at the end of the paper.
- High robustness against miss-detection and performance under inappropriate conditions - e.g. noisy communication channels.

We present in this paper an auction method that is capable of fulfilling the above requirements. The auction algorithm presented is able to collaboratively form a group based on a feature (e.g. a sensor value), or an interval of a feature. In our above example, such characteristic feature (fingerprint) is generated out of movement data. The group is initiated when first bringing in the file into the archive by detecting the same movement fingerprint through the sensors built into each of the documents and the file wrapper. The file wrapper's sensor node is then responsible for holding and comparing the documents belonging to his group and to trigger an alarm when one of the group members leave.

2.1 Potential of Group Formation Algorithms

The above motivation demonstrates only a small part of the potential of the auction based collaborative group formation concept. The proposed method can:

- Set up, supervise and dissolve group membership
- Set up groups according to strict or fuzzy common properties
- Agreement on values (e.g. collective selection of the best sensor reading), similar to voting algorithms. The difference is that auctions algorithms do not require knowledge about the number or identity of participants.
- Epidemic forwarding of agreed values. This is done by successively building up groups that are spread over regional areas.
- Integration of all these processes within a single and minimal communication step, thus greatly minimizing the number of information to be transferred via wireless communication channels. This results in an optimized transmission time and highly reduced power consumptions making the method applicable to low-power sensor networks or passively powered RFID systems.

3 Auction-based Group Formation

Before we present the group formation method, we shortly introduce auction concepts and auction-based group formation.

3.1 Variants of Auction Methods

Auctions are market-based algorithms, where prices are used to achieve a common coordination. Classical auction methods require an auctioneer that is responsible for starting and closing an auction and one or more bidders communicating bids to the auctioneer.

There are several variants of auction methods known, and the most popular is probably the English auction variant. In an open English auction the auctioneer starts with the lowest price value and increments it step-by-step. Each price is broadcasted to all bidders, which compare it to their private maximum value and accept or close by sending an accept/close message back to the auctioneer. Each bidder follows its own bidding strategy to decide on accepting or closing. All

auction participants know all bids. Once the bidder closes the price, she drops out of the next auction step. The auction runs as long as bidders offer bids to the announced price. The final negotiated price is the highest one.

Apart from the above-described English auction, another well-known type is the Dutch auction, where prices are counted downwards. Finally, the Vickrey auction is a closed auction, i.e. only the auctioneer knows the bids. After the start each bidder has exactly one bid and the auctioneer selects the second highest bid. Ebay auctions would be Vickrey type auctions, under the condition that all users use the proxy functionality.

3.2 Auction-based Group Formation Concept

The proposed method is a variant of the English auction type. The difference to the original auction type is twofold: First, there is no dedicated auctioneer. Instead the coordination is distributed among all partners in the auction. Second, there are multiple winners of the auction possible that collectively own the property - i.e. being a member of the group. Sensor nodes take part in the auction based on so-called local belief b . A local belief represents a local and private value of each node stating in how far he believes in the auctioneered item. In the case of a Ubicomp scenario, measured sensor values - e.g. movement fingerprints from the introductory example - are belief items to be auctioneered. A belief in this case expresses the level of certainty of a sensor sampling or processing result of a reasoning process. The benefit of this approach is, that for example results of context recognition processes such as Bayesian inference processes could be directly used as context beliefs. Figure 1 shows an example belief function of a

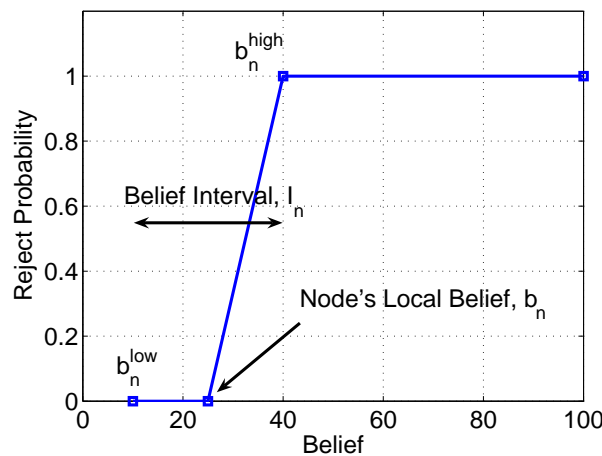


Fig. 1. A node's local auction strategy based on its belief state (example)

sensor node. The belief domain interval is normalized and fixed to $[0; 100]$ in our example. The belief b_n of a node n is surrounded by two border values: b_n^{low} and b_n^{high} . Both values determine the sensor n 's local belief interval I_n . The interval expresses the significance of a local belief. A broad interval means that the belief b_n is less significant than a very tight interval. For example, if b_n is a context belief, e.g. from a Bayesian inference process, then I_n may express the confidence interval around b_n . It suggests how exact the context recognition process is. The more significant, the more exact is the recognition process. As a consequence, the interval around b_n gets tighter. Beliefs outside this interval express then very likely a different context belief than b_n . In relation to the auction algorithm the belief interval defines an auction strategy, i.e. how a node behaves on an offer: A node will at first accept an offer at b_n^{low} and will at latest reject the offer at b_n^{high} . The sensor node may reject even a bit before, as expressed by the growing reject probability between b_n and b_n^{high} in Figure 1.

3.3 Auction-based Group Formation Algorithm

The group formation takes place using the above described believes as items to be auctioned. Each node starts bidding at the belief point b_n^{low} and ends bidding between b_n and b_n^{high} . The algorithm steps are described below and depicted in Figure 2 for a simple example with two nodes.

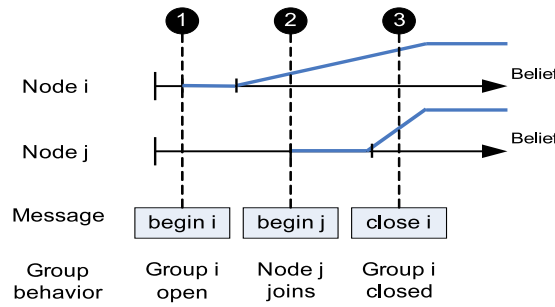


Fig. 2. Auction-based group formation. The auctioneer announces the steps on the belief scale. In step 3, node i closes the group by a close-message. Hence, nodes i and j are in the same group.

1. Node i starts the auctioning process by broadcasting a begin message. The message may contain additional information about the type of item to be auctioned. The node i is only the initiator of the process, but has not further special role.
 - The auctioning process proceeds with time, i.e. the value of the auctioneered item - the belief - increases with time. This means, there is no additional information transfer required in a minimal system. For practical reasons - i.e.

to synchronize clocks between nodes - transmission of a clock synchronization signal might be appropriate. Synchronization of clocks may be generated collaboratively like in [3], or are centralized through an selected node, e.g. the initiator.

2. If any node j detects that the belief value reaches b_j^{low} , it enters the auction by broadcasting a begin-message. This way, all participants of the auction group know the identity of all other members of the group.
 - Any node that is member of the auction group checks, if the auctioneered belief value reaches the values b_n . If this is the case, the node computes its current maximal belief b_n^{cur} using the node's belief function shown in figure 1 as input for the probability to end the auction. The node then waits until (and if) the value reaches the computed b_n^{cur} and sends a close-message.
3. If any of the nodes has send a close-message, the group is closed and all other nodes end processing with the previous step. The group members are all nodes that have sent an begin-message, but have not yet seen a close-message so far. The closing node is also a member of this group. Any node, initiator or any other node, is allowed to close a group.

It is important to see, that the above described algorithm could be implemented in a very minimal way, using e.g. short RF signals - not complete packets. An extensive discussion about the features of the proposed algorithm will be presented in the next section.

3.4 Probability of Membership

The above examples are somehow simplified descriptions of the auctioning mechanism presented in this paper. To be more general, the auction may start at any belief point b^{auct} (which was assumed to be 0 in the previous examples) and may end at any value (was 100 in the previous example). For this we are able to give the probability density function. In the algorithm, the auctioneer announces the belief b^{auct} and any node replies with a reject-message according to the belief function in Figure 1. Hence, the probability for a reject message is defined as follows:

$$p(b^{auct}) = \begin{cases} 0 & \text{if } b^{low} < b^{auct} < b \\ \frac{b^{auct}-b}{b^{high}-b} & \text{if } b \leq b^{auct} \leq b^{high} \\ 1 & \text{if } b^{auct} > b^{high} \end{cases} \quad (1)$$

We denote by $X_i = j$ the membership of node j to group i .

$$\begin{aligned} \forall j \ b_i^{low} \leq b_j^{low} \leq b_i & \Rightarrow P(X_i = j) = 1 \\ \forall j \ b_i < b_j^{low} \leq b_i^{high} & \Rightarrow P(X_i = j) = 1 - p_i(b_j^{low}) \\ \forall j \ b_i^{high} < b_j^{low} & \Rightarrow P(X_i = j) = 0 \end{aligned} \quad (2)$$

This approach can be further extended if the group contains more than just one node. Lets take for example a group containing no less than l members. Node k has sent its begin-message and therewith has opened the group k . The probability

that a new node j with a announced believe b_j^{auct} will join this group k can be calculated using:

$$P(X_k = j, l) = \prod_{i=k}^{k+l} [1 - p_i(b_j^{auct})] \quad (3)$$

This equation applies when group k has l or more nodes. The probability in (3) states: if there are more nodes within the belief interval I_k of node k , it is less likely for j to join in. The more nodes are within I_k , the more belief functions are ready to trigger a close-message before j has sent its begin-message.

Equation (3) delivers an easy way to calculate the expected value of members in a given group. The approach relies on two conditions. First, the slope in the belief interval, denoted by s , must be the same for all nodes and second all beliefs, which the auctioneer can announce in this interval, must be occupied by the nodes. We begin with the cumulative distribution function (cdf) for a belief slope s given by $\Phi_s(l) = 1 - P(X_k = j, l)$. Because of the monotonic decrease of $P(X_k = j, l)$, we have to apply $1 - P(X_k = j, l)$ in order to get the cdf. The probability density function (pdf) is obtained by the difference quotient Δ , which is the slope in $P(\cdot)$ between l and $l + 1$. We denote $\Delta\Phi_s(l) = \phi_s(l)$. We then proceed to calculating the expected value using the standard formula:

$$E_s [X_k] = \sum_{b_k}^{b_k^{high}} l\phi_s(l) \quad (4)$$

$E_s [X_k]$ denotes the expected number of group member when all nodes have the same slope s and all believes in the belief interval of a group-starting node k are taken by other nodes. The results from our extensive simulations further confirm the validity of the result from equation (4).

3.5 Discussion

This section will discuss primary features of the auction-based group formation algorithm and identify advantages of the approach.

Distribution of control, autonomous and infrastructure-free operatoin.

The proposed algorithm does not require any dedicated node to drive the auction process, nor any infrastructure support, nor any manual intervention. Although there is an initiator, this role is only defined by this object being the first one to send an open message, and is not distinguished among other nodes. In practical implementation, that initiator may additionally set the initial values and may also define parameters for embedding the algorithm in a concrete application.

Minimized number of communication messages. The number of messages is minimal for the overall process in general. This is because there is only one begin message required from each of the participants - which is minimal because else other group members do not know about the new member. It is also minimal, because there is only one close message required from one of the node to

end the process. Acceptance of group membership for all other nodes is implicit.

Minimized Message complexity. There are only two type of messages required for a minimal process run. Assuming that the communication characteristics are know beforehand, and that the belief interval starts at 0, we would require two simple signals for performing the process. The benefit of such minimal signaling is twofold: For all types of devices, minimizing signaling greatly reduces energy consumption. Compared to traditional voting algorithms and the required exchange of packets, several magnitudes of gain in energy performance can be expected. The second benefit is the simplicity of the algorithm. Because of this characteristic, the use of the algorithm in extreme simple circuitry - e.g. in ultra-low-cost RFID systems - is possible.

Data compatibility. We see believes as a kind of universal model to express a value and the trust in the value. These values are independent of the used encoding scheme, and thus appropriate for Ubicomp settings with very heterogeneous devices. The analog scheme of the communication process is even capable of including devices that communicate using analog signaling and processing.

Formation control. As presented so far, the group formation algorithm does not include leaving a formed group. Leaving a group can be implemented by a successive group formation runs, where all those nodes that are not willing to believe anymore step out by not bidding. This method has the advantage that a permanent supervision of the group membership takes place and that dynamic changes - e.g. nodes that leave the area or run out of battery or communication link breaks - are automatically taken into account. Dynamicity of the process can be controlled by either setting a fixed time interval for repeated group formation, or by letting any node that wants to propagate information starting the formation process again on demand. The benefit for the latter strategy is the highly reduced power consumption, while the benefit of the first strategy is a highly increased robustness of the process.

Multi-Group Formation. The group formation process can be used to determine group memberships according to various believes by simply successively starting an auction process with different believes. This allows to simply build up interest groups within a large set of sensor nodes.

Privacy. Although nodes are required to uncover their entry point b^{low} , to the members of the group, neither b^{high} nor b are unveiled. For the overall process, other strategies than depicted in figure 1 are thinkable. Different strategies of individual nodes do not disturb the general Auction-based Group Formation Algorithm from section 2. This way a high degree of privacy could be ensured for all participating nodes. As we expect that Ubicomp systems will be used in business and private contexts, this is a must-have property for many applications.

Robustness and Message loss. Loss of messages due to high noise level in the wireless communication channel or to changes in the environment are a common phenomenon in Ubicomp settings. Message loss will bias the auction results for the achieved group size as well as the number of formed groups. But neither a lost

begin nor close-message necessarily excludes nodes from a group. When missing a begin-message, other group members joining in will send a begin-message according to the above algorithm thus proving an additional point-of-entry for a node missing a previous begin-message. When missing an close-message, those nodes that are not part of the group yet, may enter the group and thus increase group size - instead of starting their own group. Therefore, close-loss causes less groups, but with more members due to the 'virtually' increased belief interval. In a practical setting there is a simple countermeasure to this problem: Because communication requires only short signaling, repeated run of the algorithm may be feasible even when real-time, low-power communication is required.

4 Performance Evaluation

Group formation depends on a series of parameters such as the belief interval and the number of nodes. In this section we investigate the effects on the parameter choice through simulations. We also investigate the conditions under which the performance of the algorithm declines, but will also show that this happens only under VERY extreme conditions. It should be noted that the group formation is itself a stochastic process, therefore the most results and figures contain average and/or expected values.

4.1 Node Count

The section explores the relation between node count and group formation. We make no assumption on some distinguished beliefs and beliefs are uniformly distributed among the belief scale, here $[0; 100]$. The belief interval for all nodes is equally set to $I = 20$. Figure 3 shows the histogram of the number of groups after 1000 simulations for two cases, 21 and 61 nodes respectively. The group count raises with the introduction of more nodes, which can also be observed for an increasing range of nodes in Figure 4. In the range between 20 and 50 nodes





Fig. 3. Distributions of groups with 21 nodes (dark) and 61 nodes (bright)

Fig. 4. Number of formed groups for a given number of nodes

this relation continues almost linearly. However, from point 50 it saturates due to the overpopulation of the belief range with nodes that have overlapping belief intervals. To conclude, for increasing node counts the algorithm may experience saturation, which stops the formation of new groups.

4.2 Robustness and Bias

This section investigates the robustness of the auction-based algorithm in the case of message loss. The message loss probability ranges span across the whole range between 0 (no message loss) and 1 (all messages are lost). Message loss can occur with begin-messages as well as with close-messages. The loss of messages can have a profound impact on the group formation process and especially on the count of the formed groups. The consequences of a lost close message can lead to a stretching of certain groups and delayed or no formation of groups that follow. Figure 5 shows the change in the formation of new groups in the cases of no

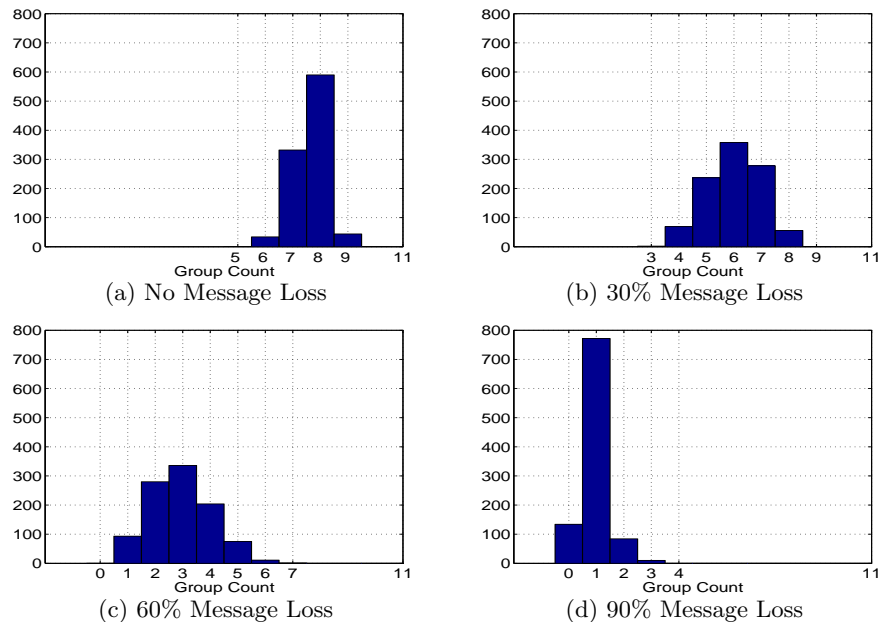


Fig. 5. Group count distribution for a setting with 21 nodes (1000 simulations)

message loss a), 30% b), 60% c) and 90% d) message loss. There exists a tendency of shifting the average group count to the left. As the message loss probability increases the formation of new groups decreases until it reaches a critical point around 90%. From this point forward the auction-based algorithm forms either one large group or no group at all. Figure 6 a) illustrates the distortion in the

group count, which is caused by message loss. The relation is approximately linear. An interesting feature is that the form of the function does not change with the increase or decrease of the node count. What changes is the steepness of the almost linear relation. The results found in figure 6 b) show impact of loss of certain messages on average group size. The auction may experience from two types of message loss: loss of begin-messages and loss of close-messages. A typical loss of a begin message must lead to a decrease in current group size. A loss of close-message may lead to an increased group size. To evaluate the impact, we count every begin-loss as a negative distortion on the group size and every close-loss is positively counted. Figure 6b) shows a negative curve indicating that begin-losses causing a dominant distortion on the average group size rather than close-losses. An explanation for such a behavior can be found in the algorithm itself. Every close-message loss can be 'replaced' by a successful close-message from the node that follows, but unsuccessful begin-messages always lead to a decrease in the group size no matter how many members there are in the current group. Another interesting feature of the auction-based algorithm is also shown

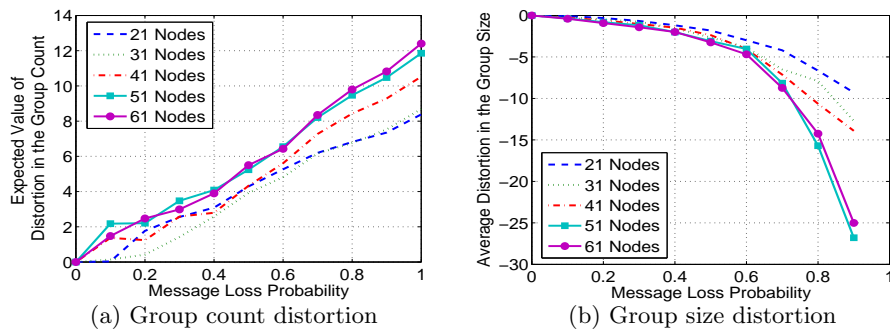


Fig. 6. Bias for group count and size caused by message loss

in figure 6b). The average distortion is almost constant up to the point of about 60 Percent for all node counts.

4.3 Accuracy

The accuracy describes how well the auction-based algorithm identifies a single group. We assume that group members follow a normal distribution $\mathcal{N}(b_\mu, \sigma)$ around a given belief b_μ . The parameter σ indicates how strongly a group is stretched. Each group member is assigned to a belief b_n selected from the normal distribution around b_μ . The accuracy depends further on the belief interval I_n of each member. It is further $I_1 = \dots = I_n = I = \text{const.}$ for all members. In this investigation we are interested in the expected number of groups $E_{\sigma, I}[G]$ for a distribution of group counts G found by the auction algorithm. The figure 7 depicts this expected value. It can be seen that the auction-based algorithm

works very stable for variances below $\sigma = 1.4$. In these cases, it always forms one single group depicted as a black surface in the figure 7. This is the ideal result. Brighter surfaces in the figure mean group counts larger than 1. The legend at the right side of the graph reflects the expected group counts. We must further note that the parameter σ has a much larger impact on the accuracy of the algorithm than the belief interval I .





Fig. 7. Accuracy of the auctioning for finding a single group for a belief interval I and a group variance σ (black = single group) **Fig. 8.** Selectivity: expected number of groups, when two groups move towards each other (1000 simulations, no loss)

4.4 Selectivity

Selectivity describes the accuracy of separating two groups. We select here two groups, which can be clearly identified as separate groups according to the result from sec. 4.3. Fig. 8 shows the expected value for the number of groups that can be formed using our algorithm. The two groups contain 10 member nodes with beliefs according to a normal distribution with $\sigma = 1.1$ and the interval length of the belief probability function $I = 15$. According to fig. 7, the influence on the accuracy of the algorithm is minimized. The belief distance describes the distance between the centered beliefs values of the two groups on the belief scale. The curve in fig. 8 shows the growing tendency in the selectivity of the algorithm. A critical point is reached at belief distance 10. From that point forward the auction-based method separates the two groups precisely.

5 A Use Case for Collaboration and Group formation

Within the research project Collaborative Business Items, CoBIs ([4], <http://www.cobis-online.de>), we explored automatic workplace safety enforcements when handling hazardous chemicals in one of BP's plants. Wireless sensor nodes were attached on chemical drums (figure 9) and collaboratively processed critical information in order to detect hazardous situations like an exceeded storage limit, prohibited storage combinations of materials or an invalid storage area. Alerts

were raised visually on the drums for notification of nearby workers. The setting required in-situ, infrastructure-less, real-time detection. CoBIs was implemented on the Particle Computer sensor node platform [5]. The configuration consists of communication board and sensor board in splash water resistant housing: 8bit PIC18LF6720 MCU, RFM TR1001 communication at 868MHz with 125kbit/s, sensors: light, temperature and acceleration, Actuators: two ultra bright LEDs plus warning light (dedicted nodes only). In [6], we explored an agent-based

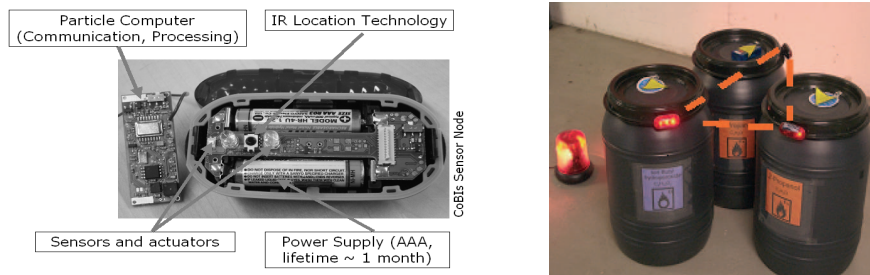


Fig. 9. Left: Particle nodes for CoBIs. Right: Detecting hazardous situations

implementation of auctions on Particle nodes. The auction mechanisms were implemented on the Particle sensor node using the Particle JavaVM [7]. The auction logic runs independently as a Java agent class on the node. Which node starts the group formation depends on who first detects a situation, e.g. on the detection of a hazardous event. In the implementation the AwareCon TDMA synchronization protocol [3] provides a ultra-low power distributed clock synchronization algorithm for exact high-reliable group formation.

5.1 Collaborative Alerting with Auction-based Group Formation

In CoBIs, nodes constantly test the current storage conditions such as amount of stored material and dangerous combinations, e.g. inflammable and oxidizing chemicals. Once a violation against pre-defined conditions is found, the nodes raise locally an alert in order to inform nearby workers. Alerts are propagated to other nodes, which also raise an alert and further propagate it. As a result, the alert is flooded through the wireless sensor network. Alerts emitted by the source of the condition violation are so-called *local alerts*, while re-emitted alerts are *remote alerts* (figure 10). Ending an alert situation is more difficult: How can the alert decay or in other words how can the nodes achieve the common state of non-alert? Most nodes are always slightly in a alert situation, also because of their high sensitivity to potential dangerous situations. Hence, remote alerts are sent repeatedly and raise other remote alerts. This motivates a belief model of the alert states as show in figure 11. The alert states are associated with the zones for non-alert, remote alert and local alert on the belief scale. The auction-based group formation detects number of groups and their size. If the alert, e.g.




Fig. 10. Propagation of a hazard alert




Fig. 11. Belief scale divided in alert zones (example); lower fig.: beliefs of nodes moving towards the non-alert zone

a storage violation is resolved, the local alert group dissolves. This is detected by other groups and the remote alert groups decay their beliefs, e.g. by halving them. Each node does this individually. As a consequence, groups move towards the non-alert zone (see figure 11, lower fig.) similar to the findings in section 4.4, while moving groups are stretched and separated in smaller groups.

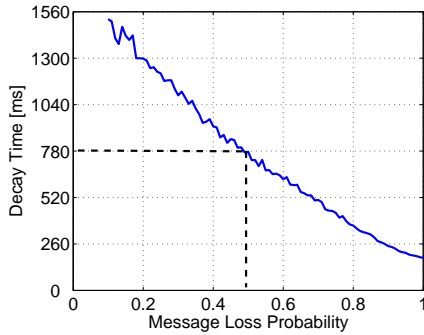


Fig. 12. Alert decay behavior for low pass filter. Message loss behaves causing a low pass effect.

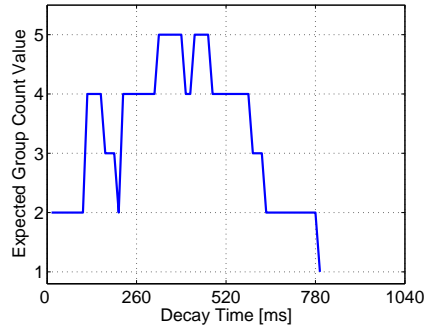


Fig. 13. Group formation behavior when converging to the non-alert zone. After 780ms all nodes achieve the non-alert state.

We compare the group formation approach with a low pass filter implementation - an application specific implementation for achieving the common non-alert state. The low pass runs on each node and decays the remote alert after a timeout. However, if it receives an alert after it exceeded the timeout, it will start over alerting. Because timeouts are slightly different, an alert will ping-pong forever when no message get lost. The degree of message loss determines overall decay time of an alert for the low pass filter (figures 12 , 13). The alert timeout was set to 130ms and blurred for individual nodes to account for slight timeout drifts. For the low pass filter in fig. 12 the message loss corresponds to the decay timeout. For the auction-based group formation the fig. 13 depicts the expected

number groups. They finally converge to exactly one group into the non-alert zone at 780ms. We found that as long as the message loss is below approximately 50%, the auction-based approach will outperform the low pass filter. Further, the ping-pong effect is resolved and convergence to the non-alert zone is guaranteed. Decay times are derived based on AwareCon's 13ms TDMA slot. Ideally, if no loss and no time drift occur, the alert will decay in one timeout period of 130ms.

6 Related Work

The XCast approach establishes a distributed shared information space for wireless sensor nodes. Nodes share common information through XCast. It was shown by the authors, that the design can reach a fix point, i.e. all nodes have a consistent view and access on a quantum of shared data. In the area of distributed systems, consensus describes a mechanism to agree on quantum of information between members of a resource population [8]. A prominent example is the Byzantine Agreement (BA), which explores consensus in a distributed voting process [9]. The goal is to agree on a given state despite faulty systems. In contrast to XCast and BA, the auctioning approach forms the group across several belief states. It *agrees on a new value*, represented by the group, which is a compromise for all members. As a consequence, private information has not to be revealed to others. Further, it relaxes the very strict assumptions on message loss and synchronization in case of BA. In the context of CoBIs, a related approach for enabling collaboration is the arteFACT framework [10]. It implements a Prolog interpreter, which uses rules in form of Horn clauses and proves for inconsistencies. Once it discovers one, an appropriate action may be raised. It presumes that all input data is available at the time of rule evaluation. In contrast, auction-based group formation may even proceed to operate when data is missing due to message loss. The investigation in section 4.2 lead even to an estimation of the resulting bias and allows instant corrective action. Finally, group management protocols for sensor networks were developed [11] to organizes them in groups. The group has one representative node and others may join in. The protocols coordinate joining and leaving of a group. The selection of which group is appropriate to join is completely left to the application. Mostly, physical proximity determines the membership. Auction-based group formation builds distributive groups according across a variety of a node's local information. It is automatically ensured that beliefs are compatible.

7 Conclusion and Outlook

We presented an auction-based negotiation mechanism for group formation of networked embedded systems for ubiquitous computing. It exhibits crucial properties for collaborative Ubicomp applications such as data compatibility, robustness against message loss, and privacy. The algorithm runs truly distributed. Auction-based group formation achieves the goal of providing a uniform abstraction layer. Groups are formed across various networked embedded systems

and their various local information. The application in the CoBIs use case and the achieved results show that it adds significantly to the performance of collaborative networked embedded systems and outperforms application specific implementations for achieving a common agreement. Future research will analyze effects of node density and large distributed networks. The research question is whether the formation process can be parallelized and local groups can be merged in larger ones. Cooperative transmission was successfully shown to work on sensor nodes [2] and might be an efficient signaling mechanism for the auction. New Ubicomp applications will be based on this group formation and the effects on their performance will be investigated in more detail.

Acknowledgments

The work presented in this paper was partially funded by the EC through the project RELATE (contract no. 4270) and by the Ministry of Economic Affairs of the Netherlands through the project Smart Surroundings (contract no. 03060).

References

1. Decker, C., Beigl, M., Krohn, A., Robinson, P., Kubach, U.: eesal - a system for enhanced electronic assertion of authenticity and integrity. In: Pervasive. (2004)
2. Krohn, A.: Optimal non-coherent m-ary energy shift keying for cooperative transmission in sensor networks. 31st IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP) (2006)
3. Krohn, A., Beigl, M., Decker, C., Riedel, T.: Syncob: Collaborative time synchronization in wireless sensor networks. In: 4th International Conference on Networked Sensing Systems (INSS). (2007)
4. Decker, C., Riedel, T., Beigl, M., de Souza, L.M., Spiess, P., Mueller, J., Haller, S.: Collaborative business items. In: 3rd IET International Conference on Intelligent Environments. (2007)
5. Decker, C., Krohn, A., Beigl, M., Zimmer, T.: The particle computer system. In: ACM/IEEE Information Processing in Sensor Networks (IPSN). (2005) 443 – 448
6. Decker, C., van Dinther, C., Miller, J., Schleyer, M., Peev, E.: Collaborative smart items. Workshop for Context-aware and Ubiquitous Applications in Logistics (UbiLog) (2007)
7. Riedel, T., Arnold, A., Decker, C.: An oo approach to sensor programming. Poster Abstract, European conference on Wireless Sensor Networks (EWSN) (2007)
8. Barborak, M., Dahbura, A., Malek, M.: The consensus problem in fault-tolerant computing. *ACM Comput. Surv.* **25**(2) (1993) 171–220
9. Lamport, L., Shostak, R., Pease, M.: The byzantine generals problem. *ACM Trans. Program. Lang. Syst.* **4**(3) (1982) 382–401
10. Strohbach, M., Gellersen, H., Kortuem, G., Kray, C.: Cooperative artefacts: Assessing real world situations with embedded technology. In: 6th International Conference on Ubiquitous Computing (UbiComp). (2004)
11. Vieira, M.S., Rosa, N.S.: A reconfigurable group management middleware service for wireless sensor networks. In: MPAC '05: Proceedings of the 3rd international workshop on Middleware for pervasive and ad-hoc computing, New York, NY, USA, ACM Press (2005) 1–8