Real-Time Auditing of Domotic Robotic Cleaners

D. Papadogkonas, M. Zoumboulakis, J. Taylor, M. Levene and G. Roussos
Birkbeck College, University of London
{dikaios, mz, jenson, mark, gr}@dcs.bbk.ac.uk

Abstract

Domotic Robotic Cleaners are autonomous devices that are designed to operate almost entirely unattended. In this paper we propose a system that aims to evaluate the performance of such devices by analysis of their trails. This concept of trails is central to our approach, and it encompasses the traditional notion of a path followed by a robot between arbitrary numbers of points in a physical space. We enrich trails with context-specific metadata, such as proximity to landmarks, frequency of visitation, duration, etc. We then process the trail data collected by the robots, we store it an appropriate data structure and derive useful statistical information from the raw data.

The usefulness of the derived information is twofold: it can primarily be used to audit the performance of the robotic cleaner –for example, to give an accurate indication of how well a space is covered (cleaned). And secondarily information can be analyzed in real-time to affect the behavior of specific robots – for example to notify a robot that specific areas have not been adequately covered.

Towards our first goal, we have developed and evaluated a prototype of our system that uses a particular commercially available robotic cleaner. Our implementation deploys ad-hoc wireless local networking capability available through a surrogate device mounted onto this commodity robot; the device senses relative proximity to a grid of RFID tags attached to the floor. We report on the performance of this system in experiments conducted in a laboratory environment, which highlight the advantages and limitations of our approach.

1. Introduction

Robotic Cleaners are commercially available household devices that are typically equipped with an array of sensors, a micro-controller and some kind of programmatic logic. These components interact with each other in order to clean a physical space. To achieve covering and cleaning a space two broad choices exist: either the robotic cleaner uses its sensors to generate a map of the room first and then clean the room according to this map, or use simple algorithms such as spiral cleaning, wall-following and random-walk angle-changing after bumping into obstacles. The main difference is the time spent cleaning a room. If a physical space is mapped first, then cleaning can be near-optimal since the path of the cleaner is determined by the map. Conversely, if random paths are taken then more time is required and some spots will be inevitably visited more than others.

Our approach aims to audit the performance of robotic cleaners by examining the *trails* they select in their cleaning tasks. To achieve this, we overlay a network of wireless proximity sensors onto the physical space. While cleaning, the robotic cleaner communicates wirelessly with the proximity sensors and generates a sequence of interactions – a trail. In this sense, the trail is a metadata-enriched path that a specific robotic cleaner took at some point in time. Each proximity node can be thought of as a *landmark* – so some typical metadata fields would be: timestamp, time spent at each landmark, proximity to the landmark, number of times a landmark was visited etc. By performing trail analysis we can then evaluate the performance of each robotic cleaner and even adapt the algorithm used by the cleaner to the specific context e.g. the size and the shape of the room.

The system proposed achieves trail analysis by using a model for the representation of trails in conjunction with suitable data structures for efficient storage, filtering and retrieval. In the sections to follow we will describe in detail the different components of the system and how they interact in order to assess a trial (and furthermore audit the performance of the cleaner). We will then present results from a laboratory experiment and we will critically discuss the advantages and limitations of our approach. We conclude by reporting on our experiences and future work and by identifying other possible areas that such an approach could be useful.

In summary, this paper makes two main contributions:

1. It presents a general-purpose machine learning framework for the representation and analysis of trail-based records on interactions between pervasive computing devices within a smart environment.
2. It demonstrates how this framework and associated algorithms can be used within a practical system for auditing domotic cleaners, with particular emphasis on efficient negation querying.

2. Motivation

The motivating factor for our work stems from the fact that the effects of cleaning a physical space are not always visible to the naked eye – at least not without close scrutiny. In other words it is often hard to judge objectively the quality of cleaning. This issue becomes more relevant in large public spaces such as airports, hospitals, train stations and supermarkets.

Typically the performance of human cleaners is evaluated by supervising staff assessing the quality of the cleaning based on some criteria. This assessment process often lacks rigor and is open to bias. Our system envisages a not-to-distant future where robotic cleaners will replace the menial task of manual cleaning. Some of the commodity robotic cleaners available in the market today employ cleaning algorithms that have random elements in their programmatic logic. In environments where cleaning is of utmost importance – for instance in hospitals – a mechanism must exist to deterministically evaluate the cleaning operation using a set of well-defined criteria.

Furthermore, it is desirable to know, even in domestic environments, how well a space has been cleaned. Often a robotic cleaner will be left to operate at a home during a time that everybody is out. Our approach can provide a reporting tool assessing how well the robotic cleaner performed. Varying levels of detail can be provided such as: *hotspots* or areas that were visited (cleaned) more than others, areas missed, time spent at each room or area within a room, and so forth.

Finally, looking further into the future of smart homes, it is conceivable that blocks of flats will benefit from sharing automated infrastructures in which case it will be constellations rather than single domotic cleaning agents that will carry out this task for a collection of associated homes. This situation will complicate the auditing task considerably and would completely remove any opportunity of human-operated supervision. In this conditions, out techniques will be particularly useful and it is only fully automated methods that can provide adequate auditing of the performance of such robotic swarms.

3. Augmented Smart Homes

Robotic cleaners are designed to operate in physical spaces such as house rooms or offices. For our context, we envisage *augmented* physical spaces that essentially encompass the installation of numerous sensing components in a home or an office [1], leading to the development of advanced smart home facilities. The context of our work is the intersection of such pervasive computing, in particular wireless sensor networks, with robotics specifically applied for the benefit of smart home maintenance.

In this context, some of the sensor components would typically assist the user in controlling the environment – e.g. ambient temperature sensors that wirelessly interact with thermostat-controlled central heating or air-conditioning systems. Other sensors would monitor and assist robotic workers in accomplishing their task. For instance we can envisage a robotic cleaner that interacts with sensors in order to fully cover a space or to compose a near-optimal cleaning roadmap.

The configuration of the sensing components in the physical space can vary in terms of density, capabilities and services offered. Some of the sensors could be capable of detecting proximity of a robot, whilst others could be capable of detecting environmental conditions such as assessing the air quality of a living room and accordingly adjusting air purifiers or other similar devices. Specifically for robotic cleaners a sensing component should be capable of detecting the cleaner when it moves into range and logging the visit with a timestamp and other relevant information such as proximity information.

In this manner a simple monolithic sensing component with some limited memory can become part of a distributed sensing environment capable of monitoring movements of mobile agents and robots. This part of the sensing environment is purely passive and its main purpose is the logging of activity. Its product is a set of activity logs and other related messages that can be communicated between the robotic cleaner and the sensors. Algorithmic components of the system (that can either be part of the robotic cleaner or the smart physical space) can take programmatic action based on this information and direct the mobile agent or robot accordingly.

Such co-operation between sensing components and robotic agents can take many forms and in this paper we propose a particular approach built on trails as the main data primitive and a centralized approach to coordination. In particular, we introduce a universal model for the capture of such interactions between sensors and robots and their representation in such a way that the complete information is organized in a manner where common patterns can be quickly and efficiently identified. Individual robots do not have to process the rapidly growing information collected by the sensing components; instead this information is centrally processed by non-constrained computing equipment (such as desktop class PCs or mini-servers) and a simple actuating or information message is then sent to the individual robot.

The central processing of trail information is performed continuously and in real-time fashion. As new readings become available trails are updated to reflect the new information. The data structures used to store and process trail information are optimized for the task in hand allowing for high throughput and performance benefits. In the sections to follow we will describe in more detail exactly the data structures and the processing methods used to extract meaningful information that can be used to primarily audit the performance of the robots and secondarily enhance their performance by making decisions based upon these audit reports.

4. Trail-based Auditing and Coordination

As already mentioned, the core ingredient of our approach is the use of trails as the principal data processing primitive used for analysis and evaluation of a robotic cleaner. We formally define a trail of an individual robotic agent as the sequence of recorded interactions between itself and the nodes of a wireless sensor network. Trails contain patterns of actions and they can be used for the provision of different services, spatial analysis or navigational assistance. A downside of trail analysis is that it requires considerable storage and computational resources to discover such patterns. Moreover no single method exists that identifies significant trails based on different metrics related to a particular application.

In this paper we introduce a trail-based analysis approach, an associated model for the representation of trails and trail aggregates and suitable data structures for efficient storage, filtering and retrieval.

To identify specific types of interactions we introduce the notion of the landmark as the position of a significant object in a mixed pervasive-robotic environment. A sequence of interactions recorded in trails represent wireless communications between a robot and the sensor node located at a specific landmark which are associated with both spatial and semantic attributes induced by the embedding of sensor components in physical space.

The choice of trails as our representation of a sequence of interactions is not coincidental. Indeed, trail records have been used as the basis for coordination between humans for centuries in different forms. For example, navigation trails provide route information and record information about paths to potential destinations. Aggregating multiple trails acquired over time across a particular environment is the technique humans often use to develop complete maps of a particular landscape and subsequently assist navigation, especially in the context of exploration [2]. Oral trails are also quite common in human coordination and are best represented as narratives which are replayed and recast repeatedly to incorporate new knowledge [3].

Trails have also been used with great success in assisting coordination of information seeking activities on the Web [4]. Our work is intimately related to this approach and proposes extensions to this model of collective experience aggregation to cater for the distinct requirements of environments that mix wireless sensor networks and robotics.

Sometimes trails are constructed based on imperfect or ambiguous information, and in literature they are often referred to as *tracklets* (although the term tracklet can refer to other related notions too). In this context, a tracklet is often an estimate of a trail and it is based only on a few measurements or interactions that are available rather than the full trail or path that the agent or robot has followed. One of the uses of tracklets is in mobile target recognition [5] and path disambiguation.

Our work is based on a model that organizes the complete information captured by one or more members of a robotic cleaning team into a unified representation. We then propose a number of metrics that can be used to identify the most important of these trails that can be of value when carrying out a particular task such as cleaning. Such so-called significant trails are stored, disseminated and used to provide guidance and assistance in a variety of coordinated tasks.

5. High-level description of the coordination engine

In order to construct the coordinator we need to introduce the interaction network that is a directed graph where vertices represent sensor network nodes and edges represent paths between these nodes. Two vertices (nodes) are said to be connected when there is a corresponding interaction record that indicates that the two landmarks have been visited in sequence by at least one of the robots. For the remainder of the discussion we will use the terms landmarks, nodes and vertices interchangeably. Similarly we will use the terms edges, paths, links or trails to represent the connection between landmarks.

The links between landmarks are always directed and the landmarks are weighted with different usage metadata. Example of the metadata fields can include a unique id for the robot, a timestamp representing when the robot came into range, a timestamp representing when the robot went out of range, a positive integer representing frequency of visitation, the distance between the robot and sensor node during their interactive session, the orientation of the robot in relation to the sensor node, etc. Higher order metadata can also be used, for instance the compound probability that a link will be followed given the robot has arrived to a specific node following a particular path of fixed length within the network.

It should be noted that not all sensor nodes will be capable of providing all this information, nor does it make sense to store all the possible metadata fields for all types of landmarks. In any case, calculating the weights from the raw system logs requires considerable computational effort and poses several challenges in reconciling and ordering the log records. A few examples of the challenges involved, stem from the fact that the sensor network is likely to be heterogeneous and with only approximate time synchronization. Different nodes will have varying capabilities in identifying a specific robot and expectedly will be recording different information or use different format to record the same interaction information. Similarly, robots may vary in their logging capabilities too. It is the purpose of our system to disambiguate the information received from the robots and pre-process it in a manner that no conflicting transactions exist. This is somehow semantically related to concurrency control protocols and serialisable transaction schedules in distributed database environments.

To explain this with an example if a robot visits nodes *A, B* and *C* in this order, there should never be any margin of computational or communication error that would allow this sequence of interactions to be recorded in any other order (such as *A, C, B*) and inadvertedly change the structure of the graph.

In terms of this graph which is central to our approach trails are represented as sequence of nodes. For efficiency the graph is stored as a probabilistic suffix tree [6] enhanced with metadata needed to encapsulate different information and metrics relevant to each interaction [7]. An example of such a representation is shown in Figure 1. The choice of this data structure has been guided by our desire to develop a system that can maintain all captured information while being able to rapidly respond to a great variety of queries, virtually being capable of responding to requests about any number of possible time, space and semantics related criteria.

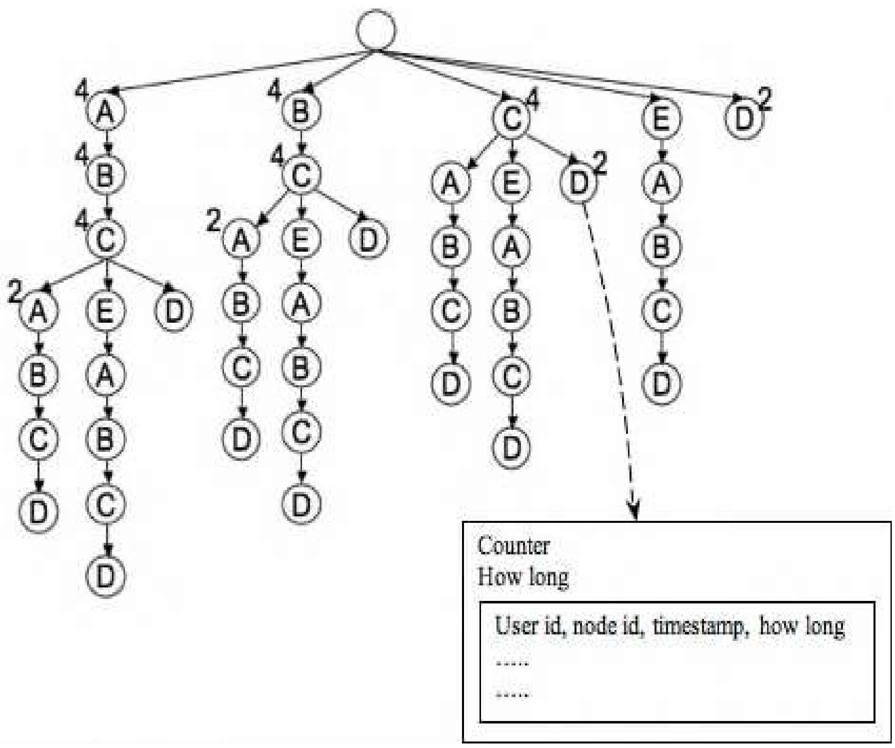

Fig.1 - Suffix tree containing two trails, namely ABCABCD and ABCEABCD.Each letter represents a wireless sensor nodes with which at least one robot has interacted with. Each node is also associated with additional metadata used to answer and rank significant trail queries.

The complete process by which interactions are logged by robot team members and aggregated to the data structure is shown in Figure 2.

A trail represented by a sequence of interactions between nodes is said to be significant if it satisfies one or more of the following criteria:

- It is one of the top *n* trails in respect of trail popularity.

- It is one of the top *n* trails in respect of average time (or some other temporal statistical measure) spent interacting with the sensor nodes (landmarks) in the trail.
- It is one of the top *n* trails in respect of the relevance of the landmarks to some chosen semantics (for example related to a specific spatial sub-area of the physical space that carries some possibly arbitrary user-defined significance).
- It is one of the top *n* trails in respect of one of the above criteria for a chosen team sub-grouping.

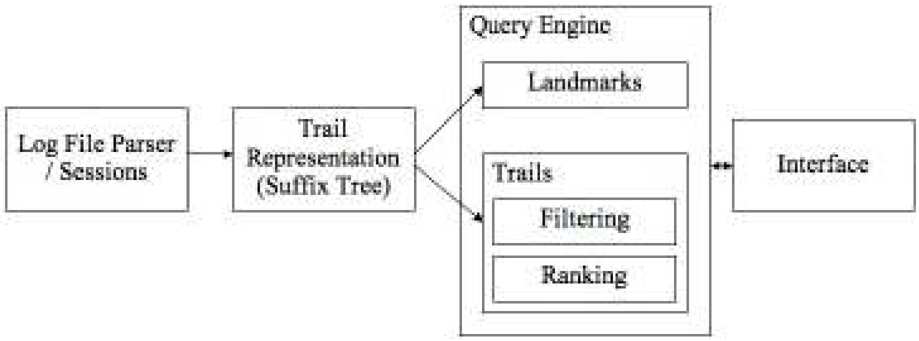

Fig 2 – Coordination engine architecture. From left to right, a log file is parsed for correctness and in order to ensure that there are no violating transactions. Then the parsed interaction logs are transformed to our representation – namely the Suffix Tree. The next step is the filtering and ranking of the trails. On the far right, the User Interface can be used to submit arbitrary queries that will in turn consult the data structure to formulate the results.

In addition, the above criteria can be combined and weighted for example one of the top *n* trails in respect of trail popularity that had a duration that exceeded a given threshold. Significant trails are identified within a particular class of trails – for example within the class of trails that share a specific landmark as their starting point and another one as their ending point (we assume that many possible paths exist between two given landmarks). Sub-classes can be formed by refining the selection to extract all the trails that satisfy a relational condition for some fixed period of time (find a trail that lasted more than or less than a given time value). Other classes can be formed by finding cycles e.g. trails that started and ended at the same landmark.

To extend the context to other applications outside cleaning, classes of trails can be found for a mobile agent or robot by selecting trails that occur at a specific time – for example in the morning, the afternoon, when the house is empty, or during a period that unusually high temperatures are recorded. Significant trails can be inferred from the interaction network using a variety of tree traversal methods. We are proposing mechanisms to efficiently compute significant trails based on a fine grain heuristic defined on the interaction network, extending similar techniques already developed for web navigation.

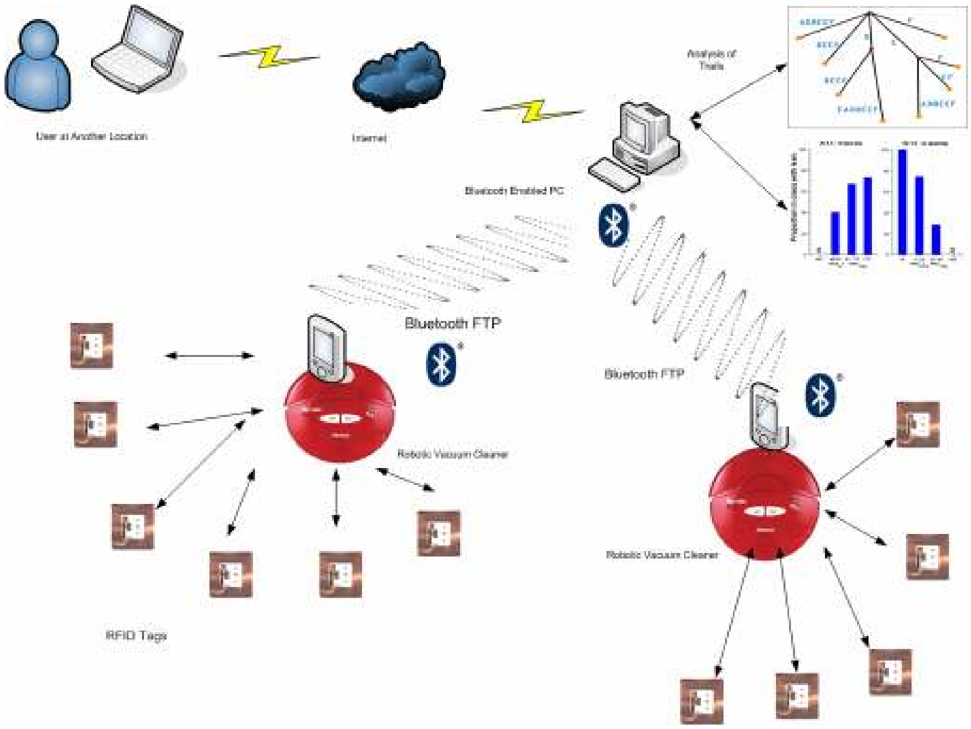

Fig. 3 – High-level depiction of a version of our system. The round objects represent the robotic cleaner which is augmented with the surrogate device that reads and logs interactions with the RFID sensors (square objects in the figure). Information is transferred wirelessly via Bluetooth (although other alternatives may be used e.g. Wi-Fi, Zigbee, etc) to a desktop-class PC or a mini server. The PC or mini-server continuously processes and filters the information, stores in the probabilistic suffix tree and computes significant trails and other statistical information that can then be accessed by any user (e.g. a user at his office can monitor how "well" a robotic cleaner is cleaning her home).

6. Experimental Evaluation

In order to test the feasibility of this proposal we have implemented a simple system to collect and analyze robotic trails interacting with simple wireless proximity sensors (Figure 3). In this case, we only analyzed the aggregate of the collected data and we have not used it to modify the behavior of the individual robots in response to significant trails. We have observed that this approach is effective and efficient and we believe it can be easily extended to other types of wireless sensor networks.

The sensing component of our prototype (Figure 4) is provided by simple ISO 14443 RFID proximity tags. For our robotic platform we have selected the iRobot Roomba cleaner [17] primarily due to its cost advantages. As the capabilities of the Roomba are not adequate for this experiment we have augmented it with a surrogate device based on our mobile sensor node prototype [18] instrumented with an ACG multi-protocol RFID reader [19]. In more detail, the different components of the system play the following roles:

- *Surrogate device*. The Shared Memories surrogate device sits on top of the Roomba robotic cleaner and interacts with the proximity tags. Each interaction is timestamped, recorded and transmitted wirelessly to the coordination server in real-time or batch mode in case when the robot is not within range of an access point.
- *Proximity Sensing Tags*. We use proximity tags in a form of printable adhesive roll that can be stuck to or placed on almost any surface. This allows for rapid deployment of a system in virtually any environment.
- *Coordinator.* The Shared Memories coordinator runs on a server with both wired and wireless connectivity. The analysis software also provides a graphical user interface to display on the fly significant trails for observation and experimentation.

Fig. 4 – iRobot Roomba Robotic Cleaner with Surrogate Device mounted on top of it. Also visible in the picture, are the RFID tags attached to the floor, forming a proximity grid.

The experiments were conducted in the first floor of London Knowledge Lab [20]. The main obstacles for the Roomba robotic cleaner were the desks, chairs and other office fixtures and fittings. Initially, the robot was placed in the center of the room as it is described in the user instructions and it was let to run for approximately one hour. We have used the time frame of one hour based on the size of the room and the approximate time it takes a human operator to clean the room (approximately 20 minutes).

One of the aims of our experiment was to evaluate the claims made by the iRobot corporation that Roomba covers (e.g. cleans) physical spaces almost in their entirety (excluding staircases, split-levels, etc). The focus of our investigation is based on the most visited landmarks - they can be thought of as *hotspots* of activity – and their relationship to trails followed and the corresponding popularity of each trail (e.g. how often it was followed).

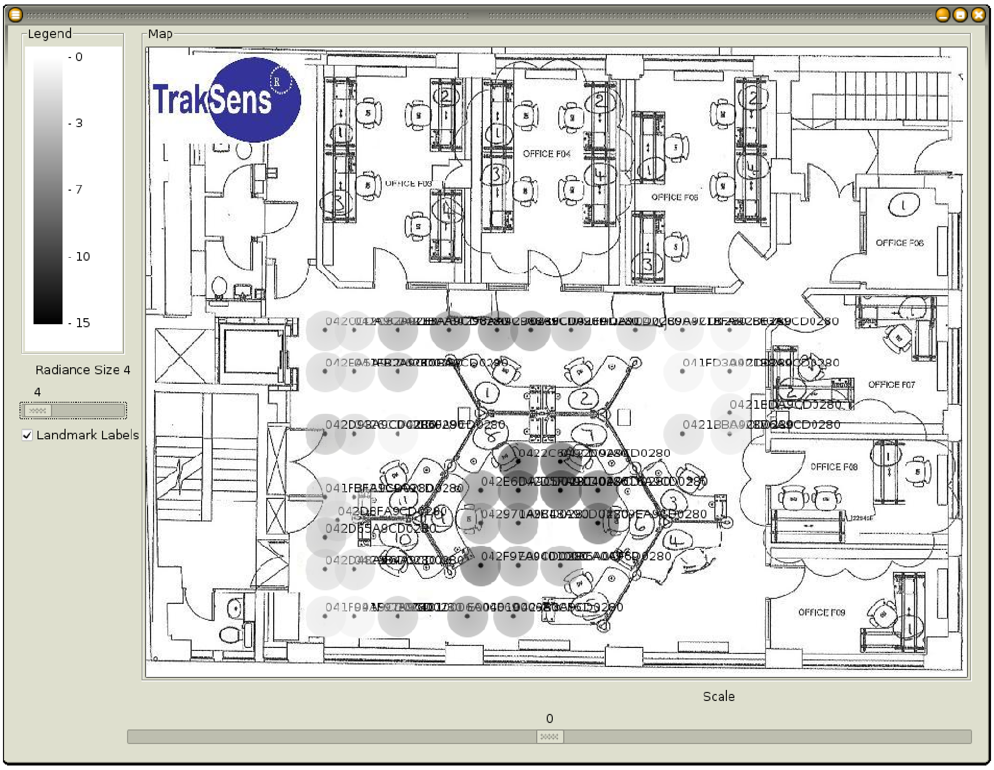

Fig. 5 – Screenshot of the GUI of our system – the office plan where the robotic cleaner was deployed is shown together with hotspots of activity (represented by circles). Darker circle colors represent landmarks that were visited more often than others.

Figure 5 depicts the top landmarks – the *hotspots* of activity – that were visited by the robot and Figure 6 shows the overall trail followed. In the overall view of the dataset it is fairly easy to spot the flaw of the robotic cleaner: namely the failure to exit the cyclic arc-shaped enclosure formed by the offices, chairs and other obstacles.

In addition, we have extracted from the data set the best trails by the popularity of trails visited in order to investigate the flaw of the robot's navigation algorithm. In Figure 7 we present the most popular landmarks visited by the robotic cleaner.

4. Related Work

Robotic cleaners have many potential applications both for domestic and commercial environments. Large physical spaces such as supermarkets and airports can benefit from automating the cleaning process. Naturally with such growing interest from both industry and general public, the research area surrounding robotic cleaners has been very active for some time.

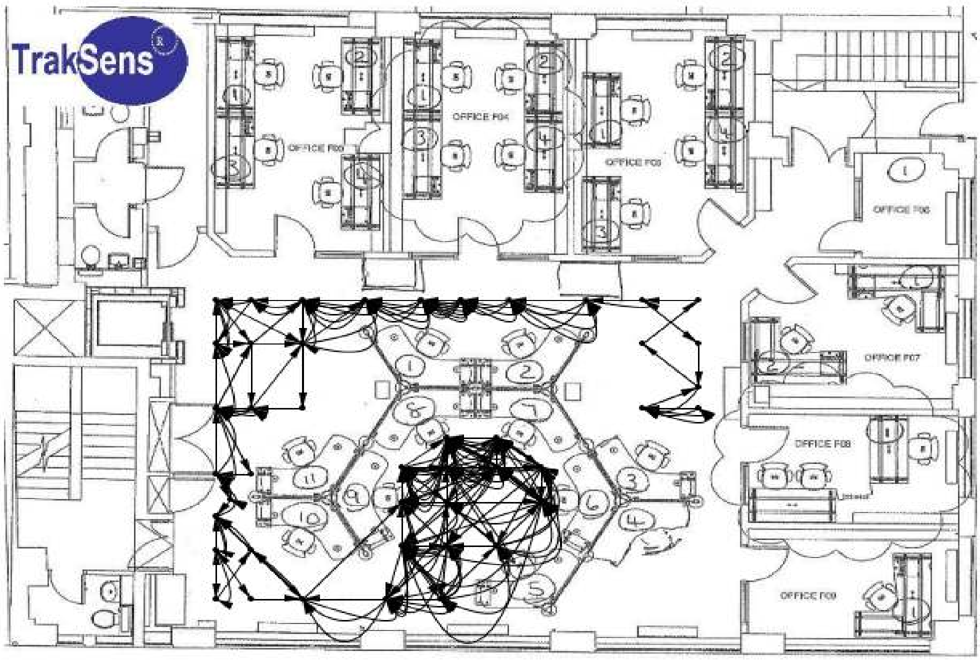

Fig. 6 – This figure shows the trails followed by the robotic cleaner – represented by the arcs and arrows. It is clear that within the time frame given (one hour) many spots within this space were not covered at all.

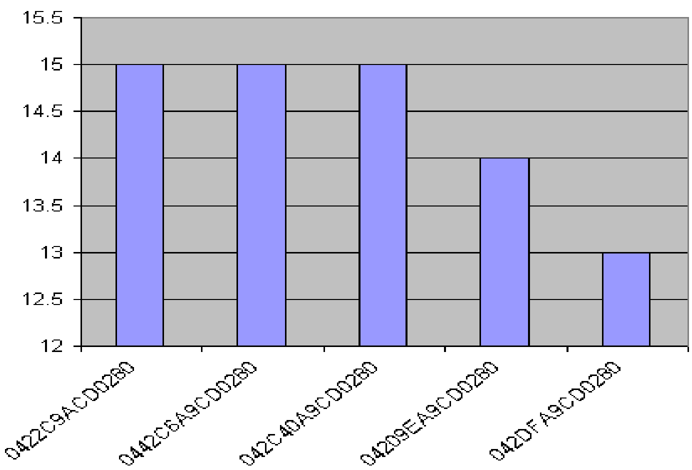

Fig. 7 – This bar chart shows the top five (most popular) landmarks. The X axis shows the unique ID number of the RFID tag and the Y axis shows the number of times that each landmark was visited.

In particular, the marriage between wireless sensor networks and robotics has attracted considerable attention [8]. The majority of this work is focused on two main problems: first, the introduction of mobility within the sensor network as a means to improve robustness in the presence of node failures and the increase of routing performance and coverage; and second, the local use of the sensing elements by robots as a means of improvement of the information they hold about their immediate environment. To the best of our knowledge there are only a few attempts to employ the wireless sensor network as a way of coordinating teams of robotic agents by employing its ad-hoc communication capability in addition to its sensing. Such an approach is presented by [15] where the authors describe their probabilistic navigation model based on the information collected by a pre-deployed sensor network; the advantage of this approach is that it can cater for dynamically changing environments (such as airports).

Notably, rather than construct a universal model of the aggregate experience gained by a team, the approach by [9] proposes the use of embedded RFID sensor nodes as a means to store discoveries made by specific individuals to be employed as pointers by subsequent visitors to these sites. This approach has the distinct advantage that a robot team can coordinate in an entirely decentralized manner. Nevertheless, it fails to capitalize on ample ad-hoc networking opportunities.

In the more general research area of robotic navigation, there is plenty of work that mainly focuses on near-optimal space coverage. Some of the techniques proposed are employing complex machine learning algorithms such as the neural network approach proposed by [10-11] and the fuzzy logic approach proposed by [12]. Other approaches such as [13] support dynamic assignment of cleaning polygons in a physical space; such an approach has the advantage of resilience; if one robot breaks down, its assigned polygon can be given to another robot. Remaining in the research area of coverage, the approach by [14] suggests navigation should be based in the pre-construction of a cell-based map where each cell is a triangle. In this manner, a specific robot has more flexibility since it can move in more directions and avoid obstacles.

Within our own group, a similarly decentralized coordination algorithm has been proposed [16], with a view to provide coordination opportunities in search and rescue operations conducted by robot teams. Similar to [9], this technique uses RFID tags embedded in the environment to relay information to other team members regarding the history of operation within some particular vicinity.

5. Future Work and Conclusions

We introduce a trail-based approach to auditing the performance of robotic cleaners; our system can be extended to monitor other robotic or mobile agents that are designed to operate unattended. Central to our approach is the trail Coordination Engine that makes use of ad-hoc wireless networking opportunities and the availability of wireless sensor networks embedded within the environment of operation. We have outlined the operation of our coordination engine which processes the interaction histories of individual team members namely domotic robotic cleaners, with the sensor network to construct probabilistic representations of the collective experience that can be used to provide the analysis for any query related to the team activities.

Such patterns of past experience can be used to identify areas that require further attention or optimize the collective behavior of the team. We are currently working on developing this passive coordination approach into an active coordination mechanism. In particular, we are investigating classification algorithms that allow the prediction of individual behaviors and thus provide a way to steer the team into areas that require additional attention in a manner that is both effective and efficient. From an application perspective, we are particularly interested in combining out techniques with delay-tolerant networking approaches to improve the robustness and accuracy of our trail collection performance, which is expected to have a significant effect on the capability to successfully anticipate coordinator: most significant path computed, and per-landmark visitation statistics.

Our proposed model will have direct applicability to any domestic or commercial robotic or otherwise unattended cleaning operation. We envisage that our system will be particularly useful to large smart homes and blocks of flats with shared cleaning infrastructures, but its usefulness will extend to other smart building environments such as hospitals, airports, and supermarkets. The effects of cleaning are not always perceptible through visual inspection and it often requires close scrutiny to determine how well a space has been covered. Moreover, this process will always be prone to human error and bias. We aim to remove the need for human supervision by providing the system described that will automatically audit and deterministically report on the quality of a cleaning operation. Finally, our approach can be extended to other areas outside cleaning where robots or mobile agents may operate.

References

[1]. MIT, *House_ n Project*, http://architecture.mit.edu/house_n/

[2]. R. Golledge, *Wayfinding Behavior: Cognitive Mapping and Other Spatial Processes*, The John Hopkins University Press, 1998.

[3]. M. Mateas and P. Sengers, *Narrative Intelligence*, John Benjamins Publishing, 2003.

[4]. J. Borgers and M. Levene, *Evaluating Variable Length Markov Chain Models for Analysis of User Web Navigation Sessions,* IEEE Transactions on Knowledge and Data Engineering, Vol. 19, Issue 4, April 2007.

[5]. Drummond, Oliver E.; Blair, W. D.; Brown, George C.; Ogle, Terry L.; Bar-Shalom, Yaakov; Cooperman, Robert L.; Barker, William H., *Performance Assessment and Comparison of Various Tracklet Methods for Maneuvering Targets,* in Proceedings SPIE, Volume 5096, pp. 514-539, 2003.

[6]. P. Weiner. *Linear Pattern Matching Algorithms*. In Proceedings of the. 14th IEEE Annual Symp. on Switching and Automata Theory. pp. 1-11. 1973.

[7]. D. Papadogkonas, G. Roussos and M. Levene. *Discovery and Ranking of Significant Trails.* 2nd Int. Workshop Expl. Context History in Smart Environments (ECHISE 2006). Irvine, CA, 17 September, 2006.

[8]. H. Chen, B-C. Cheng, C-C. Cheng, L-K. Tsai. *Smart Home Sensor Networks Pose Goal-Driven Solutions to Wireless Vacuum Systems*. In Proceedings of the 2006 Int. Conference in Hybrid Inf. Tech. (ICHIT'06). Vol. 2. pp. 364-373. 2006.

[9]. M. Mamei and F. Zambonelli. Pervasive *Pheromone-Based Interaction with RFID Tags. ACM Trans. Autonomous and Adaptive Systems*. 2007 (in press).

[10]. S. X. Yang and C. Luo, *A Neural Network Approach to Complete Coverage Path Planning*, in IEEE Transactions on Systems, Man, and Cybernetics—Part B: Cybernetics, Vol. 34, No. 1, February 2004.

[11]. C. Luo, S. X. Yang, and X. Yuan, *Real-time area covering operations with obstacle avoidance for cleaning robots*, in Proceedings of Int. Conference on Intelligent Robots and Systems, October 2002.
[12]. Y. Fu and S. Y. T. Lang, *Fuzzy-Logic based mobile robot area filling with vision system for indoor environments*, in Proceedings of the IEEE Int. Symposium on Computational Intelligence in Robotics and Automation, pp.326-331, 1999.
[13]. M. Jager and B. Nebel, *Dynamic Decentralised Area Partitioning for Cooperating Cleaning Robots,* In Proceedings of IEEE Int. Conference on Robotics and Automation, Washington DC, 2002.
[14]. J. Oh, Y. H. Choi, J. B. Park and Y. F. Zheng, *Complete Coverage Navigation of Cleaning Robots Using Triangular Cell-based Map,* in IEEE Transactions on Industrial Electronics, Vol. 51, No.3, June 2004.
[15]. M. A. Batalin, G. S. Sukhatme and M. Hattig, *Mobile Robot Navigation using a Sensor Network,* in Proceedings of IEEE Int. Conference on Robotics and Automation, LA, 2004.
[16]. E. Ferranti, N. Trigoni and M. Levene. *Brick & Mortar: An Online Multi-Agent Exploration Algorithm.* IEEE Int. Conf. Robotics and Automation (ICRA). 2007.
[17]. T.E. Kurt, Hacking Roomba: *ExtremeTech*. Wiley. 2006
[18] G. Lane, C. Brueton, D. Diall, D. Airantzis, N. Jeremijenko, G. Papamarkos, G. Roussos and K. Martin. *Community-based Public Authoring with Mobile Chemical Sensor Networks*. Proc. IEE Intelligent Environments. 5-6 July, Athens, Greece. 2006.
[19]. ACG. *ACG Multi-ISO RFID Reader Manual*. Available via http://www.acg.de/
[20]. *The London Knowledge Lab,* http://www.lkl.ac.uk/
[21]. How Stuff Works, http://www.howstuffworks.com/

Authors

Dikaios Papadogkonas. Dikaios is a PhD candidate at Birkbeck College, University of London. His main research involves deriving useful information about users' activities in physical spaces with the aim of improving navigation.

Michael Zoumboulakis. Michael is a PhD candidate at Birkbeck College, University of London. His main research involves the investigation of data management algorithms for wireless sensor networks.

Jenson Taylor. Jenson is a PhD candidate at Birkbeck College, University of London. His main research explores participatory sensing in ubiquitous environments and the intersection of robotic and pervasive environments.

Mark Levene. Mark is a Professor at Birkbeck College, University of London. His main research interests involve Web Search and Navigation, Web Mining and more generally Web Interaction.

George Roussos. George is a Senior Lecturer at Birkbeck College, University of London. His main research interests involve experimental systems in ubiquitous computing and the effects of such systems on human activity.