

A Taxonomy of Collaborative Context-Aware Systems

As'ad Salkham, Raymond Cunningham, Aline Senart and Vinny Cahill

Distributed Systems Group
Department of Computer Science
Trinity College Dublin, Ireland,
{Salkhama, Raymond.Cunningham, Aline.Senart, Vinny.Cahill} -at-
cs.tcd.ie,
WWW home page: <http://www.dsg.cs.tcd.ie/>

Abstract. Context awareness is a vital element in pervasive and ubiquitous systems. While most existing research has focused on designing context-aware systems to integrate into the environment, less attention has been placed on the interoperability among the entities comprising such systems. In this paper, we consider how the components of a context-aware system can collaborate to achieve a common goal. We provide a taxonomy of such Collaborative Context Awareness (CCA) based on three axis, i.e., goal, approaches and means. We also discuss a number of context-aware systems from different domains, i.e., augmented artefacts, robotics and sensor/(actuator) networks that exhibit some form of collaboration. Finally, we classify the different studied systems according to our taxonomy.

Keywords : taxonomy, collaboration, Collaborative Context Awareness, context awareness

1 Introduction

Context-aware systems are an emerging genre of computer systems that help add some forms of intelligence to our surroundings. It is well-established that context-aware (sentient) systems should address three basic requirements, i.e., sensing, inference and actuation [1]. Furthermore, a number of ongoing research efforts have been targeting the definition and classification of context-aware systems; the most recent is a survey on context-aware systems [2] that suggests a set of common design principles and assesses a set of available context-aware middlewares and frameworks against those principles. In this paper, we do not provide a classification nor a survey of context-aware systems. Readers are encouraged to refer to [2–6] for such systems. Instead, we focus on collaboration and its intrinsic relation to context awareness.

To our knowledge, there is no existing research providing a taxonomy or a concrete definition for collaborative context-aware systems that range from small augmented

The authors are grateful to Science Foundation Ireland for their support of the work described in this paper under Investigator award 02/IN1/I250 between 2003 and 2007.

artefacts to large-scale and highly distributed sensor(/actuator) networks. Mäntyjärvi *et al.* [7] emphasise reliability in what they refer to as *collaborative context recognition*. Additionally, they presume that context-aware devices within a certain area have common views of the context and can agree on a time- and space- dependent *collaborative context* through short-range communication. They describe *collaborative context* as the “summary of the situation of the other devices in the local range corrected by the local context” providing an update strategy for these devices and associated trigger conditions. In [8], a context-aware communication platform to support smart objects is described. The platform emphasises the importance of a distributed tuplespace-based communication model to support inter-object collaboration. The model allows for smart objects within a vicinity to share a distributed tuplespace, broadcast their data and contribute with an equal amount of memory.

An interesting taxonomy for coordination in Multi-Robot Systems (MRS) was presented by Farinelli in [9]. The MRS taxonomy is divided into four levels namely, *cooperation*, *knowledge*, *coordination* and *organisation*. However, this taxonomy considers only cooperative systems and defines them as those systems composed of “robots that operate together to perform some global task” [9].

In [7] and [8] for instance, the focus is solely on collaborative context recognition and how reliable and consistent the outcome of this recognition is. Neither actuation nor decision making is taken into account and only handheld devices are considered. The MRS taxonomy [9] does not address context awareness but rather cooperation of robots. The robots are said to be *aware* if they could have some knowledge of their team members [9], this view is not well-defined and lacks a definition of the relevant knowledge nature. Furthermore, the MRS taxonomy only considers a coordination protocol-based/protocol-free decision making process and does not include data and context sharing.

We define a collaborative context-aware system *as a system that comprises a group of entities, capable of sensing, inferring, and actuating that communicate in order to achieve a common goal*. We have identified that collaboration among context-aware entities may not only be based on communicating contextual information but also sensed and fused data in addition to possible next actions to perform. Such communication supports efficient collaboration to occur as a result of more precise inference, decision making and awareness. Moreover, collaborating components can follow a *consensus-based* or a *consensus-free* approach in which negotiated decisions or local decisions are taken respectively. *Delegation* may also be used to achieve optimal behaviour in collaborative context-aware systems. We see *delegation* as the ability of entities to pass tasks to neighbours depending on their estimation of the best option to achieve a specific common goal. Our contribution is presented in the form of a taxonomy for Collaborative Context Awareness (CCA) that encompasses the characteristics we highlighted above.

The remainder of the paper is structured as follows: Section 2 describes the context-aware systems from different domains that we studied while emphasising their components collaboration. Section 3 presents our taxonomy for CCA. In Section 4, we provide a classification of the studied systems according to our taxonomy. We conclude in Section 5.

2 Studied systems

Our interest in the following systems stems from the fact that they directly address collaboration while also exhibiting context awareness. However, the definition of collaboration in different systems may result in a philosophical debate. Some researchers may prefer to divide collaboration into cooperative (i.e., negotiated decision making through communication) and coordinated (i.e., local decision making through communication) solutions, as in [10] for instance. Others consider coordination as cooperation in which an agent performs actions while taking into account the actions performed by other agents [9]. We tend to see collaboration as a synonym for cooperation and we do not focus on the differences among collaboration, cooperation and coordination. Nevertheless, our CCA taxonomy is flexible enough to encompass various views and opinions. In this section, we span several domains that are the most representative of CCA. In parallel, we present the systems we studied and provide an analysis of their characteristics, in particular their collaboration models.

2.1 Augmented artefacts

One of the best-known projects in the augmented artefacts domain is Smart-Its [11] which was part of the EU-funded Disappearing Computer Initiative. The idea behind the project is to develop smart small-scale embedded devices known also as Smart-Its that are able to sense, actuate, compute and communicate. These smart devices are introduced to help develop and study collective context awareness and increase the widespread deployment of ubiquitous computing systems [12]. Furthermore, the Smart-Its project encompasses different everyday artefacts (e.g., cups), enabling greater and more user-friendly perception of the surrounding environment.

Smart-Its are generic and exhibit a modular design; in a single Smart-It, an application-specific processing module forms a bridge between the communication module and the different sensing and actuation modules. In addition, the inter-communication of Smart-Its is based on a stateless peer-to-peer protocol where local broadcast of context and other information within a certain proximity is supported, (i.e., a direct communication scheme) [11]. Application-dependent fusion and inference techniques in a Smart-It are likely to occur in the processing module while a lower level of fusion could be carried out in the sensor module(s). Also, the generic nature of Smart-Its enables them to form different kinds of systems, e.g., common goal-oriented systems that might implement a consensus-based and/or a consensus-free approach.

A spin-off of Smart-Its was Smart-Its Friends [13, 14] that emerged from applying the concept of context proximity to connect Smart-Its. This concept enables Smart-It devices within range and experiencing similar situations or conditions (e.g., same shaking pattern), to be considered near to each other in context (i.e., have a common context perception) and to be connected or friends. Upon the reception of broadcast data and the ID of the source, a Smart-It tests the data against a predefined threshold and declares the source as a friend if the test is passed. Subsequently, a Smart-It will remain identified as a friend regardless of the connection breaking while continuing to respect the friendship expiry constraints [13].

Another related project is Cooperative Artefacts [15]. The concept is based on an infrastructure-less approach to allow easier deployment of these artefacts when cooperatively assessing specific situations in the environment. The cooperation among artefacts is solely based on sharing their knowledge through protocol-based direct communication. The knowledge base in each artefact comprises three types of knowledge; domain knowledge, observational knowledge and inferred knowledge. In an artefact, facts are defined as “the foundation for any decision-making and action-taking within the artefact”, while rules allow inference of advanced or upgraded knowledge based on facts and other rules [15]. Key knowledge is inferred knowledge that is the knowledge inferred from previous facts. These facts are based on the previously mentioned three local types of knowledge and/or the knowledge shared by cooperating artefacts. It is worth mentioning that there are also actuation rules that are responsible for triggering a corresponding action. In one application [15], a group of chemical containers were modelled as Cooperative Artefacts equipped with infrared light sensors and ultrasonic sensors to ensure that all artefacts (containers) are within an approved safety area and that certain artefacts stay within an acceptable distance of each other. Actuation is simplified in this application to the control of LEDs to raise an alarm if safety constraints are not met. Clearly, each Cooperative Artefact is capable of sensing, perception (fusion), inference, actuation, direct communication and sharing of the knowledge.

In [16], Ricci *et al.* present an idea that is partially inspired from the stigmergic form of communication in nature [17]. This idea is based on Coordination Artifacts that are defined as “entities used to instrument the environment so as to fruitfully support cooperative and social activities of agent ensembles” [16], for instance, street semaphores, blackboards and maps. Coordination is also seen in its very general concept as the “management of dependencies among separate activities” [16]. On the other hand, no distinction is made between coordination and cooperation. In contrast, the distinction is made between what is referred to as *subjective* and *objective* coordination. In the *subjective* form, coordination is perceived as an individual activity where the environment is not part of the coordination and the coordination aims to achieve a subjective goal through direct inter-entity communication. *Objective* coordination uses mediators that are part of the environment to decouple communication between entities to enable these entities to achieve a common or global goal. *TuCSon* [18] is an open source coordination infrastructure that uses the Coordination Artifacts concept. Examples of these artefacts in *TuCSon* are, mailboxes and blackboards for communication, tuple centre (i.e., programmable tuple spaces) for knowledge mediation and resource sharing. We can see these artefacts as communication mediators for a number of agents or entities which as a whole can form a context-aware system.

2.2 Robotics

The robotics domain is a very important source of inspiration for collaboration schemes and context-aware systems. In this section, we discuss a number of robotic systems that emphasise their collaboration aspects.

In the cooperative sensing field, Grocholsky *et al.* [10] propose a scheme for anonymous cooperation in robotic sensor networks. The scheme involves a decentralised architecture that enables entities to globally and anonymously cooperate in sensing

without the need for global knowledge. The idea is based on a Decentralised Data Fusion (DDF) algorithm. The DDF is seen as a decentralised alternative approach to the typically centralised Kalman filtering and Bayesian estimation techniques. It provides means for fusing information in a distributed network of sensors. Moreover, a DDF node depends on data gathered from a group of sensors to generate estimates of some time varying state that may then be propagated. Aggregation of information in a single node is the fusion of local sensor data, local predictions and the directly communicated information (estimates) from other nodes. Based on this fused information, a subsequent decision is taken locally by the node. Actuation is not clearly described in this architecture but can be presumed to exist since its robotic nature involves mobility that is likely to be controlled by inferred actions.

The authors break what they describe as *Anonymous Collaborative Decision Making* into a *coordinated* and *cooperative* solutions providing definitions for both. They perceive a *cooperative* solution “to be a predictive jointly optimal negotiated group decision in the sense of a Nash equilibrium¹” [10]. On the other hand, “in *coordinated* solutions there is no mechanism for this negotiated outcome. Decision makers act locally but exchange information that may influence each others’ subsequent decisions” [10]. Cooperative and coordinated solutions can be seen here as a consensus-based and a consensus-free solutions respectively.

Millibots [19] is a research project at the Robotics Institute in Carnegie Mellon University that, as the name conveys, has designed “a team of heterogeneous centimetre-scale robots which coordinate to provide real-time surveillance and reconnaissance” [19]. Millibots exhibit a modular design in sensing, processing (including inference) and mobility that implies actuation. The sensor modules that a team of Millibots can utilise may include short/long range sonar, directional infrared and vision. Subsequently, each team member specialises based on its sensor module(s). Hence, *collaboration* is seen as the consequence of distributing functionality and resources (i.e., specialised sensing and processing). This *collaboration* is defined as “the explicit exchange of information between members of a team”, this is clearly done using direct communication. Also, *collaborative sensing* is presented as being “where the sensing process itself is distributed between one or more robots”; this view is believed to be a consensus-free one that does not involve negotiated decisions. One of the interesting applications of the Millibots team is collaborative mapping, i.e., team members are able to collaborate to collect and fuse sensory information (based on a Bayesian technique) in order to create a map of the encompassing area.

As for schemes inspired from nature, swarming has inspired the creation of many systems. For instance, Parunak *et al.* [20] provide their view of collaborative sensing through swarming of multiple Unmanned Aerial Vehicles (UAV) noticeably for military imaging applications. They define different types of coordination that can occur within a group of UAVs: spatial, temporal, and team coordination. Spatial coordination is concerned with efficiently distributing UAVs over an observed area while temporal

¹ Named after Nobel Laureate (in economics) and mathematician John Nash. We quote Roger McCain’s definition: “If there is a set of strategies with the property that no player can benefit by changing her strategy while the other players keep their strategies unchanged, then that set of strategies and the corresponding payoffs constitute the Nash Equilibrium”.

coordination ensures the timeliness of all UAVs' behaviour and information exchange. Team coordination is basically inspired from natural systems (e.g., colonies of social insects), and aims at optimising distribution of roles among UAVs and managing their formation, maintenance and dispersion.

There are three principles and techniques that are needed to achieve collaborative sensing in this context. First is team and role coordination which comprises dynamic entity classification and dynamic role activation. The dynamic classification enables adding or removing roles to or from an entity while dynamic activation enables changing roles over time within an entity. Second is local optimisation where each UAV is assumed to make local decisions (i.e., consensus-free), in order to accomplish the overall designated mission on time, i.e., UAVs are capable of reconfiguration based on their perceived quality in which the goal is achieved, for example the image quality. Third are the techniques inspired from natural systems, for instance, stigmergy [17]. Basically, the scheme described employs the idea of digital pheromones from which maps are formed to enable real-time path planning.

2.3 Sensor(/actuator) networks

Sensor networks are a rich domain for studying and experimenting with new collaboration models for context awareness. This domain becomes more challenging when actuators are also involved, such as in [21]. Sensor(/actuator) networks applications span many fields for instance, surveillance/tracking systems usually for the military [22, 23] and Intelligent Transportation Systems (ITS).

Melodia *et al.* [21] provides a framework in which a *sensor-actor* coordination model in Wireless Sensor and Actor Networks (WSAN) is specified. This model is based on *event-driven clustering* of sensors and actors, i.e., a cluster is created on-the-fly after being triggered by an event. Besides this model, the framework encompasses an *actor-actor* coordination model. The *sensor-actor* coordination occurs whilst establishing data paths between sensors and actors. On the other hand, *actor-actor* coordination occurs when actors coordinate to make an optimal decision to perform the action, i.e., consensus-based. Furthermore, a cluster emerges only in a single event area where a group of sensors send their data to the same actor/collector that, as part of its role, centrally fuses the gathered data. A notion of reliability is also introduced in terms of reliable packets. This notion depends on a latency bound and a reliability threshold [21]. Moreover, sensor-actor coordination is based on a distributed protocol of localised routing decisions, i.e., consensus-free. The protocol assumes that the sensor is aware of its position, neighbours' positions, and actors' positions. Each sensor node is also governed by a multi-state protocol for optimal operation, i.e., energy consumption, reliability, etc. Concerning actor-actor coordination, a localised auction protocol is proposed. The protocol is inspired from a real-time auction protocol that defines the behaviour of actors participating in transactions as buyers and sellers. The protocol is consensus-based and designed to deal with selecting the best actor in an overlapping area of actors. This model implies some form of delegation in which the most suitable actor is assigned the task to perform. Furthermore, sensors and actors communicate directly among themselves in both models.

The *CoSense* project [22, 23] developed in the Palo-Alto Research Centre (PARC), has aimed at providing a collaborative sensing scheme for target recognition and condition monitoring. Moreover, the focus is on energy-constrained environments filled with low-observable targets. An energy-efficient sensor collaboration is presented in [22, 23]. This collaboration is information-driven, i.e., which dynamically determines who should sense, what to sense and to whom the sensed information must be passed. An assumption is made that each sensor has its communication and local sensing range. Furthermore, a sensor node is assumed to have local estimation capabilities of the cost of sensing, processing, and direct data communication to another node in terms of its power usage. The Information-Driven approach subsequently enables each sensor node to efficiently manage its communication and process resources. This entails that sensor selection is based on a local decision thus exhibiting a consensus-free approach. Also, a leader node holds the current belief and receives all passed information for fusion for a certain period of time. This leader node may then act as a relay station. Otherwise, the belief can travel through the network where the leadership is changed dynamically.

3 Collaborative Context Awareness (CCA) taxonomy

Our methodology in designing the CCA taxonomy is a result of inspiration from B. Randell's [24] and J.C. Laprie's work in defining Dependability. In addition, Roy Sterritt's method of defining Autonomic Computing [25, 26] was an additional influence. We designed our CCA taxonomy based on the commonalities of different

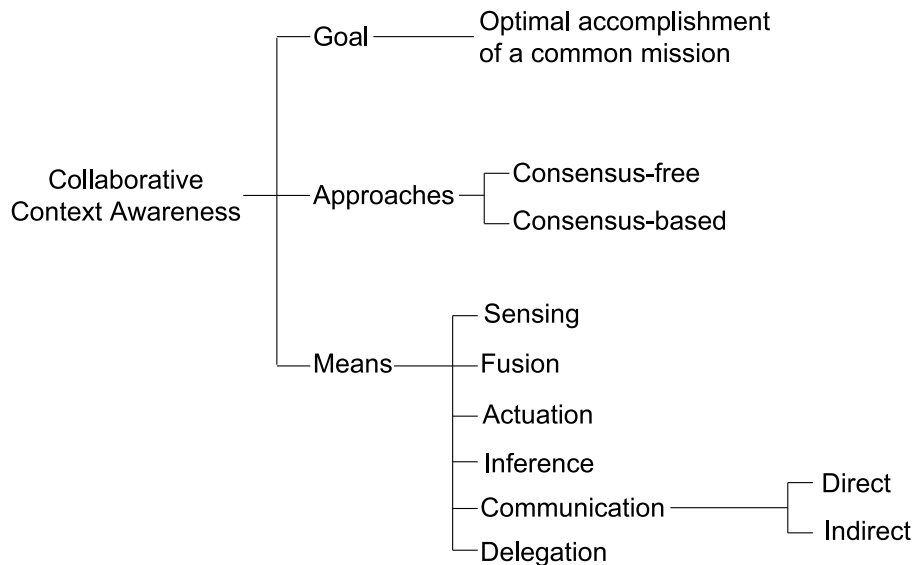


Fig. 1. CCA Tree

context-aware systems that emphasised collaboration. The structure of the CCA taxonomy illustrated in Figure 1 is based on three axis: *Goal*, *Approaches* and *Means*:

Goal : every collaborative (context-aware) system/sub-system aims at achieving a common *goal*, e.g., to optimally accomplish the assigned mission through some form of collaboration, cooperation and/or coordination among the comprising entities. We believe this must be achieved through information exchange and possibly some delegation techniques.

Approaches : depending on the application requirements, collaboration may follow a *consensus-free* and/or a *consensus-based* approach.

Consensus-free — entities may need to take local decisions hence they do not negotiate a common decision and communicate different information, i.e., sensory data, fused information, context/sub context and next action(s) in order to aid this local decision making.

Consensus-based — a system/sub-system may need negotiated outcomes, consequently entities are compelled to communicate different information in order to take a common decision.

Means : we identified seven means that may be used for collaborative context-aware systems to function.

Sensing — observing the environment typically entails the ability to receive different kinds of stimulus.

Fusion — the usual presence of numerous data sources in context-aware systems justifies the need to gather different low level pieces of data and information and the ability to build more reliable higher levels of knowledge.

Actuation — adjusting the system behaviour needs the realisation of inferred action(s), whether physically applied on the environment or not.

Inference — knowledge is a very crucial element in context-aware systems; hence the ability to build, update and reason about this knowledge is vital. In addition, deciding upon needed action(s) is important for context-aware systems' adaptability.

Communication — components/entities *must* exhibit a form of communication, i.e., indirect and/or direct, in order to realise their collaboration.

Direct — depending on the system architecture and application; entities may communicate using a dedicated channel for peer-to-peer, multicast or broadcast communication.

Indirect — stigmergic communication inspired from nature may help more efficient collaboration through the ability of communicating by changing and then sensing the shared environment.

Delegation — optimality is normally an important characteristic of context-aware systems, hence an entity capable of estimating for instance, computational and/or power needs, could decide that it is more efficient to delegate a task to a neighbouring entity if found more capable of handling it.

	Domain			Approaches		Means						
	Augmented Artefacts	Robotics	SAN	Consensus-free	Consensus-based	Fusion	Sensing	Actuation	Inference	Delegation	Communication	
											Direct	Indirect
Smart-Its	✓			✓	✓	✓	✓	✓	✓	✓	✓	
Smart-Its Friends	✓			✓	✓	✓	✓	✓	✓	✓	✓	
Cooperative Artefacts	✓			✓		✓	✓	✓	✓	✓	✓	
Anon. collaborative decision making [10]		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	
Millibots		✓		✓		✓	✓	✓	✓	✓	✓	
UAV collaborative sensing		✓		✓		✓	✓	✓	✓	✓	✓	
WSAN coordination framework			✓	✓	✓	✓	✓	✓	✓	✓	✓	
CoSense			✓	✓		✓	✓	✓	✓	✓	✓	

Table 1. Evaluation Results {SAN: *Sensor/Actuator Networks*}

The CCA taxonomy is flexible enough to encompass a large number of collaborative context-aware systems. For instance, a system of UAVs may be classified from the taxonomy as a consensus-free collaborative context-aware system that exhibits sensing (vision), actuation (manoeuvring), direct communication, fusion and inference. The *goal* of such a system could be drawing a map of a certain terrain with the best quality possible. In the next section, we classify the studied systems against the CCA taxonomy.

4 Evaluation

We provide in Table 1 the evaluation resulting from applying the CCA taxonomy to the studied systems. We believe that the relevant systems share the same *goal* of accomplishing an application-specific mission, hence we omitted the *goal* criterion from the evaluation table. We discuss each system and justify the relevant classification.

Smart-Its

Smart-Its exhibit a very generic design and application-specific behaviour. A system comprising Smart-Its could follow a consensus-free or a consensus-based approach or both (on different levels). A Smart-It explicitly provides dedicated sensing and actuation modules along with a processing module that provides inference and possibly fusion (depending on the implementation). The communication scheme among Smart-Its is direct at the moment. Indirect communication would also be possible if the application design benefits from stigmergy [17], i.e., Smart-Its understand/interpret

each others' physical actuation on the encompassing environment.

Smart-Its Friends

Smart-Its Friends are identical to Smart-Its but they exhibit a specific form of connection establishing.

Cooperative Artefacts

These artefacts cooperate by sharing a knowledge base through a query/response technique. They follow a consensus-free approach where they take decisions locally. The Cooperative Artefacts structure clearly exhibits sensing, actuation, fusion, and inference. Although, their current application as an alert system for storing hazardous chemical material containers shows limited actuation, i.e., switching LEDs on and off. As for fusion, it can be seen in the dedicated perception component that can produce location and proximity information for instance. A simple Prolog interpreter-like inference engine is also provided [15]. Communication is direct and through a short range wireless link.

Anonymous collaborative decision making

This scheme allows a consensus-free and a consensus-based approach through its coordinated and cooperated solutions respectively [10]. Sensing and fusion are exhibited in the Decentralised Data Fusion (DDF) technique. Furthermore, nodes/robotic sensors communicate directly and propagate their information in the whole network.

Millibots

Millibots directly communicate diverse sensory data depending on the type of sensors each Millibot is equipped with. They also do not negotiate decisions and hence follow a consensus-free approach. Each Millibot is responsible for local data fusion and inference depending on its view and the communicated information from other team members. Furthermore, Millibots typically exhibit actuation through mobility and the ability to command special type of adjustable sensors, i.e., a servo motor based Directional Infrared Detector Module (DIDM) [19].

UAV collaborative sensing

The communication scheme among the collaborating UAVs is believed to be direct despite the inspiration from swarming intelligence. This is because each UAV has a view based on its map of digital pheromones however, it wirelessly receives pheromone information from other UAVs in range and alters its map accordingly. Furthermore, a UAV is a local decision maker that is capable of sensing, e.g., imaging and actuating through manoeuvring. Inference in UAVs can be seen in a fitness evaluation procedure, i.e., quality of imaging [20].

WSAN coordination framework

The means of communication in both the sensor-actor and actor-actor models are direct. The framework exhibits sensing and actuation within the event-driven clusters. Furthermore, each sensor performs intermediate local data fusion and take local

decisions while forwarding event information to the designated actor in the cluster that has emerged. The actor gathers, processes and reconstructs event data while a consensus has to be reached afterwards among actors within the action/event area to select the best actor suitable to perform the action. The actor-actor consensus takes action completion time and/or energy consumption into account [21]. Interesting to note is that some form of delegation is present in the actor-actor coordination model.

CoSense

Sensors communicate directly and there is no means of actuation in the described research. Moreover, indirect communication is not possible since there is no actuation. The data fusion process centralised in a leader node. Also, each sensor node takes local routing decisions based on the cost of sensing, communication and processing estimations. Inference is implicit in the sensor selection process for specific target surveillance.

From the evaluation above we see that the CCA taxonomy succeeds to classify the diverse number of context-aware systems from different domains. We also see that certain systems provide for consensus-free and consensus-based approaches simultaneously or for a single one at a time. Most of the systems exhibit sensing, fusion, actuation and inference at different levels since these characteristics are normally intrinsic to context-aware systems. Communication is typically direct; this is justified by the possible difficulties in adopting a fully stigmergic communication paradigm. Also, a form of delegation is present in one system, namely, the WSA coordination framework; this could be a good motive to start adopting forms of delegation in other context-aware systems to improve overall system performance.

5 Conclusion

In this paper, we provided a taxonomy for Collaborative Context Awareness and investigated a number of context-aware systems and projects that focussed on collaboration. Based on the evaluation of these systems against the CCA taxonomy; we believe that the taxonomy is sufficiently generic to encompass a diverse number of collaborative context-aware systems from different domains. We also believe that the CCA taxonomy is a corner-stone in organising the concept of collaboration in context-aware systems by specifying a goal, two main approaches, i.e., consensus-free and consensus-based, and a set of concrete means.

We envisage delegation to be an important aspect in collaborative context-aware systems that seek optimality, e.g., traffic control, surveillance and UAV systems. In addition, we believe that the means for indirect communication should be provided besides the normal direct communication scheme. Finally, we do not encourage philosophical debates about terminologies such as collaboration, cooperation and coordination. We believe researchers should be more concise about their usage of such terms.

Our future efforts will focus on designing and implementing a CCA middleware that will support the development of collaborating intelligent context-aware entities for scenarios ranging from augmented artefacts to WSANs.

References

1. Biegel, G., Cahill, V.: A framework for developing mobile, context-aware applications. In: 2nd IEEE Conference on Pervasive Computing and Communications, Percom 2004, Orlando, Florida (2004)
2. Baldauf, M., Dustdar, S., Rosenberg, F.: A survey on context-aware systems. *International Journal of Ad Hoc and Ubiquitous Computing* (2006)
3. Chen, G., Kotz, D.: A survey of context-aware mobile computing research. Technical Report TR2000-381, Dept. of Computer Science, Dartmouth College (2000)
4. Dey, A.K., Abowd, G.D.: Towards a better understanding of context and context-awareness. GVU Technical Report GIT-GVU-99-22, Georgia Institute of Technology, Atlanta, GA, USA 30332-0280 (1999)
5. Pascoe, J.: Adding generic contextual capabilities to wearable computers. In: Second International Symposium on Wearable Computers, Pittsburgh, Pennsylvania, USA (1998) 92–99
6. Schilit, B., Adams, N., Want, R.: Context-aware computing applications. In: IEEE Workshop on Mobile Computing Systems and Applications, Santa Cruz, CA, US (1994)
7. Mäntyjärvi, J., Himberg, J., Huuskonen, P.: Collaborative context recognition for handheld devices. In: PERCOM '03: Proceedings of the First IEEE International Conference on Pervasive Computing and Communications, Washington, DC, USA, IEEE Computer Society (2003) 161
8. Siegemund, F.: A context-aware communication platform for smart objects. In: Pervasive Computing: Second International Conference, PERVASIVE 2004. Number 3001 in LNCS, Linz/Vienna, Austria, Springer-Verlag (2004) 69–86 Springer-Verlag.
9. Farinelli, A., Iocchi, L., Nardi, D.: Multirobot systems: a classification focused on coordination. In: Systems, Man and Cybernetics, Part B, IEEE Transactions on. Volume 34. (2004) 2015–2028
10. Grocholsky, B., Kumar, V., Durrant-Whyte, H.: Anonymous cooperation in robotic sensor networks. In: American Association for Artificial Intelligence, AAAI-04 Workshop on Sensor Networks, San Jose, California, USA (2004)
11. Beigl, M., Gellersen, H.: Smart-its: An embedded platform for smart objects. In: Proc. Smart Objects Conference (SOC 2003), Grenoble, France (2003)
12. : (<http://www.smart-its.org/>)
13. Holmquist, L.E., Mattern, F., Schiele, B., Alahuhta, P., Beigl, M., Gellersen, H.W.: Smart-its friends: A technique for users to easily establish connections between smart artefacts. In: Proc. Ubicomp 2001. Number 2201 in LNCS, Springer-Verlag (2001) 116–122
14. Gellersen, H.W., Schmidt, A., Beigl, M.: Multi-sensor context-awareness in mobile devices and smart artifacts. *ACM journal of Mobile Networks and Applications (MONET)* 7(5) (2002) 341–351
15. Strohbach, M., Gellersen, H.W., Kortuem, G., Kray, C.: Cooperative artefacts: Assessing real world situations with embedded technology. In: UbiComp 2004: Ubiquitous Computing: 6th International Conference, Proceedings, Nottingham, UK (2004) 250–267
16. Ricci, A., Viroli, M., Omicini, A.: Environment-based coordination through coordination artifacts. In: Environments for Multi-Agent Systems, First International Workshop, E4MAS 2004, New York, NY, USA, Springer (2004) 190–214
17. Dorigo, M., Bonabeau, E., Theraulaz, G.: Ant algorithms and stigmergy. *Future Gener. Comput. Syst.* 16(9) (2000) 851–871
18. Omicini, A., Zambonelli, F.: TuCSon: a coordination model for mobile information agents. In Schwartz, D.G., Divitini, M., Brasethvik, T., eds.: 1st International Workshop on Innovative Internet Information Systems (IIS'98), Pisa, Italy, IDI – NTNU, Trondheim (Norway) (1998) 177–187

19. Navarro-Serment, L., Grabowski, R., Paredis, C., Khosla, P.: Millibots. *IEEE Robotics and Automation Magazine* **9**(4) (2002) 31–40
20. Parunak, H., Brueckner, S., Odell, J.: Swarming coordination of multiple uav's for collaborative sensing. In: 2nd AIAA "Unmanned Unlimited" Systems Technologies and Operations Aerospace Land and Sea Conference, Workshop and Exhibition, San Diego, California, USA (2003)
21. Melodia, T., Pompili, D., Gungor, V.C., Akyildiz, I.F.: A distributed coordination framework for wireless sensor and actor networks. In: *MobiHoc '05: Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing*, New York, NY, USA, ACM Press (2005) 99–110
22. Zhao, F., Liu, J., Liu, J., Guibas, L., Reich, J.: Collaborative signal and information processing: An information directed approach. *Proceedings of the IEEE* **91**(8) (2003) 1199–1209
23. Zhao, F., Shin, J., Reich, J.: Information-driven dynamic sensor collaboration. *IEEE Signal Processing Magazine* **19**(2) (2002) 61–72
24. Randell, B.: Turing memorial lecture facing up to faults. *Computer Journal* **43**(2) (2000) 95–106
25. Sterritt, R., Bustard, D.W.: Autonomic computing - a means of achieving dependability? In: 10th IEEE International Conference on Engineering of Computer-Based Systems (ECBS 2003), Huntsville, AL, USA (2003) 247–251
26. Sterritt, R., Bustard, D.W.: Towards an autonomic computing environment. In: 14th International Workshop on Database and Expert Systems Applications (DEXA'03), Prague, Czech Republic, IEEE Computer Society (2003) 699–703