

A swarm of Mini-MEs: reasoning and information aggregation in ubiquitous multi-agent contexts

Floriano Scioscia, Michele Ruta, and Eugenio Di Sciascio

Politecnico di Bari, via E. Orabona 4, I-70125 Bari, Italy
E-mail: (floriano.scioscia, michele.ruta, eugenio.disciascio)@poliba.it

Abstract. Collaborative paradigms are essential in ubiquitous computing. *Swarm intelligence* can deal with complex problems deriving from the intrinsically unpredictable nature of mobile and pervasive scenarios. This paper introduces a novel reasoning framework for information fusion in multi-agent systems, encompassing a data dissemination protocol and non-standard, non-monotonic inference services. It enables agents to integrate detected and received information in a coherent view. Its properties include reconciliation of inconsistencies, quick adaptation to changes and robustness against spurious events. The paper presents the above generalized framework along with an experimental evaluation of a prototypical implementation supported by a swarm of pico-reasoners.

1 Introduction and motivation

The increasing effectiveness and computing capabilities of mobile and embedded devices allow to process and exchange rich and structured information in a novel range of applications. More and more effort is being spent to enhance flexibility and autonomy of ubiquitous and pervasive systems for information management, dissemination and discovery. Nevertheless, both studies and technological issues reveal that classical approaches based on centralized control and processing are troublesome in such scenarios [11]. So Multi-Agent Systems (MAS) are often targeted to provide the needed organizational support. This could be particularly true for the so-called *Semantic Web of Things* (SWoT) [12], a novel paradigm reflecting the convergence of Semantic Web and the Internet of Things. In an articulated effort aiming to tie metadata to ordinary objects and environments, MAS allow flexible interaction patterns and distributed information management, able to drive complex behaviors through *swarm intelligence* [6]. SWoT powers a plethora of motley micro-devices to convey small amounts of information each. High-level semantic languages cope with interoperability issues, also enabling machine-understandability. Non-monotonic reasoning is exploited to deal with the evolving and even partially conflicting knowledge. Unfortunately, no frameworks still exist for semantic-based information fusion integrating agent interactions suitably for ubiquitous contexts.

This paper introduces a novel semantic-based approach for information fusion in ubiquitous MAS. It exploits non-standard, non-monotonic reasoning allowing

agents to integrate captured and received knowledge in a coherent view. A new complex inference service has been defined for this purpose, capable of managing inconsistencies in information acquired by different agents. The framework leverages and extends the Mini-ME mobile matchmaker reasoning services [10]. A data dissemination protocol enables agents to receive, augment and propagate knowledge in a chain, increasing accuracy of situation awareness progressively. Reconciliation of conflicts, quick adaptation to changes and robustness against spurious events are among its properties. An experimental performance evaluation was executed on a network simulator to early assess the feasibility of the proposal.

The remainder of the paper is organized as follows. Next section discusses related work. Section 3 describes in detail the proposed framework, while Section 4 reports on the performance tests. Final remarks are in Section 5.

2 Related work

Multi-agent information fusion. Several approaches exist for distributed information aggregation in agent networks and MAS. When dealing with numerical data only, *aggregation functions* [3, 7, 4] provide a way to fuse measurements of individual agents and to reduce the size of exchanged data. In this family, *distributed particle filtering* [5] is one of the most popular and versatile techniques.

Semantic-based multi-agent intelligence relies on operators and inference services for information fusion on formulas. Peer-to-peer information integration in [2] was based on epistemic logic, allowing both a parallel and a sequential processing. Nevertheless, the approach focused on schema mappings in order to query a heterogeneous P2P system as a single entity, while the idea proposed here concerns distributed intelligence, with each peer having some relevant information to act upon and to exchange with the others. Such contexts require the capability to retract consequences when new conflicting knowledge becomes available. Therefore Subsumption and Satisfiability classical inference services are not enough: non-monotonic reasoning is needed. Belief merging and revision [7] have been studied for a long time, with many relevant theoretical results.

Reasoning services. In particular, in [8] the *Concept Abduction Problem* (CAP) has been defined for Description Logics (DLs) to provide an explanation when Subsumption does not hold. A CAP is presented as in what follows: given a DL language \mathcal{L} , an ontology \mathcal{T} in \mathcal{L} and two concepts C and D satisfiable in \mathcal{T} , if $\mathcal{T} \not\models C \sqsubseteq D$, then find a concept H (for *Hypothesis*) such that $\mathcal{T} \models C \sqcap H \sqsubseteq D$. That is, H is a possible explanation why resource features do not imply requested ones. On the other hand, if the conjunction $C \sqcap D$ is unsatisfiable w.r.t. the ontology, one can retract some requirements G (for *Give up*) from D to obtain a concept K (for *Keep*) such that $K \sqcap C$ is satisfiable in \mathcal{T} . This is called a *Concept Contraction Problem* (CCP) [8]. For both CAP and CCP, some minimality criteria must be defined [8], since one usually wants to hypothesize and give up as few things as possible. Abduction and Contraction were studied and applied to semantic matchmaking, which is basically the process of finding best matches

of a *request* (named D in the above definitions) among available *resource ads* (named C above), where both the request and the ads are annotated w.r.t. a reference ontology [8]. In such scenarios, under the Open World Assumption (OWA: *what is not specified has not to be interpreted as a constraint of absence*), the result is a list of ads ranked by semantic proximity. Interestingly, ads could contain features which the user did not know or consider when composing her request, but which could be added in a query refinement process to improve the outcome. For this purpose, the **Bonus** inference service was defined [8] to extract a concept B from D which denotes something that the ad offers even though the request did not ask for it. Finally, the proposed information fusion framework requires a way to subtract information in a description from another one. This is accomplished by means of the Concept **Difference** reasoning service [14]. Also for Difference many valid solutions exist, but a maximality criterion should be adopted in this case (*i.e.*, subtract as much as possible).

3 Reasoning for multi-agent system intelligence

The proposed framework refers to a swarm of independent agents exchanging annotated *envelopes*. Each envelope contains the sender’s current knowledge of the context, expressed in a DL language w.r.t. a shared reference ontology. Basically, envelopes are composed of a timestamp and a 4-tuple of annotations, with meaning as in Table 1. The juxtaposition of the four elements allows information aggregation with reconciliation of inconsistencies through a collaborative protocol for information dissemination, exploiting a properly devised inference service. The protocol assigns a *Time To Live* (TTL) to all envelopes. An envelope *expires* and an agent discards it when the sum of its timestamp and TTL exceeds the agent’s current time.

When an agent completes a data gathering and annotation round, it has available a semantic expression N representing newly detected information. Furthermore, a cache stores the most recent incoming envelope and, when a new envelope is received, the former one is discarded. In the protocol, three basic cases are distinguished:

1. *Generation*: it occurs if an agent generates an annotation and it has no *fresh* (*i.e.*, non-expired) envelope in its cache, that is when either it has received no envelopes or all received envelopes have expired.
2. *Relay*: it occurs if an agent has an envelope in its cache and it is not able to

Table 1: Meaning of the annotations in exchanged envelopes

Name	Description
C (Confirmed)	elements observed by both the sender and other agents
X (Clash)	elements observed by the sender, inconsistent with observations by others
M (My)	elements observed by the sender, but not by other agents
E (External)	elements observed by other agents, but not by the sender

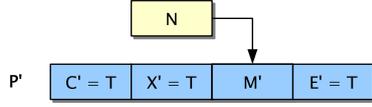


Fig. 1: New envelope generation

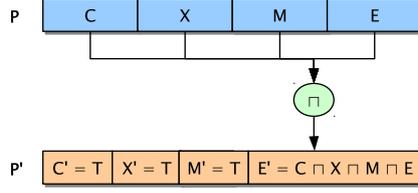


Fig. 2: Relay of an envelope

produce a semantic annotation of its own.

3. *Integration and relay*: it occurs when an agent has a received envelope in its cache, which must be fused with self-detected information before transmitting the updated snapshot of the environment.

In any case, the result is a new envelope P' , which will be sent in broadcast.

In the above case 1, the resulting P' is sketched in Figure 1. Since everything is the outcome of observations of the sending agent, the annotation N will be placed in the envelope field M , while other fields will be set to \top (Top or Thing) to denote the most generic possible observation. Conversely, in case 2 the agent has to relay a received envelope P , therefore the result P' will contain only observations made by other agent(s). The outgoing envelope is structured as in Figure 2: only the E field will be populated with the conjunction of the four annotations in P ; the remaining ones will be set to \top .

The last case requires the fusion of generated and received information. For this purpose, a novel inference service was devised, so as to preserve the semantics of the fields of both input and output envelopes. The logic problem to solve is named *Concept Integration* and is formalized as:

Definition 1. Let \mathcal{L} be a DL language, \mathcal{T} a set of axioms in \mathcal{L} , N a concept expression in \mathcal{L} and P a 4-tuple of concept expressions $\langle C, X, M, E \rangle$ in \mathcal{L} , where N is the current knowledge of agent A' , P is knowledge accumulated by agent A —whose components have the semantics in Table 1— and N, C, X, M, E are satisfiable in \mathcal{T} . A *Concept Integration Problem (CIP)*, identified by $\langle \mathcal{L}, N, P, \mathcal{T} \rangle$, is finding a 4-tuple of concept expressions $P' = \langle C', X', M', E' \rangle$, which is the integration of N and P , having the semantics in Table 1. Then P' is the integration of N and P and represents knowledge accumulated by agent A' .

Algorithm 1 solves a CIP exploiting the Concept Contraction, Abduction, Difference and Bonus inference services. At line 3, $J_1 = C \sqcap X \sqcap M \sqcap E$ is the “resource” and N the “request”, therefore G_N is the part of the observation of agent A' which contradicts previous observations accumulated by A , while K_N contains the part that is in accordance. Analogously for the CAP at line 4, result H_N is something only A' has detected and previous observations had not reported. To determine what is shared by A and A' , it is needed to subtract H_N from K_N , computing the Concept Difference at line 5. Now G_N and H_N clearly have the required meaning of elements X' and M' of the new envelope, respectively (lines 6-7). Finally, in order to compute what other agents had seen but A' has not, the Bonus service is adopted (line 8): it is important to note that in this case the role of “resource” is played by $J_2 = C \sqcap X \sqcap M$ only. Including E in J_2 would imply a perpetual propagation of concepts in E when

Algorithm 1 Concept Integration Problem solution

Algorithm: $Integrate(\langle \mathcal{L}, N, P, \mathcal{T} \rangle)$
Require: $\langle \mathcal{L}, N, P = \langle C, X, M, E \rangle, \mathcal{T} \rangle$ with and $N, C, X, M, E \in \mathcal{L}$ satisfiable in \mathcal{T}
Ensure: $P' = \langle C', X', M', E' \rangle$ with $C', X', M', E' \in \mathcal{L}$ satisfiable in \mathcal{T}

- 1: $J_1 := C \sqcap X \sqcap M \sqcap E$
 - 2: $J_2 := C \sqcap X \sqcap M$
 - 3: $\langle G_N, K_N \rangle := contract(N, J_1, \mathcal{T})$ // [8]
 - 4: $H_N := abduce(K_N, J_1, \mathcal{T})$ // [8]
 - 5: $C' := difference(K_N, H_N, \mathcal{T})$ // [14]
 - 6: $X' := G_N$
 - 7: $M' := H_N$
 - 8: $E' := computeBonus(K_N, J_2, \mathcal{T})$ // [8]
 - 9: **return** $P' = \langle C', X', M', E' \rangle$
-

applying CIPs in cascade, which is not compliant with the purpose of the inference. CIP allows to integrate an “older” snapshot with “new” information to produce an “updated” snapshot, which reflects the viewpoint of the agent that has performed the integration.

It is important to highlight that $\forall P = \langle C, X, M, E \rangle, C \sqcap X \sqcap M \sqcap E$ is consistent. Indeed, C and X are two parts of the produced annotation N , which is always consistent by assumption. M is derived from C and E is the result of Bonus, which is always consistent with its input.

The CIP-based protocol has a further interesting property: it reaches a steady state in just two observation rounds after every variation in information. Formally:

Theorem 1. *Given a reference TBox \mathcal{T} , let A_1, A_2, \dots, A_n be n agents which receive, integrate and relay envelopes P_1, P_2, \dots, P_n in cascade, with $A_i \neq A_{i+1} \forall i$. Let N_1, N_2, \dots, N_n be the concept expressions –all satisfiable in \mathcal{T} – of the information they detect respectively. Then $\forall i = 1, \dots, n-2$ and $\forall S$ satisfiable concept in \mathcal{T} : $(C_i \sqsubseteq S) \wedge (N_{i+1} \sqsubseteq \neg S) \wedge (N_{i+2} \sqsubseteq \neg S) \Rightarrow C_{i+2} \sqsubseteq \neg S$.*

Proof. If $C_i \sqsubseteq S \wedge N_{i+1} \sqsubseteq \neg S$, then C_i and N_{i+1} are inconsistent and lines 3 and 6 of Algorithm 1 ensure that $X_{i+1} \sqsubseteq \neg S$. In the subsequent CIP between N_{i+2} and P_{i+1} , $\neg S$ will not be a source of inconsistency and $K_{N,i+2} \sqsubseteq \neg S$ (line 3). Since $\neg S$ is in both N_{i+2} and $K_{N,i+2}$, it will not be part of $M_{N,i+2}$ and line 5 of CIP algorithm ensures that $C_{i+2} \sqsubseteq \neg S$.

This means a MAS adopting the proposed protocol can (i) follow rapid changes in environmental conditions closely and (ii) recover from a detection mistake quickly. Furthermore, the approach is grounded on the OWA, which makes it well suited for incomplete information, while Theorem 1 allows it to cope with noisy data. Finally, Algorithm 1 clearly shows that CIP complexity depends directly on the one of Abduction, Contraction, Bonus and Difference inference services. In particular, the experimental evaluations with Mini-ME worked on the \mathcal{ALN} (Attributive Language with unqualified Number restrictions) DL with *simple-TBoxes* [9], for which structural algorithms exist with PTIME complexity for all these services, so CIP has PTIME complexity too.

4 Experiments and performance evaluation

A prototypical implementation of the proposed framework was developed for swarm intelligence scenarios with the *NTCUns* network simulator [15]. The testbed simulates a VANET (Vehicular Ad-hoc NETwork) for cooperative monitoring of driving risk factors, where each agent runs on a smartphone in a car and gathers data from phone micro-devices as well as from the OBD-II (On-Board Diagnostics) port integrated in the vehicle. Raw data are analyzed and translated in concept expressions referred to an ontology modeling road features, traffic, weather and driving style [13].

The early performance test reported here¹ aimed to evaluate the sustainability of the proposed approach in pervasive contexts. In the simulations, agents move according to the Manhattan mobility model [1] and exchange envelopes through wireless ad-hoc connections. Reasoning is supported by the Mini-ME matchmaker [10], extended with Concept Difference and Integration inference services. Knowledge capturing is simulated via pre-defined annotations pertaining to different zones of a map. The reference \mathcal{ALN} ontology includes 99 classes, 18 object properties and 272 axioms. Each annotation contains observations of either 5 or 6 properties, each expressed through an existential-universal restriction pair; atomic concepts are the fillers of universal restrictions. Two scenarios were simulated, differing in the number and density of nodes, as well as in environmental conditions. Tests were executed on an Asus UX31A notebook with Intel Core i7 3517U processor (dual core, 1.9-3.0 GHz), 4 GB DDR3 1600 MHz SDRAM, 256 GB SATA III SSD disk and GNU/Linux Fedora 16 operating system. Each scenario was run 3 times and average values were taken. Table 2 summarizes simulation parameters, while Table 3 reports on results. The greater node count in Scenario 2 leads to a higher number of envelope exchanges and *Integration* resolutions. Average times are below 10 msec: even taking into account an order of magnitude performance penalty on mobile devices or single-board computers w.r.t. the testbed hardware [10], the CIP inference service appears efficient enough for swarm intelligence in pervasive computing scenarios.

Table 2: Simulation parameters

Scenario	1	2
No. of agents	60	80
Sense & process period (s)	2	2
Max agent speed (m/s)	36	50
Map size (km)	5×2	10×2
No. of map zones	9	10
Duration (s)	300	300

Table 3: Test results

Scenario	1	2
CIP resolutions	9330	16444
Avg. CIP time (ms)	9.8	8.1

¹ See <http://sisinflab.poliba.it/swottools/minime/download/ore2015exp.zip>

5 Conclusion

A novel extensive semantic-based framework has been proposed for information fusion in pervasive multi-agent contexts, based on non-standard, non-monotonic inference services. It allows to reconcile inconsistent observations and makes the MAS able to reach quickly a steady state with a coherent situation awareness. Performance evaluation with the optimized Mini-ME mobile matchmaker suggests adequate efficiency. Future work includes integration in an off-the-shelf message-oriented middleware for swarm intelligence and in VANET scenarios. Extensive experiments to evaluate the effectiveness of the proposed approach w.r.t. the state of the art will concern: reasoning performance on very resource-constrained devices (*e.g.*, sensor motes); scalability w.r.t. the size and complexity of managed ontology and expressions, the number of agents and monitored parameters; communication performance of the devised protocol; quality of the disseminated information with respect to application goals in realistic scenarios. Finally, theoretical investigation is underway to extend the approach to multi-item fusion.

Acknowledgments

The work was supported by the Italian PON project PLATINO (PLATform for INnOvative services in future internet) and the ETCP project ARGES (pAssen-gerS and loGistics information Exchange System).

References

1. Bai, F., Sadagopan, N., Helmy, A.: IMPORTANT: A framework to systematically analyze the Impact of Mobility on Performance of Routing protocols for Adhoc Networks. In: INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies. vol. 2, pp. 825–835. IEEE (2003)
2. Calvanese, D., De Giacomo, G., Lenzerini, M., Rosati, R.: Logical foundations of peer-to-peer data integration. In: Proceedings of the twenty-third ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems. pp. 241–251. ACM (2004)
3. Calvo, T., Beliakov, G.: Aggregation functions based on penalties. *Fuzzy sets and Systems* 161(10), 1420–1436 (2010)
4. Castanedo, F., García, J., Patricio, M.A., Molina, J.M.: Data fusion to improve trajectory tracking in a Cooperative Surveillance Multi-Agent Architecture. *Information Fusion* 11(3), 243–255 (2010)
5. Hlinka, O., Hlawatsch, F., Djuric, P.M.: Distributed particle filtering in agent networks: A survey, classification, and comparison. *Signal Processing Magazine, IEEE* 30(1), 61–81 (2013)
6. Kennedy, J., Kennedy, J.F., Eberhart, R.C.: *Swarm intelligence*. Morgan Kaufmann (2001)
7. Konieczny, S., Pérez, R.P.: Logic based merging. *Journal of Philosophical Logic* 40(2), 239–270 (2011)

8. Ruta, M., Di Sciascio, E., Scioscia, F.: Concept abduction and contraction in semantic-based P2P environments. *Web Intelligence and Agent Systems* 9(3), 179–207 (2011)
9. Ruta, M., Scioscia, F., Di Sciascio, E., Gramegna, F., Loseto, G.: Mini-me: the mini matchmaking engine. In: *OWL Reasoner Evaluation Workshop (ORE 2012)*. vol. 858, pp. 52–63 (2012)
10. Ruta, M., Scioscia, F., Loseto, G., Gramegna, F., Ieva, S., Di Sciascio, E.: Mini-ME 2.0: powering the Semantic Web of Things. In: *3rd OWL Reasoner Evaluation Workshop (ORE 2014)*. *CEUR Workshop Proceedings*, vol. 1207, pp. 8–15 (2014)
11. Ruta, M., Sciascio, E.D., Piscitelli, G., Scioscia, F.: A ubiquitous knowledge-based system to enable RFID object discovery in smart environments. *Pacific Asia Journal of the Association for Information Systems* 2(3) (2010)
12. Ruta, M., Scioscia, F., Di Sciascio, E.: Enabling the semantic web of things: framework and architecture. In: *2012 IEEE Sixth International Conference on Semantic Computing*. pp. 345–347. IEEE (2012)
13. Ruta, M., Scioscia, F., Gramegna, F., Di Sciascio, E.: A mobile knowledge-based system for on-board diagnostics and car driving assistance. In: *UBICOMM 2010, The Fourth International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies*. pp. 91–96. IARIA (2010)
14. Teege, G.: Making the Difference: A Subtraction Operation for Description Logics. *Proceedings of the Fourth International Conference on the Principles of Knowledge Representation and Reasoning (KR94)* pp. 540–550 (1994)
15. Wang, S., Chou, C.: NCTUns tool for wireless vehicular communication network researches. *Simulation Modelling Practice and Theory* 17(7), 1211–1226 (2009)