# A Semantic-based Approach for Resource Discovery and Allocation in Distributed Middleware

Michele Ruta
DEI - Politecnico di Bari
Via E. Orabona 4
Bari, Italy I-70125
michele.ruta@poliba.it

Floriano Scioscia
DEI - Politecnico di Bari
Via E. Orabona 4
Bari, Italy I-70125
floriano.scioscia@poliba.it

Eliana Bove
DEI - Politecnico di Bari
Via E. Orabona 4
Bari, Italy I-70125
eliana.bove@poliba.it

Annarita Cinquepalmi
DEI - Politecnico di Bari
Via E. Orabona 4
Bari, Italy I-70125
annarita.cinquepalmi@poliba.it

Eugenio Di Sciascio
DEI - Politecnico di Bari
Via E. Orabona 4
Bari, Italy I-70125
eugenio.disciascio@poliba.it

## ABSTRACT

This paper presents a knowledge-based approach for resource discovery, allotment and sharing in distributed pervasive scenarios. The proposed framework enables semantic-based resource retrieval exploiting non-standard inference services and a novel method for ontology dissemination and on-the-fly reconstruction. The approach can augment any publish/subscribe message-oriented middleware. A prototype was implemented and tested to prove correctness of the approach and get early performance evaluations.

## Keywords

Mobile Resource Discovery; Internet of Things; Semantic-based Matchmaking; Message-oriented Middleware; Swarm Intelligence

## 1. INTRODUCTION

The *Internet of Things* (IoT) vision is a world where information could be really scattered in a given environment in the form of atoms which deeply permeate the context [1]. Most current middleware platforms are topic-based [2], relying on a set of predefined subjects and on trivial string matching on topics. Those static configurations reveal a not-negligible limit in the exploitation of service-oriented approaches in fully autonomic application scenarios. Semantic-enabled service and request specification allows more accurate and flexible characterization of requesters' needs and providers' capabilities. Furthermore, mobile and pervasive computing contexts require decentralized paradigms and dynamic lightweight middleware. Authors in [3] took into account only full matches, which are quite rare in complex domains with detailed descriptions. Furthermore, the problem of distributed ontology management cannot be overlooked. Typical ontologies for service annotation are too large and complex to be processed by any single node in pervasive computing contexts, therefore strategies for modularizing terminologies are necessary. This paper presents a knowledge-based approach for resource discovery, allocation and sharing in distributed scenarios, augmenting the publish/subscribe (pub/sub) Message-Oriented Middleware (MOM) architectural model. The proposed framework interconnects distributed components along three different levels ranging from the external interface to the low level communication primitives: semantic-enhanced resource discovery; ontology partitioning and on-the-fly reconstruction; pub/sub message exchange.

The above theoretical architecture has been implemented in a working prototype which has been used for an early experimental campaign devoted to ensure correctness of the approach and to perform a preliminary performance evaluation. Next Section 2 discusses the proposed framework, while Section 3 reports on the ongoing experimental evaluation.


DOI: 10.1145/1235

## 2. INJECTING SEMANTICS IN DISTRIBUTED MIDDLEWARE

The proposed framework enhances the application level of classical pub/sub MOM infrastructures exploited for interconnecting large numbers of loosely coupled components on independent subsystems and devices. Interactions are based on the exchange of *messages* labeled with a *topic* string,Each node can act as a *publisher* to emit messages with a particular topic and/or as a *subscriber* to receive all messages with the subscribed topic. In standard pub/sub MOM architectures, resource discovery occurs through syntactic match of topics. The proposal enhances off-the-shelf middleware with support for dynamic semantic-based service/resource discovery. The proposed architecture comprises three layers:

**1. Data Distribution Service**. It is an off-the-shelf pub/sub MOM which provides services for real-time data distribution by including Data Local Reconstruction Layer (DLRL) and Data Centric Publish/Subscribe (DCPS).

**2. Ubiquitous Knowledge Base** (u-KB).A KB is a pair $K = <\mathcal{T}, \mathcal{A}>$. $\mathcal{T}$ is the terminological box (TBox or ontology), *i.e.*, a formal representation of the conceptual model of a domain. $\mathcal{A}$ is the assertion box (ABox), specifying the factual information of a particular scenario within the domain as a set of *individuals*. In classical knowledge representation approaches KBs are centralized and monolithic, stored in a given server location reachable by several clients interested in operations on it. This kind of approach is not realistic for IoT contexts featured by volatile nodes interacting in an opportunistic fashion. The implemented approach finds its fulfillment in the u-KB paradigm by allowing KB partitioning in a decentralized way, scattering on multiple nodes. Multiple u-KBs can co-exist within the same physical and logical environment, in order to model different resource domains. The proposal introduces a novel method for ontology partitioning based on associating each class with a unique ID, computed from its position in the taxonomy. The most generic class, always named *Thing* in OWL 2 (*a.k.a. Top* or $\top$ in DL notation), takes ID 1. Each nesting level adds a further numerical suffix, separated by a . (dot). As an optimization, the $N$ topmost levels in the class hierarchy (which specify the most general concepts of a domain), can be disseminated among all nodes as an *Upper Ontology* (UO) chunk. Basically, any node in a network can manage any domain ontology, even using multiple ontologies in order to cover several application domains. Furthermore, the use of unique ontology URIs (Uniform Resource Identifiers) ensures that all objects working with the same reference ontology can share parts of the u-KB dynamically without requiring preliminary agreement among them. When a node performs resource discovery, as a preliminary task it must rebuild a subset of the ontology in order to guarantee correctness of reasoning procedures. To do so, it publishes a message with the *BuildTBox* middleware topic, which all semantic-enabled nodes must be subscribed to. The message contains: (i) the unique ontology URI, (ii) the list of requested class IDs, and (iii) the topic name (*e.g.*, *MergeOnto*_NodeID) to be used in reply messages. If a node has one or more requested class IDs in its cache, it will register a publisher on the above topic and send the compressed ontology chunk(s) containing those classes. Requester node is subscribed to topic *MergeOnto_NodeID* to receive the ontology chunks and merge them.

**3. Resource/Service Discovery**. A semantic-based resource request consists of a logic-based annotation expressed w.r.t. a reference ontology. The discovery phase exploits a general topic named *Discovery*, which all nodes in the network are subscribed to. The requester starts inquiry by sending a *Discovery* message containing: (i) the URI of the ontology the request refers to and (ii) the topic name (by design, *SemAnn*_NodeID) to be used in reply messages; also in this case, the reply topic is node-specific. Through the *Discovery* topic, other nodes receive the request and check whether they own services/resources related to the same domain. Only in this case, nodes become publishers on the reply topic and send back the related compressed annotations; each annotation is associated with a service-specific topic. The requester collects all descriptions and compares them with its request through the semantic matchmaking process described in [4], with support for approximate matches, resource ranking and formal explanation of outcomes (details not recalled here due to space constraints). The outcome of the match is a ranked list of resources which best satisfy the request. Finally, the requester uses the topic(s) associated to the selected resource(s) in order to start fruition.

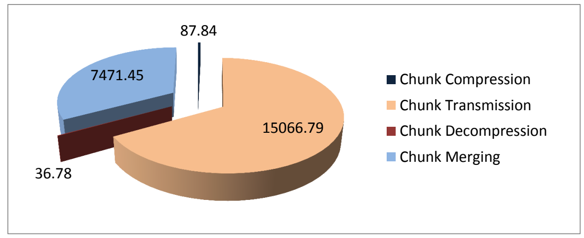


**Figure 1: Time for ontology rebuilding**

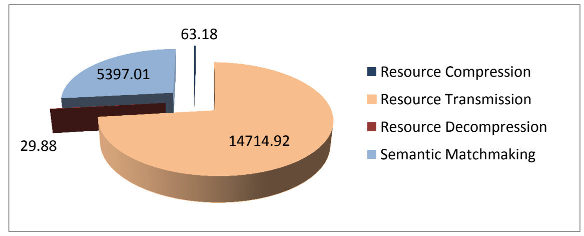


**Figure 2: Time for resource allotment**

# 3. RESULTS AND CONCLUSION

The proposed framework was implemented in Bee Data Distribution System (Bee-DDS, http://sine.ni.com/nips/cds/view/p/lang/en/nid/211025), an off-the-shelf middleware, in order to evaluate both feasibility and performance. Nevertheless, the approach is completely general and can be implemented on top of any pub/sub MOM.

Performance was evaluated simulating a workload of 500 resource requesters and 10 providers, connected through a 100 Mb/s IEEE 802.3 network. The execution time for the ontology rebuilding phase was about 22.67 s on average (Figure 1), while the resource allotment stage took an average time of 20.2 s (Figure 2). For both phases, message transmission was the longest sub-task. RAM occupancy for provider nodes was lower than for requesters (data not shown) since the latter are equipped with the reasoner to execute service matchmaking. Overall, the above tests validated the correctness of the proposal. Nevertheless, performance optimizations and scalability evaluations are needed.

# 4. REFERENCES


[1] L. Da Xu, W. He, and S. Li. Internet of things in industries: a survey. *Industrial Informatics, IEEE Transactions on*, 10(4):2233–2243, 2014.

[2] G. Fortino, A. Guerrieri, M. Lacopo, M. Lucia, and W. Russo. An agent-based middleware for cooperating smart objects. In *Highlights on Practical Applications of Agents and Multi-Agent Systems*, pages 387–398. Springer, 2013.

[3] S. B. Mokhtar, D. Preuveneers, N. Georgantas, V. Issarny, and Y. Berbers. EASY: Efficient semAntic Service discoverY in pervasive computing environments with QoS and context support. *Journal of Systems and Software*, 81(5):785–808, 2008.

[4] M. Ruta, E. Di Sciascio, and F. Scioscia. Concept abduction and contraction in semantic-based P2P environments. *Web Intelligence and Agent Systems*, 9(3):179–207, 2011.