

Architecture Design Principles to Support Adaptive Service Orchestration in WSN Applications

Kirill Mechitov
University of Illinois at
Urbana-Champaign
Urbana, IL, USA
mechitov@cs.uiuc.edu

Reza Razavi
University of Luxembourg
FSTC, Luxembourg
razavi@acm.org

Gul Agha
University of Illinois at
Urbana-Champaign
Urbana, IL, USA
agha@cs.uiuc.edu

Abstract

Our goal is to facilitate the development of sensor network applications in an open system, where applications arrive and leave dynamically and execute concurrently. We identify design principles that govern the creation of these systems, such as having a network-wide programming model, late binding and global resource management. In accordance with these principles, we assume that an application is modeled as a composite service, and propose an architecture for its adaptive orchestration on a WSN. Adaptivity here refers to automatic runtime selection of service implementations and network resources to execute the application specification in a resource-efficient and context-aware manner.

1 Introduction

Current practice considers wireless sensor networks (WSNs) in the context of a single application, *e.g.*, a WSN for target tracking or a WSN for environment monitoring. This model of application development, together with the small scale of most experimental sensor network deployments, has led to the design of middleware services that are highly efficient but often tightly coupled or customized to a particular application. This practice hampers service portability and reusability, such as when a data aggregation service is designed to work only with a specific routing protocol.

Some recent work has proposed supporting several concurrent applications on a sensor network [13]. As WSN deployments become more numerous and their scale increases, we envision sensor networks becoming a *computing platform* used concurrently by multiple users and multiple applications for different and uncoordinated activities. For instance, as illustrated in Fig. 1, middleware services should be shared among unrelated applications. In this context, efficient customized middleware services specific to each application are a poor solution, as common functionality is needlessly replicated.

In our view, these requirements imply the need for a software architecture that provides a looser coupling between services and applications, and among the services themselves, in a resource-efficient and context-aware manner. We consider applications that make use of a number of general middleware services, such as routing, localization, and time




Figure 1. Middleware services are shared among concurrently executing applications, resulting in a many-to-many relationship between applications and services.

synchronization. In order to accommodate the vast collection of services and protocols already developed by the sensor network community, we adopt a very broad definition of a middleware service, concerning ourselves only with their interfaces to applications and other services and not their internal semantics or implementation method. Specifically, we propose a dynamic service composition-based architecture for WSN applications, based on the principle of *self-mediated execution*, with the dual goals of facilitating large-scale application development and enabling global, network-wide optimization, rather than application-focused local optimization of the constituent middleware services. Our approach is based on postponing the binding of applications to specific network resources and implementations of middleware services from design- or compile-time to the runtime. Appropriate service implementations are chosen at runtime and deployed on demand.

By dissociating middleware services from the application context and from each other, we give up some possible performance advantages due to explicit customization and tight coupling. In return, we provide a more scalable software development process, support for multiple concurrent applications, and the possibility of global resource management across applications.

The remainder of the paper is organized as follows. Section 2 states our overall design principles. Sections 3 and 4

then describe a dynamic service composition-based architecture implementing these principles and Section 5 illustrates its use. Section 6 discusses the properties of our architecture, and Section 7 relates it to other work in this area.

2 Design Principles

Service- and component-based architectures are widely used, providing greater ease and scalability to the software design and implementation process. We aim to apply the same approach to the sensor network domain, adapting to its unique limitations and requirements. We identify the following key principles for the design of scalable, resource-efficient WSN applications as a composition of middleware services:

1. *Network-wide programming model.* The sensor network is treated as a collaborative distributed computing platform. Applications are specified as a collection of network-wide tasks and not as a unique program image per sensor node.
2. *Sharing and reuse.* Multiple uncoordinated applications and middleware services need to coexist in the network without prior knowledge of each other. Therefore, both network resources requiring exclusive access (sensors, actuators, etc.) and middleware services are shared among several applications. Resource management cannot be relegated to each application individually, it must be performed globally.
3. *Late binding.* Application specification is sufficiently flexible to allow run-time adaptivity in selecting the services and resources to be used. We do not know in advance which services or resources will be used by which application, or when. Postponing the choice of which service or resource best fits the application opens up more opportunities for optimization.

In the following section we present a service composition-based software architecture that follows from these design principles.

3 Architecture Overview

Our architecture leverages the concept of dynamic service composition to support application development for open WSN systems. We adopt a two-level architecture, separating the two major concerns: that of controlling the execution process, including strategic decision making and adaptation, and that of the execution itself. First, we restate our assumptions about the problem more formally.

3.1 Assumptions

We consider applications specified in terms of a composition of calls to middleware service interfaces, and we refer to the service interface specification as a *contract* and each call to a service a *service request*. A repository of available services for a given WSN or application domain is provided.

To facilitate the use of a large number of pre-existing middleware services within our architecture, we choose not to constrain the model of a service's behavior, *e.g.*, whether it is distributed, centralized, single-threaded, *etc.* Since services and applications need to interact and coordinate, however,

we fix a model for their interaction. We use the Actor model of computation [1] to represent service interfaces connecting services to each other and to the application. Thus, services are used by our system as if they were implemented as *actors*: concurrent active objects interacting via asynchronous messaging. We distinguish between the actors representing the service itself from *meta-actors*, which are control threads supervising deployment and execution of the services.

Responsibilities of the meta-actor include controlling the lifecycle of a service (deploying, starting, stopping and disposing of the service) and interaction with other services. Note that once the appropriate services are deployed, they may choose to interact directly, rather than through their corresponding meta-actors. Interaction then occurs through the actor interface specified in the service contract. Only interactions through actor interfaces are mediated by our architecture; any side effects are not captured by this model.

We further assume the existence of a functional service composition language, where service requests are *self-sufficient* and *minimally constrained*. The service composition language is functional in that (1) the control flow between service requests is partially ordered and driven by data dependencies, and (2) it allows for a recursive graph traversal to autonomously process each service request in the specification. Self-sufficiency refers to the fact that each individual service request is provided with the required knowledge about the arguments, resources, context and method required for its execution. Minimally constrained refers to delaying as long as possible placing constraints necessary to execute a specific instance of the service, in other words, the service instance does not refer to information that can be computed or supplied to it at run-time.

The last requirement is a fine-grained runtime code deployment method, such as a mobile agent system like ActorNet [7] or Agilla [5].

3.2 Example

Consider how a typical localization service request is represented in our architecture. To be self-sufficient, the contract includes a reference its execution method, *e.g.*, a compiled library implementation of the localization algorithm, the type of sensors used, such as distance measuring or angle of arrival, and data types for the output (locations and error intervals). To be minimally constrained, it must not specify a deployment location (node ids) or method (a specific range measurement service), referring instead to the contracts in the repository. Execution-specific information is filled in at run-time based on the specified constraints.

4 Architecture Components

Given an application comprising a composition of middleware service requests represented in such manner, its execution consists of a self-decomposition and self-deployment process. This results in a system of distributed interacting meta-actors responsible for handling the interaction among the services. Execution proceeds concurrently and asynchronously as the preconditions for the deployment of each service request are satisfied. We call this process *self-mediated execution*.

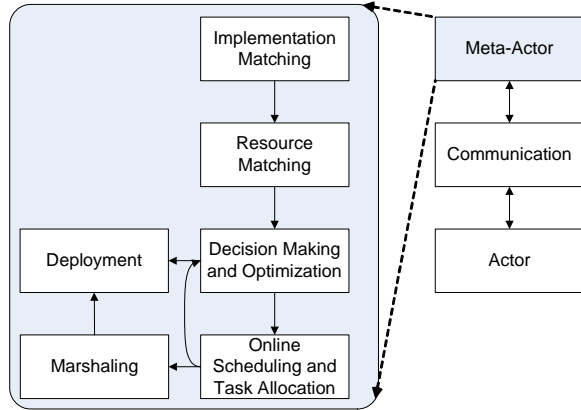


Figure 2. Self-mediated execution architecture for middleware services.

Let us now focus on the role of the meta-actors in this process. Fig. 4 highlights the governing behavior of a meta-actor in processing service requests. Due to service request self-sufficiency, each meta-actor can decide *how*, *where* and *when* to execute its associated service. We now explain the function of each component of this architecture and their interactions.

4.1 Choice of Implementation

Deciding *how* to execute the service request involves matching a particular service implementation to an interface from the service contract repository, and then finding the network resources required by that implementation.

Implementation Matching.

This component finds all implementations that match the constraints of a given service request. For example, we might search for implementations of a ranging service with a `MeasureDistance` method that also satisfy a maximum distance constraint. This is done by querying the contract repository and filtering the results according to the constraints specified in the service request. Pattern matching or a linear constraint solver may be used to filter the available service implementations.

Resource Matching.

Likewise, the resource matching component finds all suitable resources for a given service implementation. Matching algorithms used by this service depend on the resource description language employed by the system. Several methods are available for indexing a dynamic set of geographically distributed resources, including a yellow pages service, tuple spaces and actor spaces. For instance, if tuple spaces are used, sensor nodes entering the system can publish their resource descriptions in the tuple space, and the resource matching component performs a search in the form of pattern matching [2]. Caching and prefetching techniques can make the process more efficient, eliminating the need to scour the network for each query. Due to the location-dependent nature of most WSN computations, we expect most queries to be limited geographically, avoiding the need to flood the network even in cases when cached information is unavailable.

4.2 Location and Deployment

Second, the meta-actor needs to decide *where* to execute the service request. For the sake of efficiency, deployment and invocation are treated separately. As such, code deployment starts as soon as possible, while the invocation is delayed by the scheduling component until the necessary resources become available.

Decision Making and Optimization.

Given a list of possible resources and implementations, this component chooses which implementation/resource combination best fits the application requirements or system performance considerations. The output of this service is a platform-specific executable code segment, along with a list of its required resources, which dictate where in the WSN the service must be located. This component comprises the core of the self-mediated execution approach. Choosing an appropriate option from a list of resources and service implementations is critical to efficiently executing composite service-based applications.

Deployment.

This component is responsible for transporting the executable code segment to the destination platform, thereby making the service available to other services and applications. If an implementation of the service is already available at the destination platform, the code deployment step is skipped entirely, and the service request is sent to the deployed service.

4.3 Scheduling

Third, the meta-actor decides *when* to execute the service request. This is accomplished by the scheduling and task allocation component.

Online Scheduling and Task Allocation.

The goal of this component is to decide when the service instance can be deployed and executed. If the resources required by the service instance are not immediately available, its execution is postponed, along with all services that depend on it. Shared resources requiring exclusive access, *e.g.*, certain types of sensors and actuators, must be scheduled globally, since service implementations may not be aware of each other. An up-to-date resource use schedule is provided to the decision making and optimization component to facilitate the selection of less-utilized resources whenever possible, and a repository of active services is maintained to keep track of all service instances currently deployed in the system. This is also used by the implementation matching component to check if an already-deployed component may satisfy a service request.

4.4 Invocation and Execution

Finally, the service request is ready to be deployed and executed on the target platform. This step includes marshaling and remote invocation.

Marshaling.

The marshaling component packages the service request for transport and deployment on the destination platform, using the deployment component. The method is platform-dependent. In our system, this involves wrapping the service

invocation code in a mobile agent, which can move to the destination node without relying on an external routing service.

The service request is then handed off directly to the run-time environment to launch or query the selected implementation of the service. From this point onward, the service instance interfaces via its actor interface with its meta-actor and with other services in the WSN by means of asynchronous message passing, implemented by the communication component. Asynchronous messaging is used both to deliver computation results and error notifications from the executing services and to deliver control messages from the meta-actor.

5 Illustration

We now demonstrate how a sensor network application can be executed by our self-mediated execution architecture. As an example, we consider a distributed target tracking service similar to one proposed by Liu *et al.* [8].

Distributed target tracking is one of the canonical problems in sensor networks. Target tracking algorithms typically consist of detecting a signal emitted by the target, identifying or classifying the target by its type or signature, and once detected and classified, keeping track of its position as it changes over time. We assume that the tracking application is provided to us as a composition of Signal Detection, Target Classification and Track Maintenance application-level services, along with Localization, Time Synchronization, Routing and Group Formation middleware services, whose dependency graph is shown in Fig. 5. In this figure, Forever Do and For All Nodes Do are special control constructs, which are executed entirely by meta-actors.

Let us look at how this composite service is deployed and executed. In response to a request, the self-mediated execution architecture creates a meta-actor for the composite service, and recursively for its individual subcomponents, made possible due to the functional nature of the service composition (see Section 3.1).

Consider a request to the Signal Detection service, which is the first service instance ready to execute, due to having no dependencies. The target tracking service meta-actor requests to deploy a Signal Detection service on all nodes in the network. The Signal Detection service contract specifies that it needs a certain type of sensor, say a magnetometer, to detect the target. The resource and implementation matching components will locate a suitable implementation by pattern matching the request with service and resource descriptions.

At this point we have an executable code segment that is ready to be transported to the destination node. After the scheduling process is completed, the service request is also marshaled and transported. This in effect creates a platform-specific relocatable executable.

The only resource used by the Signal Detector service is the magnetometer; however, since multiple uncoordinated applications may be concurrently executing on the WSN, the magnetometer at the target node may currently be in use by another service. It is the responsibility of the scheduling component of the architecture to control its invocation time, such that the required resource is available prior to request

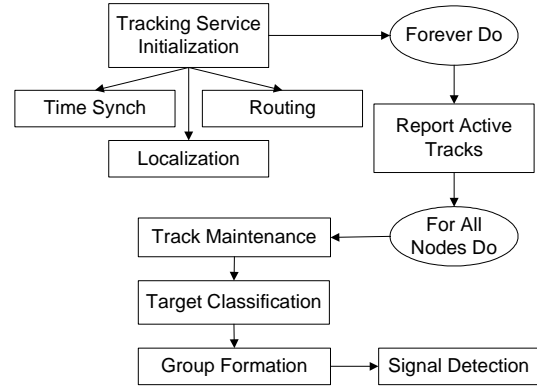


Figure 3. Composite target tracking application represented as a service dependency graph.

deployment. This means that the Signal Detector service request may be blocked from deployment until the magnetometer at its destination node becomes available.

Now consider a scenario where after the target tracking service starts executing, an intrusion detection application enters the system, ready to be executed. It is also represented as a composition of services, and happens to rely on the same target tracking service in its computation. However, its specification contains additional constraints on the Target Classification service, *e.g.*, requiring a higher confidence threshold before a target is positively identified.

Due to our design choices (dynamicity and late binding), we have an opportunity for run-time optimization. When this new application starts the self-mediated execution process, the implementation matching service lists the instances of the already-deployed services as matching the requested service contracts. This is suitable for Signal Detection and Track Management services, but the Target Classification service will fail a constraint check. With negligible incremental deployment cost, the former two service instances will be reused by the system and linked to a newly instantiated Target Classification service instance meeting the more stringent requirements of the new application.

6 Discussion

To summarize our approach, applications represented as a functional composition of services with well-defined interfaces are executed in a concurrent and distributed manner by the self-mediated execution architecture. Service implementations fitting application requirements are found and deployed on demand, sharing or reusing already-deployed implementations whenever possible. Invocation requests to these services are also generated on demand. Let us first address the benefits of taking this approach to building WSN applications.

6.1 Benefits

Late binding of service implementations and network resources is a key distinguishing feature of our architecture. By postponing the explicit identification of methods and resources until the point when they are actually used, we avoid

the problem of *overspecification*. Overspecification occurs when the programmer implicitly or explicitly supplies constraints on execution beyond what is strictly necessary to specify the desired behavior. Sampling a sensor at a *given* node within a region of interest, where sampling a sensor at *any* node within that region would have been sufficient is an example of overspecification. With late binding, we postpone the decision-making process as to which method or resource to employ from design-time to run-time, thus allowing the scheduler or optimizer components more freedom.

We also argue that service abstraction, a reusable service composition machinery, and fine-grained code deployment and execution allow creating more dynamic, maintainable and customizable applications for WSNs. Code mobility also enables predictive behavior or system-directed load balancing: a service may decide to move from one node to another to better achieve its goal, or to do so more efficiently.

6.2 Requirements

Our self-mediated execution architecture requires the application specification to be provided in the form of a composition of service descriptions. This specification may or may not be immediately executable, as not all elements are fully specified. For example, the composition may not contain a reference to a specific Target Classification service implementation, but rather to a Target Classification service contract. It is up to the mediated execution architecture to identify an appropriate implementation or resources matching the contract.

We require all composable services to conform to such a contract specification. This translates to a substantial amount of work on the service designer's part to supply a sufficiently rich service contract to turn an existing middleware service into a composable service usable by our architecture. Fortunately, the transition process can be facilitated by starting with a very rigid constraint on the interface (*e.g.*, it is only usable by the service it was originally designed for) and gradually relaxing it as a more comprehensive service contract is constructed.

The dynamic service deployment and execution process relies on the availability of a fine-grained code deployment method for the WSN, meaning that it should be possible to deploy a service to a single node or to a subset of nodes in the network at runtime.

6.3 Feasibility

We have a prototype implementation of an architecture supporting a subset of the described functionality in the context of dynamic application deployment on WSNs for Ambient Intelligence applications, called *Ambiance* [10]. This system reuses Dart [11] at its knowledge level, which is an example of a service composition framework in alignment with our design principles, for both representing applications and supporting the self-mediated execution process. Additionally, Dart supports creating intuitive Web interfaces for interactive specification of applications by multiple uncoordinated end-users at run-time. At the operational level, *Ambiance* deploys the ActorNet mobile agent platform. The interactions between these two levels conform to the logical

architecture described in prior sections.

Our architecture makes use of a service composition framework, online resource scheduling and task allocation algorithms, fine-grained runtime code deployment, and implementation- and resource-matching methods. Several approaches to these tasks have been proposed:

An extensive body of distributed resource scheduling and task allocation research is available from the real-time and parallel processing communities, and may be applied to the WSN domain given allowances for limited bandwidth, memory and processing capabilities and high likelihood of failures of typical sensor nodes.

Mobile agent platforms such as ActorNet and Agilla [7, 5] or virtual machine-based code migration systems such as Melete [13] satisfy our requirement for a fine-grained runtime code deployment method.

We consider the Decision Making and Optimization component to be one of the most challenging aspects in the implementation of our architecture. While a simple heuristic-based approach is sufficient for a prototype implementation, achieving efficient resource utilization is vital to making WSNs a suitable platform for deploying large, concurrent applications. Developing novel algorithms for this task is an important direction for future research. We believe that the clean separation of request processing and execution aspects in our architecture facilitates the integration of these components.

6.4 Potential Applications

We see a number of application opportunities for this architecture. In [10], we have described a possible application to a query processing engine for end-user defined concurrent queries integrating with sensor networks.

Another promising possibility is sensor-rich business processes, where sensors are attached to "smart items," and the interactions between these items is modeled within the business process. The goal then consists of enabling the execution on the sensor nodes of that part of the business logic. For example, in a safety process, smart chemical containers collaboratively ensure continuous compliance with certain storage regulations. Any violation of these rules results in local alerts, as well as reporting to the back-end systems [4].

Such processes are considered to increase visibility, enable real-time decision making and business process adjustment, and thus allow responding to situations more efficiently, with a higher degree of quality and end-user satisfaction. They also allow for management by exception, where the relocated processes only notify the back-end system of extraordinary situations, increasing scalability and speed of detecting situations that require action (avoiding latency of control loop), and does not require a constant connection to the back-end [12].

7 Related Work

The Melete system provides a method for concurrently executing uncoordinated applications in a sensor network [13]. Melete applications are written in the TinyScript language and executed by a virtual machine on an arbitrary

subset of nodes in the network. We propose a more comprehensive method of executing concurrent applications in WSNs, which allows global resource management and a higher level of optimization. In fact, Melete may be used as part of our architecture, acting as the code deployment method for service instances.

The Tenet architecture enables service composition for multi-tiered applications incorporating WSNs [6]. Most of the coordination and processing functionality is relegated to more powerful tiers, while the WSN nodes are used primarily to retrieve sensor data. Our approach differs in that we treat the sensor network as a collaborative distributed computing platform. By associating asynchronously interacting, autonomous actors to service instances on sensor nodes we make possible *in situ* collaborative problem solving.

The SONGS architecture and programming model considers sensor network applications as a composition of semantic services [9]. Semantic services are a type of semantic data transformation functions, and do not correspond to what we call services in this paper. We are interested in facilitating composition of less structured infrastructure and middleware services, a vast quantity of which has already been developed for wireless sensor networks.

The European-funded project Collaborative Business Items (CoBIs) [3, 4] is also concerned with *in situ* processing and coordination for embedded devices, while integrating them in the context of higher-level business processes.

8 Perspectives

An expanded version of the current architecture prototype [10, 11] is under development. The primary focus is on incorporating the decision making and optimization components. This includes a study of which aspects of low-level service optimization and control decision can be externalized.

We believe that the design principles and architecture defined in this paper have wider implications beyond the adaptive execution of composite middleware services in WSNs. We are specifically interested in coordination behaviors within the WSN as well as its relation to outside platforms and applications. We are thus investigating the scalability of our architecture in the context of complex hierarchical scientific and business processes running in a pervasive computing environment, which also includes sensor networks.

9 Acknowledgments

This work is partially funded by the University of Luxembourg in the framework of the Ambiance project (R1F105K04), by NSF under grant CNS 05-09321 and by ONR under DoD MURI award N0014-02-1-0715 to UIUC. The authors would also like to acknowledge the valuable collaboration of A. Cardon, N. Bouraqadi, P. Bouvry, Ch. Dony, V. Ginot, R. Kumar, R. Johnson, M. Malvetti, T. Odenwald, J.-F. Perrot, D. Riehle, S. Sundresh, and J. Yoder.

10 References

- [1] G. Agha. *Actors: A Model of Concurrent Computation in Distributed Systems*. MIT Press, 1986.
- [2] N. Carriero and D. Gelernter. Linda in context. *Communications of the ACM*, 32(4):444–458, 1989.
- [3] L. W. F. Chaves, J. Anke, L. M. S. de Souza, and J. Muller. Service lifecycle management infrastructure for smart items. In *MidSens'06*, November 2006.
- [4] C. Decker, P. Spiess, L. M. sa de Souza, M. Beigl, and Z. Nochta. Coupling enterprise systems with wireless sensor nodes: Analysis, implementation, experiences and guidelines. In *Pervasive Technology Applied @ PERVASIVE*, May 2006.
- [5] C. L. Fok, G.-C. Roman, and C. Lu. Mobile agent middleware for sensor networks: An application case study. In *4th International Conference on Information Processing in Sensor Networks (IPSN'05)*, pages 382–387, April 2005.
- [6] O. Gnawali, B. Greenstein, K.-Y. Jang, A. Joki, J. Paek, M. Vieira, D. Estrin, R. Govindan, and E. Kohler. The TENET architecture for tiered sensor networks. In *ACM Conference on Embedded Networked Sensor Systems (SenSys)*, November 2006.
- [7] Y. Kwon, S. Sundresh, K. Mechitov, and G. Agha. ActorNet: An actor platform for wireless sensor networks. Technical Report UIUCDCS-R-2005-2595, University of Illinois at Urbana-Champaign, 2005.
- [8] J. Liu, J. Liu, J. Reich, P. Cheung, and F. Zhao. Distributed group management for track initiation and maintenance in target localization applications. In *2nd International Workshop on Information Processing in Sensor Networks (IPSN)*, 2003.
- [9] J. Liu and F. Zhao. Towards semantic services for sensor-rich information systems. In *International Workshop on Broadband Advanced Sensor Networks (Basenets 2005)*, October 2005.
- [10] R. Razavi, K. Mechitov, G. Agha, and J.-F. Perrot. Dynamic macroprogramming of wireless sensor networks with mobile agents. In *2nd Workshop on Artificial Intelligence Techniques for Ambient Intelligence*, January 2007.
- [11] R. Razavi, J.-F. Perrot, and R. Johnson. Dart: A meta-level object-oriented framework for task-specific, artifact-driven behavior modeling. In *Proceedings of DSM'06*, pages 43–55, 2006.
- [12] P. Spiess, H. Vogt, and H. Jutting. Integrating sensor networks with business processes. In *Real-World Sensor Networks Workshop at ACM MobiSys*, June 2006.
- [13] Y. Yu, L. J. Rittle, V. Bhandari, and J. B. LeBrun. Supporting concurrent applications in wireless sensor networks. In *4th International Conference on Embedded Networked Sensor Systems (SenSys)*, pages 139–152, 2006.