

Enhanced Coordination in Sensor Networks through Flexible Service Provisioning

Chien-Liang Fok, Gruia-Catalin Roman, and Chenyang Lu

Dept. of Computer Science and Engineering
Washington University in St. Louis
Saint Louis, MO, 63105, USA
[liang, roman, lu]@cse.wustl.edu

Abstract. Many applications operate in heterogeneous wireless sensor networks, which represent a challenging programming environment due to the wide range of device capabilities. Servilla addresses this difficulty in developing applications by offering a new middleware framework based on service provisioning. Using Servilla, developers can construct platform-independent applications over a dynamic and diverse set of devices. A salient feature of Servilla is its support for the discovery and binding to local and remote services, which enables flexible and energy-efficient in-network collaboration among heterogeneous devices. Furthermore, Servilla provides a modular middleware architecture that can be easily tailored to devices with a wide range of resources, allowing resource-constrained devices to provide services while leveraging the capabilities of more powerful devices. Servilla has been implemented on TinyOS for two representative hardware platforms (Imote2 and TelosB) with drastically different resources. Microbenchmarks demonstrate the efficiency of Servilla's implementation, while an application case study on structural health monitoring demonstrates the efficacy of its coordination model for integrating heterogeneous devices.

1 Introduction

Wireless sensor networks (WSNs) [17] are becoming increasingly heterogeneous due to two primary reasons. First, heterogeneity allows a network to be both computationally powerful and deployed in high densities. Powerful devices can perform complex operations, but are more expensive and power-hungry. Conversely, weak WSN devices enable higher deployment densities and increase network lifetime as they are cheaper and consume less power. By integrating devices with different resources and capabilities, a heterogeneous WSN can combine the advantages of both powerful and weak devices. Second, network heterogeneity follows from the natural evolution of WSNs. WSN devices can be embedded in the environment and remain operational for a long time. For example, due to its high deployment cost, a WSN embedded in civil infrastructure for structural health monitoring must operate over several years to be economically acceptable [24]. Similarly, many urban sensing systems [48] must also remain operational for multiple years. During the lifetime of a WSN, new devices may be developed and deployed, resulting in network heterogeneity.

Network heterogeneity presents a formidable problem for application developers. Since the target platform may consist of many different devices, the application must

be platform-independent to avoid having to custom-tailor it to each device. Yet, the application must still be able to access platform-specific capabilities like sensing and computing to make full use of the underlying hardware. Furthermore, the application must accommodate diverse device capabilities and resources. These seemingly contradictory requirements complicate application development and motivates a new programming model.

To address the challenges of programming heterogeneous WSNs, we developed **Servilla**, a middleware framework supporting a novel coordination model. Servilla's coordination model makes three important contributions. First, applications are structured in terms of *platform-independent* tasks and expose *platform-specific* capabilities as services. This ensures that applications remain platform-independent, which is critical as WSNs become increasingly heterogeneous. It also enables applications to access resources on a device without having active processes, or agents, on the device itself. This reduces the system's minimum resource requirements, increasing the range of devices that can be supported. Second, Servilla provides a specialized service description language, which enable application tasks to selectively but flexibly access services that exploit the capabilities of the hardware available at a particular time and place. This allows better adaptation to network heterogeneity, and facilitates for the first time in-network collaboration between heterogeneous WSN devices, achieving higher levels of efficiency and flexibility. Finally, Servilla provides a modularized middleware architecture and enables asymmetry in the middleware among WSN devices. This widens the scope of hardware devices that can be integrated.

Servilla's coordination model is inspired by the concept of *Service-Oriented Computing* (SOC) [45], which provides loose and flexible coupling between application components. It is used on the Internet and has recently been explored in the context of WSNs. Two systems in particular are Tiny Web Services (TWS) [47] and PhyNetTM [6]. TWS implements an HTTP server on each device and enables applications outside of the WSN to invoke services over the Internet using HTTP requests. PhyNetTM provides a central gateway that exposes WSN capabilities as web services. Unlike these systems, Servilla uniquely takes the SOC programming model *inside* a WSN. It exploits the loose coupling between service consumers and providers to separate application-level platform-independent logic from the low-level software components that exploit platform-specific capabilities. Furthermore, by allowing application logic to execute inside a WSN, higher levels of efficiency are obtainable via in-network coordination and collaboration [28]. For example, in a structural health monitoring application, a low-power device may use a simple threshold-based algorithm to detect shocks that are potentially damage-inducing, and only activate more powerful devices that perform the complex operations to localize damage when necessary [24]. Or, in a surveillance application, low-power devices may sense vibrations from an intruder and activate more powerful devices with cameras [26]. The ability to support collaboration among heterogeneous devices *inside* a WSN is a key feature that distinguishes this work from existing SOC middleware for WSNs.

The remainder of the paper is organized as follows. Section 2 presents Servilla's programming model. Section 3 presents Servilla's programming languages. Section 4 presents Servilla's middleware architecture and implementation. Section 5 presents an

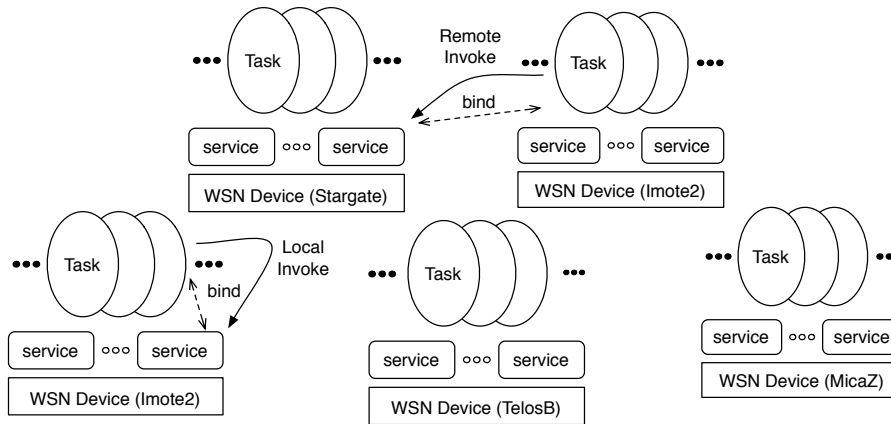


Fig. 1. Servilla targets heterogeneous WSNs in which different classes of devices provide services that are used by application tasks either locally or remotely. Services are platform-specific while tasks are platform-independent.

empirical evaluation on two representative sensor platforms with diverse resources. Section 6 evaluates the efficacy of Servilla by using it to implement a structural health monitoring application. Section 7 presents related work. The paper ends with conclusions in section 8.

2 Programming Model

An overview of a WSN using Servilla is shown in Figure 1. Servilla is meant for applications that run in WSNs with multiple classes of devices. It is not intended for flat WSNs composed entirely of resource-poor devices. Typical applications are long-lived, widespread, and involve many different tasks that vary in complexity and scale. They are expected to be written once, but continuously used despite hardware changes. For example, environmental monitoring or target tracking applications are long-lived and usually involve both widespread, but simple sensing tasks, and less widespread, but complex computational tasks, that process the data. By integrating both resource-poor and resource-rich devices, Servilla provides an ideal platform on which to build these types of applications. Specifically, resource-poor devices are less costly and more energy efficient, meaning they can be deployed in greater numbers at higher densities. Meanwhile, resource-rich devices are more expensive and limited in quantity, but offer computational power and advanced sensing capabilities.

Applications are implemented as tasks, which are platform-independent application processes that contain code, state, and service specifications. To ensure platform-independence, the code cannot directly access platform-specific capabilities like sensors. Instead, these capabilities are accessed as services that are provided by a *service provisioning framework*. The service provisioning framework takes a task's service specifications and finds services that match them. The service specifications describe

both the service's interface and non-functional properties like energy consumption. This enables tasks, for example, to selectively use the most energy-efficient sensors.

Services expose platform-specific capabilities, are implemented natively, and can, thus, be fine-tuned for maximum efficiency. They provide a description that can be compared with a task's service specification. Services are able to maintain state, provide multiple methods, and have their own thread of control, enabling them to operate in parallel with tasks. This enables higher degrees of concurrency and efficiency. For example, in a structural health monitoring application, a service provided by a low-power device can continuously monitor an accelerometer and set a flag if the vibrations exceed a threshold. A task executing on a more powerful device can mostly remain asleep, only periodically checking for potential damage.

Tasks communicate via localized tuple spaces that are structured in the same manner as that in Agilla [19]. For brevity, they are not shown in Figure 1 since service provisioning is the focus and main contribution of this paper. Tuple space coordination facilitates decoupled communication, allowing better adaptation to a changing network. They serve as a flexible means of communication between application processes and are orthogonal to service provisioning. While service provisioning messages could be sent using tuple spaces, they are sent in an RPC-like fashion in the current implementation.

Tasks remain platform-independent by delegating all platform-specific operations to services. There are two essential steps for this to occur: *binding* and *invocation*. Binding is the process of discovering and establishing a connection to the service. Invocation is the process of accessing a service.

Service binding consists of three-steps: discovery, matching, and selection. Discovery involves finding available services. In many traditional SOC frameworks, this is done by querying a central service registry. While this is sufficient in traditional networks, it is not appropriate in WSNs for a couple reasons. First, since most WSN devices operate on batteries, accessing a distant registry is not energy efficient and can unacceptably reduce network lifetime. Second, the spatial aspect of WSNs are relevant since closer services are usually preferred, e.g., if a task wants to know the temperature, it usually wants to know the ambient temperature rather than a distant location's. For these reasons, Servilla is optimized for *localized* coordination and does not rely on a centralized service registry. Instead, each device has its own registry containing only the services that it provides.

During the service discovery process, the local registry is first checked for a match. If no match is found, neighboring devices are checked. This increases a network's flexibility by allowing tasks to run on devices that do not fully satisfy the service requirements, since missing services can be provided by neighboring devices. Furthermore, although accessing a remote service requires wireless communication, energy efficiency can be increased overall by allowing high-power devices to use low-power ones, enabling the high-power devices to remain asleep longer.

Service matching involves finding a service that fulfills a task's requirements. Recall that tasks include specifications that can be compared to descriptions provided by services. The matching process must be flexible since the service and tasks are usually developed separately. Yet, it must be semantically correct to ensure that the service behaves in a predictable manner. A service is minimally described by its interface. Ide-

ally, the names of the methods, the order, number, and types of their parameters, and even the return types should not require an exact match for service binding for maximum flexibility. To achieve this, large amounts of meta-data must be included in the specification that describe the method names, input parameters, and return values. Unfortunately, such a specification is verbose and requires a complex parser, both of which consume sizable computational resources that are not available on many WSN devices. To account for this, Servilla compromises by dividing specifications into functional and non-functional properties. Functional properties include the interface and require an exact match. Nonfunctional properties describe attributes like power consumption and do not require an exact match. For example, suppose a FFT-calculating service has a non-functional attribute specifying that it is version 5. Such a service can be bound to a task that specifies it requires *at least* version 4. By enforcing an exact match between functional properties and an inexact match between non-functional ones, Servilla provides a degree of flexibility when binding services while still maintaining reasonable resource requirements.

Once a matching service is found, the binding process is completed by selecting it. Selection consists of informing the task of the chosen service, and is accomplished by informing the task of the provider's network address. Once done, the task can access the service by invoking it. Note that the provider's address is hidden from the application developer, who is able to invoke the service based on its name, a process that is described next.

Service invocations are analogous to remote procedure calls (RPCs). The task provides the name of the service, the method to execute, and the input parameters. After the service executes, the results are returned to the task. Since the task and service may be located on different devices, the process may fail, e.g., due to message loss. To account for this, Servilla provides a mechanism that notifies a task when and why an invocation fails. This is necessary because service invocations may fail in many ways depending on whether the service is local or remote, and tasks may want to handle various error conditions differently. For example, local invocations may fail because the service is busy, in which case the task may try again later, while remote invocations may fail due to disconnection, in which case the task may want to abort.

3 Programming Language

Servilla provides two light-weight programming languages tailored to support service provisioning in WSNs. The first, *ServillaSpec*, is used to create service specifications and descriptions that enable flexible matching between tasks and services. The second, *ServillaScript*, is used to create tasks and is compiled into bytecode that runs on a Virtual Machine. Services are implemented in NesC [21] on TinyOS [27] and compiled into native binary code for run-time efficiency. Servilla's specialized languages are now described.

3.1 ServillaSpec

ServillaSpec is used to describe services and is needed to match services required by tasks to those provided by devices. To support resource-constrained devices, the ser-

```
NAME = fft
METHOD = fft-real
INPUT = {int dir, int numSamples, float[] data}
OUTPUT = float[]
ATTRIBUTE Version = 5.0
ATTRIBUTE MaxSamples = 5000
ATTRIBUTE Power = 10
```

Fig. 2. A specification describing a FFT service

```
1. uses Temperature; // declare required service
2.
3. void main() {
4.     int count = 0; float temp;
5.     bind(Temperature, 2); // bind service within 2 hops
6.     while(count++ < 10) {
7.         temp = invoke(Temperature, "get"); // invoke service
8.         send(temp);
9.     }
10.    unbind(Temperature);
11. }
```

Fig. 3. A task that invokes a temperature sensing service 10 times

vice specification language must be compact and should not require an overly complex matching algorithm. As such, standard specification languages used on the Internet like WSDL [49] are avoided due to their relative verbosity and highly complex parsers. ServillaSpec avoids verbose syntax and limits the types of properties that can be included in a service specification. An example is shown in Figure 2. The first line specifies the name of the service. It is followed by three-line segments each specifying the name, input parameters, and output results of a method provided by the service. The remainder of the specification is a list of attributes that specify non-functional properties of the service. They enable flexibility in matching by defining a name, relation, and value. Using attributes, a task can, for example, require a floating point FFT service that consumes *at most* 50mW. Such a specification would match a service whose description is shown in Figure 2.

By limiting the property types to be only the five shown in Figure 2 (i.e., NAME, METHOD, INPUT, OUTPUT, and ATTRIBUTE), and arranging them to always be in the same order, the specification can be greatly compressed. For example, since the service's NAME property always appears first, the property's identifier, NAME, can be omitted. Thus, the NAME property in the specification shown in Figure 2 can be compressed to just 4 bytes, "fft" followed by a null terminator. This compression saves memory and enables greater matching efficiency.

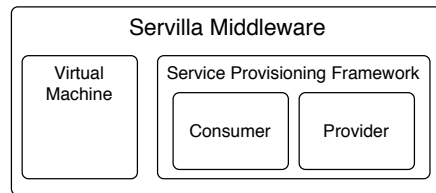


Fig. 4. Servilla’s middleware consists of a virtual machine and a service provisioning framework (SPF). The SPF consists of a consumer and provider.

3.2 ServillaScript

ServillaScript is used to create application tasks. Its syntax is similar to other high level languages like JavaScript [18], but with key extensions for service provisioning. An example, shown Figure 3, implements an application that periodically takes the temperature and sends the reading to the base station. It declares the name of the file containing the specification of the required service on line 1, which in this case is a temperature sensing service. The task initiates the service binding process on line 5. The first parameter specifies which service to bind and the second specifies that registries within two hops should be searched. The task then loops ten times invoking the service on line 7 and sending the temperature to the base station on line 8. The task ends by disconnecting from the service on line 10.

The example above illustrates how ServillaScript enables tasks to 1) indicate which services are needed, 2) initiate the service discovery process, 3) invoke services, and 4) disconnect from services. Aspects not shown for brevity include checking whether a service is bound, and, if so, how far away the service is in terms of network hops. This will allow the task to throttle how often it invokes the service based on its distance. Another aspect not shown is error handling code. If an error occurs due to a service becoming unavailable, the invocation will return an error indicating the cause, as discussed in Section 2.

4 Middleware

Servilla’s middleware architecture, as shown in Figure 4, consists of a virtual machine (VM) and a service provisioning framework (SPF). The VM is responsible for executing application tasks. The SPF consists of a consumer (SPF-Consumer) that discovers and accesses services, and provider (SPF-Provider) that advertises and executes services.

A VM is used because WSN devices contain processors with varying instruction sets. Application tasks are compiled into the VM’s instruction set, which is uniform across all hardware platforms, ensuring that tasks are platform-independent. Furthermore, the VM enables the dynamic deployment of application tasks, justifying the need for dynamic service binding. The VM is based on Agilla [19] though with major extensions to support services and the SPF. Specifically, whenever a task performs an operation involving a service, the VM passes the task to the SPF-Consumer, which is described next.

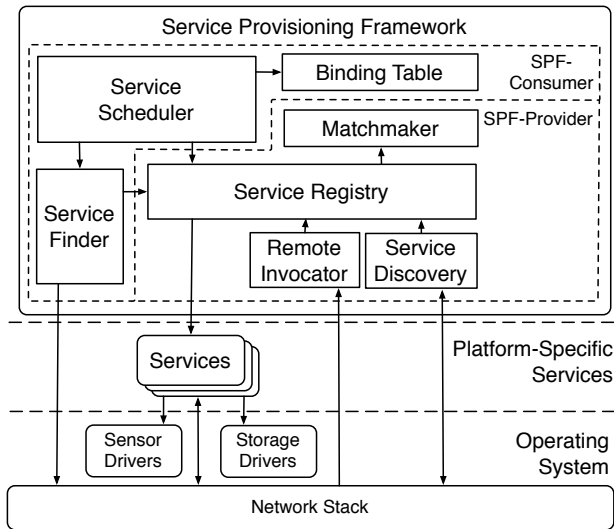


Fig. 5. The detailed architecture of the Service Provisioning Framework.

4.1 SPF-Consumer

The SPF-Consumer is responsible for discovering, matching, and invoking services on behalf of tasks. As shown in Figure 5 the SPF-Consumer consists of a Service Finder, Binding Table, and Service Scheduler. The Service Finder is responsible for finding services that match a task's specifications. It first searches locally and, if no matches are found, searches one hop neighbors. Note that while this increases the likelihood of selecting a local service, it does not necessarily select the most energy efficient provider. If a task wanted to bind to an energy efficient provider, it can include an energy attribute in its service specification, thus enabling energy efficient service provisioning. When a provider is selected, its address is stored in the Binding Table. The Binding Table maps the task's service specification to the provider that will perform the service. It is updated when the Service Finder discovers a better provider and when a task explicitly unbinds from a service. A task can query a Binding Table to determine whether it has access to a particular service.

The Service Scheduler carries out the actual invocation. It takes the input parameters provided by the task, sends them to the provider, and waits for the results to arrive. Once the results arrive, it passes them to the task which can then resume executing. If the results do not arrive within a certain time, the Service Scheduler aborts the operation and notifies the task of the error.

4.2 SPF-Provider

The SPF-Provider is responsible for providing and executing services. Its architecture, shown in Figure 5, consists of a Service Registry, Matchmaker, Remote Invocator, and Service Discovery component. The Service Registry contains the specifications

	TelosB	Imote2
Processor	8MHz 16-bit TI MSP430	13-416MHz 32-bit Intel PXA271 XScale
Radio	IEEE 802.15.4	IEEE 802.15.4
Memory	48KB Code, 10KB Data	32MB Shared
Price	\$99	\$299

Table 1. WSN devices vary widely in computational resources.

of all locally-provided services. The Matchmaker is used to determine whether a service meets the requirements of a task. When the SPF-Consumer tries to find a service, the Matchmaker is used to determine whether a matching service exists. Note that in this architecture, the task’s specification must be sent from the SPF-Consumer to the SPF-Provider. This is because the Matchmaker is located on the SPF-Provider. Alternatively, the Matchmaker can be moved onto the SPF-Consumer to reduce the footprint of the SPF-Provider. However, this requires that all specifications be sent to the SPF-Consumer, a process that may incur higher communication cost.

4.3 Middleware Modularity

WSNs are becoming extremely diverse consisting of devices with resources that differ by several orders of magnitude [46, 15]. This is true even as technology improves, since cost constraints ensure the continued presence of resource-limited devices. To accommodate the wide range of devices, Servilla’s middleware is modularized and configurable such that a device need not implement every module to participate in the network. For example, the middleware can be configured in the following ways:

- **VM + SPF:** The full Servilla framework.
- **VM + SPF-Consumer:** Executes tasks and provides access to remote services only.
- **SPF-Provider:** Provides services for neighboring tasks to use.

A detailed analysis of the memory consumed by each configuration is given in Section 5.1. The configuration containing only the SPF-Provider is particularly interesting because it allows resource-weak, but energy efficient, devices to provide services to more powerful devices. This can result in greater overall energy efficiency and, assuming the weak devices are less costly and more numerous, increase sensing density while achieving greater sensing coverage.

The various middleware configurations are *transparent* to tasks due to the decoupled nature of the SOC model. For example, a task need not know whether there is a local SPF-Provider. If a task requires a service, it will be bound either locally or remotely depending on availability.

4.4 Implementation

Servilla has been implemented on TinyOS 1.0 and two representative hardware platforms shown in Table 1. It is divided into two levels as shown in Figure 6: a lower

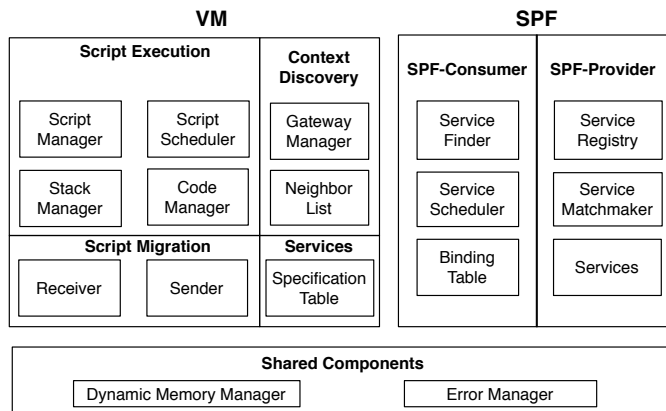


Fig. 6. Servilla's middleware components.

level consisting of shared components and a higher level consisting of Servilla's VM and SPF. This section first discusses the lower level followed by the upper level. It ends with a discussion of Servilla's programming languages.

The shared components implement low-level mechanisms needed by most high-level components. The dynamic memory manager makes more efficient use of memory. This is important because Servilla has several components that require varying amounts of memory over time. The dynamic memory manager provides just enough memory for each higher-level component to complete their function and reclaims the memory when it is no longer needed. It is shared by most components in Servilla's middleware, maximizing the flexibility of memory allocation. To aid in debugging, Servilla provides an error manager that detects and sends summaries of problems to the base station. The error manager is shared by all other components in Servilla's middleware.

The SPF is implemented natively using NesC and is divided into two modules, the SPF-Consumer and SPF-Provider, as shown in Figure 6. In the SPF-Consumer, the implementation of the Service Scheduler is simplified by serializing service invocations. This has the added benefit of avoiding saturating the wireless channel. To increase energy efficiency, the Service Finder first searches the local Service Repository, if one exists, before searching one-hop neighbors. In the SPF-Provider, the Service Registry is able to support up to 256 local services.

Servilla's compiler can compile ServillaScript and ServillaSpec into a compact binary format. For example, the task shown in Figure 3 is compiled into 181 bytes of code and 30 bytes of specifications, and the specification shown in Figure 2 is compiled into just 64 bytes. Both the Servilla middleware and compiler have been released as open-source software at <http://mobilab.wustl.edu/projects/servilla/>.

5 Evaluation

This section presents empirical measurement of the code size and performance overhead of Servilla on both the TelosB [46] and Imote2 [15] platforms. The efficacy of the

Servilla programming model is demonstrated through an application case study in the next section.

5.1 Memory Footprint

An Imote2 has sufficient memory (32MB) to hold the entire Servilla middleware. Compiled for the Imote2, the total size of the middleware without services is a mere 318KB. This is only about 1% of the total, leaving plenty for services. In contrast, TelosB devices only have 48KB of code memory. While TelosB does not have enough memory to hold every component, it can support the SPF-Provider configuration which only consumes 32 KB of code memory. This capability allows TelosB to join and contribute to a WSN as providers of services to more powerful devices. As shown in previous work [22] and our case study presented in Section 6, effective integration of resource-constrained and more powerful devices can combine the advantages of pervasive low-power sensing and computational resources, and enhance energy efficiency. This example shows how Servilla’s modular architecture enables support of diverse hardware platforms.

5.2 Efficiency of Service Binding

Service binding consists of three parts: discovery, matching, and selection. This study first focuses on discovery followed by matching and selection. Recall that the current implementation requires the Service Finder to query each neighbor individually for a match. This is because Servilla uses a reliable network interface that does not support wireless broadcasts. To optimize the selection, the Service Finder first searches locally before remotely. Since the latency of a local search is negligible, we evaluate the latency of a remote search.

The latency of a remote search depends on the number of neighbors, the percentage of them that provide a matching service, and the order in which they are queried. At a minimum, one neighbor will be queried. This study evaluates only a single query, since each additional query will proportionately increase the latency. An Imote2 is used to query a TelosB to determine whether the TelosB provides a particular service. In this case, the service being queried is `FFT` and the specification is shown in Figure 2. It is compiled into 64 bytes, which must be sent from the Imote2 to the TelosB. Due to various bookkeeping variables, the size of the query message is 72 bytes, and the reply message is 16 bytes. The time between sending the query to receiving a reply is measured by toggling a general I/O pin before and after the query, and capturing the time between toggles using an oscilloscope. Averaged over 100 trials, the latency and 90% confidence interval is $245.6 \pm 1ms$. This latency is acceptable to many WSN applications. Moreover, it may be amortized over multiple invocations of the same service after it is bound to the task.

To evaluate the efficiency of service matching, the Matchmaker is used to compare two copies of `FFT`, shown in Figure 2. This incurs the worst-case latency since every property within the specification must be compared. Each experiment is repeated twenty times on both TelosB and Imote2 platforms running at all possible CPU speeds

device	CPU Speed	Bus Speed	Sig.	Attr. 1	Attr. 2	Attr. 3	Other	Total	Units
TelosB	8MHz	8MHz	18	14	24	29	8	92	ms
Imote2	13MHz	13MHz	1569	1421	2642	3272	784	9688	μs
Imote2	104MHz	104MHz	198	180	330	408	94	1209	μs
Imote2	208MHz	208MHz	99	89	165	204	47	604	μs
Imote2	416MHz	208MHz	71	62	113	136	31	413	μs

Table 2. Service matching latency when comparing two FFT-real service specifications

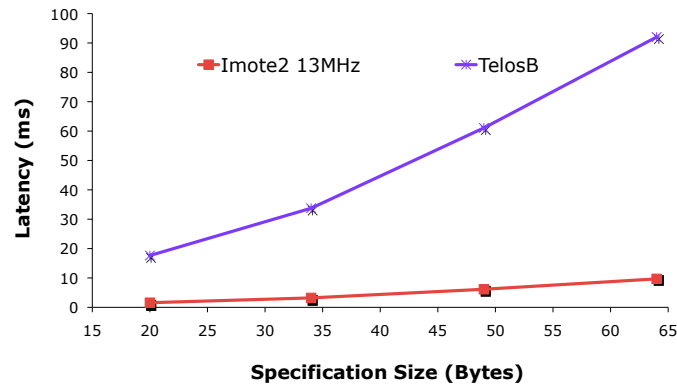


Fig. 7. The latency of comparing a specification vs. its size.

and the average latency is calculated. The results are shown in Table 2.¹ They indicate that the TelosB takes about 92ms to perform a match, while the Imote2 is at least ten times faster. The latencies are small compared to the execution times of certain VM instructions. Note that while service matchmaking does introduce overhead, it is usually done infrequently relative to service invocation.

To determine how the specification's size affects matching latency, FFT is compared to versions of itself with one, two, and all three of its attributes removed. The matching latencies is plotted against their sizes and the results are shown in Figure 7. For brevity, only the Imote2 running at 13MHz is shown. The latency when the Imote2 is running at higher frequencies is significantly lower. The results indicate that the latency is roughly proportional to its size. It is not exactly proportional because of the additional overhead incurred with the addition of each attribute, as indicated by the "other" column in Table 2.

6 Application Case Study

This section evaluates Servilla using an application case study, specifically one designed to localize damage in structures (e.g., a bridge). The application enables real-time evaluation of a structure's integrity, reducing manual inspection costs while increasing

¹ The confidence intervals are negligible since the experiment runs locally and the measurements exhibit very low variance.

```

NAME = AccelTrigger
METHOD = start
INPUT =
OUTPUT =
METHOD = stop
INPUT =
OUTPUT =
METHOD = check
INPUT =
OUTPUT =
ATTRIBUTE power = ...

```

(a) The specification of service `AccelTrigger` provided by Imote2 and TelosB devices. The power attribute specifies the amount of power the service consumes. It is 145mW on the Imote2, and 9mW on the TelosB.

```

NAME = AccelTrigger
...
ATTRIBUTE power < 50

```

(b) The specification of a low-power version of service `AccelTrigger`, which is provided by the application task. Its interface is omitted since it is the same as the one in Figure 8(a). A high-power version has attribute `power ≥ 50 mW`.

```

NAME = DLAC
METHOD = find
INPUT =
OUTPUT = float[25]

```

(c) The specification of service `DLAC` provided by Imote2 devices.

Fig. 8. The services used by the damage localization application

safety. WSNs have recently been used to successfully localize damage to experimental structures using a homogeneous network of Imote2 devices [24]. In this case, the algorithm, called Damage Localization Assurance Criterion (DLAC), was written using NesC specifically for the Imote2. The implementation using Servilla generalizes and improves upon the original by making it platform-independent and increasing its overall energy efficiency by exploiting network heterogeneity.

The heterogeneous WSN used in this study consists of TelosB and Imote2 devices. DLAC can only run on the Imote2 due to insufficient memory on the TelosB. However, Imote2 devices consume significantly more energy than TelosB devices. Thus, using Servilla, the application can combine the advantages of both platforms by using the TelosB devices to monitor the ambient vibration levels, allowing the Imote2 devices to sleep longer. Ideally, the Imote2 devices should only be activated to perform the DLAC algorithm when the TelosB devices detect that the ambient vibration levels exceed a certain damage-inducing threshold. The dual-level nature of this configuration is common to other applications like surveillance [26], and is essential for conserving energy and increasing network lifetime.

The Servilla implementation relies on two services: `AccelTrigger` and `DLAC`. Ambient vibrations are monitored by `AccelTrigger`, which sets a flag when a threshold is exceeded. Its specification is shown in Figure 8(a). The service has three methods: `start`, `stop`, and `check`. Methods `start` and `stop` control when the service monitors the local accelerometer. The status of the flag is obtained by invoking `check`.

```

1. uses AccelTiggerHP;
2. uses AccelTiggerLP;
3. uses DLAC;
4.
5. void main() {
6.     bind(DLAC, 0); // bind DLAC service
7.     if(!isBound(DLAC)) exit(); // failed to bind DLAC
8.     bind(AccelTriggerLP, 1); // bind low-power AccelTrigger service
9.     if(isBound(AccelTriggerLP)) {
10.        invoke(AccelTriggerLP, "start");
11.        waitForTrigger(1);
12.    } else {
13.        bind(AccelTriggerHP);
14.        if(isBound(AccelTriggerHP)) {
15.            invoke(AccelTriggerHP, "start");
16.            waitForTrigger(0);
17.        }
18.    }
19. }
20.
21. void waitForTrigger(int useLowPower) {
22.     int vibration = 0;
23.     while(vibration == 0) {
24.         if (useLowPower)
25.             vibration = invoke(AccelTriggerLP, "check");
26.         else
27.             vibration = invoke(AccelTriggerHP, "check");
28.         if (vibration == 1) {
29.             if (useLowPower)
30.                 invoke(AccelTriggerLP, "stop");
31.             else
32.                 invoke(AccelTriggerHP, "stop");
33.             doDLAC();
34.         }
35.         sleep(1024*60*5); // sleep for 5 minutes
36.     }
37. }
38.
39. void doDLAC() {
40.     float[25] dlac_data;
41.     dlac_data = invoke(DLAC, "find");
42.     send(dlac_data); // send DLAC data to base station
43. }

```

Fig. 9. The damage localization application task

Both the Imote2 and TelosB devices provide `AccelTrigger`. They differ in their power attribute, since the Imote2 consumes more power than the TelosB (145mW vs. 9mW).

The specification of service `DLAC` is shown in Figure 8(c). It contains a single method, `find`, that takes no parameters and returns an array of floating-point numbers that are used to localize damage to the bridge [24].

The application's task is shown in Figure 9. The first three lines specify the names of the files containing the required service specifications. The content of `AccelTriggerLP` is shown in Figure 8(b), and the content of `DLAC` is shown in Figure 8(c). Notice that `AccelTriggerLP` matches the TelosB version of the `AccelTrigger` service shown in Figure 8(a) because its power attribute is less than 50mW. `AccelTriggerHP` contains the same specification as `AccelTriggerLP` except its power attribute is ≥ 50 mW, which matches the service provided by the Imote2.

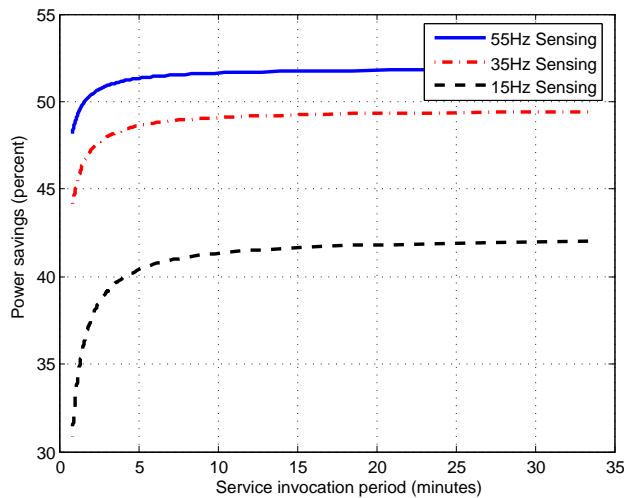


Fig. 10. Percent power savings of heterogeneous vs. homogeneous WSN.

The application attempts to reduce energy consumption by preferentially binding to an `AccelTrigger` service that consumes less power. It does this by first attempting to bind using the specification within `AccelTriggerLP` on line 8, before using the specification within `AccelTriggerHP` on line 13. Once an `AccelTrigger` service is bound, the task periodically queries it to determine if the acceleration readings are above a certain threshold (lines 21-37). If it is, `DLAC` is invoked and the results are sent to the base station (lines 39-43).

To evaluate the benefit of exploiting network heterogeneity on *Servilla*, the task shown in Figure 9 is injected into two WSNs: a homogeneous network consisting of only *Imote2* devices, and a heterogeneous network consisting of both *Imote2* and *TelosB* devices. Since the application is written using *Servilla*, it is able to run on both types of networks without modification. In both cases, `DLAC` is executed by the *Imote2*, meaning the power consumption of performing damage localization is constant. However, the power consumption of `AccelTrigger` varies because *Servilla*'s service provisioning framework enables an application to exploit more energy-efficient services when possible in a platform-independent and declarative fashion. Specifically, if *TelosB* devices are present, the service will be executed on a *TelosB* device since its `AccelTrigger` service consumes less power, otherwise it will be executed on the *Imote2*. We compare the power consumption of invoking `AccelTrigger` in different network configurations.

Since invoking `AccelTrigger` on the *TelosB* requires a remote invocation, the amount of energy saved depends on the invocation and sensing frequencies. If the service is invoked too often, more energy will be spent on wireless communication. Likewise, if the sensor is accessed very infrequently, the benefits of the *TelosB* is diminished since the devices will remain asleep a larger percentage of the time. To determine how much energy savings is possible, an oscilloscope is used to measure the time each platform spends computing, communicating wirelessly, and sensing, in both a homo-

geneous and heterogeneous network. The sensing frequency is varied between 15Hz and 55Hz (the maximum sampling frequency of the TelosB), and the service invocation frequency is varied between 50 seconds to 35 minutes. The percent savings of using a heterogeneous network relative to a homogeneous network is then calculated and the results are shown in Figure 10.

The results show that the heterogeneous implementation using Servilla achieves up to 52% power savings, and the savings increase with sensing frequency. They also show that invoking the service too frequently will reduce the amount of power saved since doing so incurs more network overhead. There is a limit to the amount of energy that can be saved as the service invocation period increases since it approaches the difference between the sensing energy consumed by the Imote2 versus the TelosB.

This case study demonstrates how Servilla enables platform-independent applications that operate over a heterogeneous WSN, and how Servilla facilitates in-network collaboration between different types of devices to attain higher energy efficiency. Moreover, it demonstrates that Servilla enables an application to bind to a more energy-efficient service through service specification.

7 Related Work

SOC has long been used on the Internet to enable independently-developed applications to interoperate. There are many SOC systems including SLP [30], Jini [31], OSGi [44], CORBA, and Web Services [2]. Servilla has three salient features that distinguish it from these SOC frameworks. First, it focuses on how service-provisioning language and middleware can be made extremely lightweight. This is necessary due to the limited resources available on many WSN devices. Many previous SOC systems have been ported to PDA-class devices, which are more powerful than the low-power sensor devices supported by Servilla. Second, Servilla is specifically designed for localized service binding which is a common case in WSNs due to limited energy resources. Finally, Servilla provides a modular middleware architecture that can be configured for devices with a wider range of resources.

SOC is a topic of interest in the coordination community. For example, new languages have been developed that enable formal reasoning about complex service interactions and compositions [8, 1, 38, 3]. Calculi have been developed to model sessions and multi-party dynamic interactions between service users and providers [39, 13]. New ways of specifying quality-of-service requirements and achieving higher levels of reliability have been proposed [4, 10, 12, 42, 9]. SOC has also been used in non-traditional environments like mobile ad hoc networks [25]. Recently, there has been increased interest in context-aware applications [43, 20, 16]. WSNs, being embedded and able to sense the environment, are inherently context-aware. This paper takes the natural next step of applying SOC principles to WSNs. The key distinguishing feature of Servilla lies in its capability to support both resource-constrained devices and more powerful devices, and its light-weight language and middleware tailored for in-network coordination among sensors.

Efforts to bring SOC technologies into the WSN domain include Tiny Web Services [47] and Arch Rock's PhyNetTM [6]. Both optimize existing Internet protocols

to function under the severe resource constraints of WSNs. Unlike Servilla, they do not provide a mechanism for service discovery or the flexible matching of service consumers to providers *within* the WSN. Instead, they enable language-independent communication between services inside the WSN and applications outside of the WSN.

In addition to SOC, Servilla shares the common approach of using scripts in a WSN, though for different reasons. Some scripting systems, including Maté [34], ASVM [35], SwissQM [40], and Agilla [19], enable reprogramming. Other systems, including Melete [51] and SensorWare [11], enable multiple applications to share a WSN. All of these systems come with different scripting languages [33, 23, 37, 50]. Servilla differs by focusing on challenges due to network heterogeneity and dynamics. Unlike other systems, Servilla allows scripts to remain platform-independent and dynamically find and access platform-specific services. One scripting system, DVM [7], explores the similar idea of integrating platform-independent scripts with native services. It features a dynamically extensible virtual machine in which services can register extensions. While this enables tuning the boundary between interpreted and native code, DVM does not support flexible matching between scripts and services.

Servilla introduces the idea of a modular and configurable platform in which extremely resource-poor devices only implement a fraction of the entire framework. This enables a hierarchy in which weak devices serve more powerful devices. The idea of having a hierarchy within a WSN is promoted by other systems. Tenet [22] creates a two-tiered WSN in which the lower tier consists of resource-poor devices that can accept tasks from higher-tier devices. It differs from Servilla in that it does not support service discovery and dynamic binding between different devices. SONGS [36] is an architecture for WSNs that allows users to issue queries that are automatically decomposed into graphs of services which are mapped onto actual devices. SONG does not provide service binding among heterogeneous devices.

8 Conclusions

The increasing difficulty of developing applications for heterogeneous and dynamic WSNs demands a new coordination model. Servilla provides this by introducing a novel service provisioning framework that enables applications to be platform-independent while still able to access platform-specific capabilities. A salient feature of Servilla lies in its capability to support coordination and collaboration among heterogeneous devices *inside* a WSN. A specialized service description language is introduced that enables flexible matching between applications and services, which may reside on different devices. Servilla provides a modular middleware architecture to enable resource-poor devices to participate by contributing services, facilitating in-network collaboration among a wide range of devices. The efficiency of Servilla's implementation is established via microbenchmarks on two representative classes of hardware platforms. The effectiveness of Servilla's programming model is demonstrated by a structural health monitoring application case study.

Acknowledgment

This work is funded by the National Science Foundation under grants CNS-0520220, CNS-0627126, and CNS-0708460.

References

1. ABREU, J., AND FIADEIRO, J. L. A coordination model for service-oriented interactions. In Lea and Zavattaro [32], pp. 1–16.
2. ALONSO, G., CASATI, F., KUNO, H., AND MACHIRAJU, V. *Web Services*. Springer, 2003.
3. ANKOLEKAR, A., HUCH, F., AND SYCARA, K. P. Concurrent semantics for the web services specification language daml-s. In Arbab and Talcott [5], pp. 14–21.
4. ARBAB, F., CHOTHIA, T., MENG, S., AND MOON, Y.-J. Component connectors with qos guarantees. In Murphy and Vitek [41], pp. 286–304.
5. ARBAB, F., AND TALCOTT, C. L., Eds. *Coordination Models and Languages, 5th International Conference, COORDINATION 2002, YORK, UK, April 8-11, 2002, Proceedings (2002)*, vol. 2315 of *Lecture Notes in Computer Science*, Springer.
6. ARCH ROCK. Arch Rock PhyNet™. <http://www.archrock.com/product/>.
7. BALANI, R., HAN, C.-C., RENGASWAMY, R. K., TSIGKOGIANNIS, I., AND SRIVASTAVA, M. Multi-level software reconfiguration for sensor networks. In *EMSOFT '06: Proceedings of the 6th ACM & IEEE International conference on Embedded software* (New York, NY, USA, 2006), ACM, pp. 112–121.
8. BETTINI, L., NICOLA, R. D., AND LORETI, M. Implementing session centered calculi. In Lea and Zavattaro [32], pp. 17–32.
9. BOCCHI, L., CIANCARINI, P., AND ROSSI, D. Transactional aspects in semantic based discovery of services. In Jacquet and Picco [29], pp. 283–297.
10. BOCCHI, L., AND LUCCHI, R. Atomic commit and negotiation in service oriented computing. In Ciancarini and Wiklicky [14], pp. 16–27.
11. BOULIS, A., HAN, C.-C., AND SRIVASTAVA, M. B. Design and implementation of a framework for efficient and programmable sensor networks. In *MobiSys '03: Proceedings of the 1st international conference on Mobile systems, applications and services* (New York, NY, USA, 2003), ACM, pp. 187–200.
12. BRAVETTI, M., AND ZAVATTARO, G. A theory for strong service compliance. In Murphy and Vitek [41], pp. 96–112.
13. BRUNI, R., LANESE, I., MELGRATTI, H. C., AND TUOSTO, E. Multiparty sessions in soc. In Lea and Zavattaro [32], pp. 67–82.
14. CIANCARINI, P., AND WIKLICKY, H., Eds. *Coordination Models and Languages, 8th International Conference, COORDINATION 2006, Bologna, Italy, June 14-16, 2006, Proceedings (2006)*, vol. 4038 of *Lecture Notes in Computer Science*, Springer.
15. CROSSBOW TECHNOLOGIES. Imote2 datasheet. <http://tinyurl.com/5jrw85>.
16. CUBO, J., SALAÜN, G., CÁMARA, J., CANAL, C., AND PIMENTEL, E. Context-based adaptation of component behavioural interfaces. In Murphy and Vitek [41], pp. 305–323.
17. CULLER, D., ESTRIN, D., AND SRIVASTAVA, M. Overview of sensor networks. *IEEE Computer* 37, 8 (2004), 41–49.
18. FLANAGAN, D. *JavaScript: The Definitive Guide, 4th Ed.* O'REILLY, Inc., 2001.
19. FOK, C.-L., ROMAN, G.-C., AND LU, C. Rapid development and flexible deployment of adaptive wireless sensor network applications. In *ICDCS '05: Proceedings of the 25th IEEE International Conference on Distributed Computing Systems* (Washington, DC, USA, 2005), IEEE Computer Society, pp. 653–662.

20. FREY, D., AND ROMAN, G.-C. Context-aware publish subscribe in mobile ad hoc networks. In Murphy and Vitek [41], pp. 37–55.
21. GAY, D., LEVIS, P., VON BEHREN, R., WELSH, M., BREWER, E., AND CULLER, D. The nesc language: A holistic approach to networked embedded systems. In *PLDI '03: Proceedings of the ACM SIGPLAN 2003 conference on Programming language design and implementation* (New York, NY, USA, 2003), ACM, pp. 1–11.
22. GNAWALI, O., JANG, K.-Y., PAK, J., VIEIRA, M., GOVINDAN, R., GREENSTEIN, B., JOKI, A., ESTRIN, D., AND KOHLER, E. The tenet architecture for tiered sensor networks. In *SenSys '06: Proceedings of the 4th international conference on Embedded networked sensor systems* (New York, NY, USA, 2006), ACM, pp. 153–166.
23. GREENSTEIN, B., KOHLER, E., AND ESTRIN, D. A sensor network application construction kit (snack). In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems* (New York, NY, USA, 2004), ACM, pp. 69–80.
24. HACKMANN, G., SUN, F., CASTANEDA, N., LU, C., AND DYKE, S. A holistic approach to decentralized structural damage localization using wireless sensor networks. In *RTSS '08: Proceedings of the 2008 Real-Time Systems Symposium* (Washington, DC, USA, 2008), IEEE Computer Society, pp. 35–46.
25. HANDOREAN, R., AND ROMAN, G.-C. Service provision in ad hoc networks. In Arbab and Talcott [5], pp. 207–219.
26. HE, T., KRISHNAMURTHY, S., LUO, L., YAN, T., GU, L., STOLERU, R., ZHOU, G., CAO, Q., VICAIRE, P., STANKOVIC, J. A., ABDELZAHER, T. F., HUI, J., AND KROGH, B. Vigilnet: An integrated sensor network system for energy-efficient surveillance. *ACM Trans. Sen. Netw.* 2, 1 (2006), 1–38.
27. HILL, J., SZEWCZYK, R., WOO, A., HOLLAR, S., CULLER, D., AND PISTER, K. System architecture directions for networked sensors. In *Architectural Support for Programming Languages and Operating Systems* (2000), pp. 93–104.
28. INTANAGONWIWAT, C., GOVINDAN, R., AND ESTRIN, D. Directed diffusion: a scalable and robust communication paradigm for sensor networks. In *MobiCom '00: Proceedings of the 6th annual international conference on Mobile computing and networking* (New York, NY, USA, 2000), ACM, pp. 56–67.
29. JACQUET, J.-M., AND PICCO, G. P., Eds. *Coordination Models and Languages, 7th International Conference, COORDINATION 2005, Namur, Belgium, April 20-23, 2005, Proceedings* (2005), vol. 3454 of *Lecture Notes in Computer Science*, Springer.
30. KEMPF, J., AND PIERRE, P. S. *Service location protocol for enterprise networks: implementing and deploying a dynamic service finder*. John Wiley & Sons, Inc., New York, NY, USA, 1999.
31. KUMARAN, I., AND KUMARAN, S. I. *Jini Technology: An Overview*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 2001.
32. LEA, D., AND ZAVATTARO, G., Eds. *Coordination Models and Languages, 10th International Conference, COORDINATION 2008, Oslo, Norway, June 4-6, 2008. Proceedings* (2008), vol. 5052 of *Lecture Notes in Computer Science*, Springer.
33. LEVIS, P. The TinyScript Manual. <http://tinyurl.com/57kycj>, July 2004.
34. LEVIS, P., AND CULLER, D. Maté: a tiny virtual machine for sensor networks. In *ASPLOS-X: Proceedings of the 10th international conference on Architectural support for programming languages and operating systems* (New York, NY, USA, 2002), ACM, pp. 85–95.
35. LEVIS, P., GAY, D., AND CULLER, D. Active sensor networks. In *NSDI'05: Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation* (Berkeley, CA, USA, 2005), USENIX Association, pp. 343–356.
36. LIU, J., AND ZHAO, F. Towards semantic services for sensor-rich information systems. In *2nd Int. Conf. on Broadband Networks* (2005), pp. 44–51.

37. MADDEN, S., FRANKLIN, M. J., HELLERSTEIN, J. M., AND HONG, W. Tag: a tiny aggregation service for ad-hoc sensor networks. *SIGOPS Oper. Syst. Rev.* 36, SI (2002), 131–146.
38. MAZZARA, M., AND GOVONI, S. A case study of web services orchestration. In Jacquet and Picco [29], pp. 1–16.
39. MEZZINA, L. G. How to infer finite session types in a calculus of services and sessions. In Lea and Zavattaro [32], pp. 216–231.
40. MÜLLER, R., ALONSO, G., AND KOSSMANN, D. A virtual machine for sensor networks. *SIGOPS Oper. Syst. Rev.* 41, 3 (2007), 145–158.
41. MURPHY, A. L., AND VITEK, J., Eds. *Coordination Models and Languages, 9th International Conference, COORDINATION 2007, Paphos, Cyprus, June 6-8, 2007, Proceedings* (2007), vol. 4467 of *Lecture Notes in Computer Science*, Springer.
42. NORES, M. L., DUQUE, J. G., AND ARIAS, J. J. P. Managing ad-hoc networks through the formal specification of service requirements. In Ciancarini and Wiklicky [14], pp. 164–178.
43. NÚÑEZ, A., AND NOYÉ, J. An event-based coordination model for context-aware applications. In Lea and Zavattaro [32], pp. 232–248.
44. OSGI. Open source gateway initiative. <http://www.osgi.org>.
45. PAPAZOGLU, M. P., TRAVERSO, P., DUSTDAR, S., AND LEYMANN, F. Service-oriented computing: State of the art and research challenges. *Computer* 40, 11 (2007), 38–45.
46. POLASTRE, J., SZEWCZYK, R., AND CULLER, D. Telos: enabling ultra-low power wireless research. In *IPSN '05: Proceedings of the 4th international symposium on Information processing in sensor networks* (Piscataway, NJ, USA, 2005), IEEE Press, p. 48.
47. PRIYANTHA, N. B., KANSAL, A., GORACZKO, M., AND ZHAO, F. Tiny web services: design and implementation of interoperable and evolvable sensor networks. In *SenSys '08: Proceedings of the 6th ACM conference on Embedded network sensor systems* (New York, NY, USA, 2008), ACM, pp. 253–266.
48. STREELINE. Parking management. <http://www.streetlinenetworks.com>.
49. W3C. Web services description language (wsdl). <http://www.w3.org/TR/wsdl>.
50. YAO, Y., AND GEHRKE, J. The cougar approach to in-network query processing in sensor networks. *SIGMOD Rec.* 31, 3 (2002), 9–18.
51. YU, Y., RITTLE, L. J., BHANDARI, V., AND LEBRUN, J. B. Supporting concurrent applications in wireless sensor networks. In *SenSys '06: Proceedings of the 4th international conference on Embedded networked sensor systems* (New York, NY, USA, 2006), ACM, pp. 139–152.