

An energy-saving model for service-oriented mobile application development

Ignacio Lizarralde (ISISTAN - CONICET), Cristian Mateos (ISISTAN - CONICET),
Alejandro Zunino (ISISTAN - CONICET)

ISISTAN Research Institute. UNICEN University. Campus Universitario, Tandil (B7001BBO),
Buenos Aires, Argentina. Tel.: +54 (249) 4439682 and Consejo Nacional de Investigaciones
Científicas y Técnicas (CONICET)

Keywords: Smartphones, SOC, Web Services, Energy save, Android

Abstract. The development of mobile applications that combine Web Services from different providers --also referred as mashup applications-- is growing as a consequence of the ubiquity of bandwidth connections and the increasing number of available Web Services. In this context, providing higher maintainability to Web Service applications is a worth of matter, because of the dynamic nature of the Web. EasySOC (1) solves this problem by decoupling mashups from application components. However, mobile devices have energy constraints because of the limitations in the current battery capacities. This work proposes a model that builds on the benefits of the EasySOC approach and improves this latter by assisting developers to select Web Service combinations that reduce energy consumption. We evaluated the feasibility of the model through a case study in which we compare the estimations provided by the model against real energy measurements. The results indicated that our model had an efficacy of 81% for the analyzed case study.

1 Introduction

The number of mobile devices is expected to exceed the world population this year (7.3 billion). As a consequence, mobile data traffic last year was nearly 18 exabytes. These facts are a consequence of the growing popularity of mobile devices (1) (2). They evolved from offering limited functional capabilities such as calendars or calculators to providing the same functionality as a small computer (3). Although these devices have enough power to be used in distributed environments (1), in general battery-life is negatively affected when resource usage grows because battery capacity has not increased accordingly to the energy consumption of recent mobile hardware (4). Thus, while 10 years ago battery duration was about one week, today is at most 2 days (4).

On the other hand, nowadays, mobile devices make intensive use of networking interfaces such as Wi-Fi, 3G and HSDPA for accessing data and resources on the Internet (5) (6). However, as mentioned earlier, battery-life is a main concern in these devices, and networking interfaces highly affects energy consumption (7) (8) (9). Thus, these concerns enforce the importance of energy consumption reduction in mobile devices.

In the last years, Service-oriented Computing (SOC) (10) started to become popular due to the ubiquity of high speed network connections. SOC is a computing paradigm that promotes the composition of external software components, called services, available and invocable through the Internet. In this context, applications are developed using services as the basic building blocks, decreasing the cost and maintainability of the development process. Additionally, development of SOC applications that combine services in mobile devices has been made attractive for developers because of the widespread availability of Web Services --the most common technological materialization of SOC-- and the popularity of social networks (i.e., Facebook, Twitter, etc) and cloud services. These applications, which rely on a combination of Web Services from different sources are called mashup applications. For example, SongDNA allows users to search for individual songs based on data collected from over 10 sources and obtain song lyrics, artist bio, links to videos and even information about what people are saying about the song in Twitter.

Broadly, despite their benefits, employing Web Services to develop mobile applications has some shortcomings because Web Service invocation is a resource demanding process. Moreover, these devices have limited battery-life and processing capabilities. In this paper, we propose an energy-aware SOC development model, which builds on the benefits of EasySOC (11), an approach to develop SOC applications that focus on maintainability by decoupling client application components from service interfaces. EasySOC isolates the user-provided (or in-house) application components at the client side from the interface(s) exposed by the provider(s) by means of service adapters which transform the provider interface into the expected client interface. Moreover, our model improves EasySOC by reducing energy consumption due to interface adaptations from applications by allowing users to select the least energy expensive service combinations. In other words, our model focus on adapters energy consumption. Additionally, because of the popularity and wide spread adoption of the Android platform, this work focuses on Android-powered devices.

The rest of this paper is organized as follows. First, in Section 2 we describe EasySOC and present the proposed estimation model. Then, in Section 3 we describe the experimental evaluations performed to asses the model. Finally, in Section 4 and Section 5 we describe the most relevant related work and the conclusions, respectively.

2 An energy-save estimation model for EasySOC

SOC promotes the reuse of software components available from the Internet, called services. Moreover, Web Services materialize SOC by prescribing languages to specify a public interface to describe the functions --operations-- and parameters required to invoke a service. Then, if a developer wants to build a mashup application he/she needs to select which available services will meet their functional requirements by inspecting their public interfaces. This ties the client application components with the interfaces of the selected services. The approach to SOC development described in (11), EasySOC, solves this problem by decoupling mashups from services and hence providing higher maintainability levels to SOC applications. The approach combines DI (Dependency Injection) and the Adapter pattern to isolate the user-provided application components at the client side from the interface(s) exposed by the

provider(s). In this context, the approach requires to firstly define an expected interface for external services at design time. Then, services adapters carry the responsibility for transforming these interfaces into the interfaces exposed by the available service providers.

The model we will present builds on the benefits of the EasySOC approach, but at the same time improves this latter by providing energy estimations to EasySOC applications. This allows users to select the least energy expensive service combination (due to interface adaptations), thus reducing the wasted energy during the invocation process and hence mashup usage. However, this work does not provide a mechanism for service selection, instead, it provides a model that will help developers in service selection by providing energy consumption estimations.

Section 2.1 describes EasySOC (11) (12). Later, in Section 2.2 we will introduce the concept of micro-operation and how it relates to the model. Finally, Section 2.3 describes the model.

2.1 EasySOC

Although SOC simplifies application development through component composition, application maintainability may be affected. SOC applications rely on services that are invoked through an interface exposed by a service provider. However, service interfaces could change affecting client applications, since some of their components need to be adapted in order to meet the new interfaces. In this section we describe EasySOC (11) (12), an approach that aims to isolate client application components from service interfaces through service adapters. The energy consumption of these service adapters is the focus of the model we will present in the following sections.

EasySOC builds on the idea that an interesting implication of using DI in SOC is that client-side application logic can be isolated from the details for invoking services (e.g., URLs, namespaces, port names, protocols, etc.). With this in mind, a developer thinks of a Web Service as any other regular component providing a clear interface to its operations. If a developer wants to call an external Web Service S with interface I_S from within an in-house component C , a dependency between C and S is established through I_S . This kind of dependency is commonly managed by a DI container that injects a proxy to S (let us say P_S) into C . At runtime, the code of C will end up calling any of the methods declared in I_S through P_S , which transparently invokes the remote service. Interestingly, this mechanism is not intrusive, since it only requires to associate a configuration file with the client application, which is used by the DI container to determine which components should be injected into other ones.

Although DI provides a fitting alternative to cleanly incorporate a Web Service into a client application, it leads to a form of coupling through which the application is tied to the invoked service interfaces (i.e., I_S). In this way, changing the provider for a service requires to adapt the client application to follow the new interface. This is illustrated from an architectural perspective in the left side of Fig. 1, in which the layer containing the business logic depends on particular proxies (grey layers are coupled to the underlying ones). At the implementation level, this means to rewrite the portions of the application that depend on I_S , which includes operation signatures that are likely to differ from that of the new interface.

To overcome this problem, EasySOC combines the DI and Adapter patterns to introduce an intermediate layer that allows developers to seamlessly shift between different interfaces. Conceptually, instead of directly injecting a layer of service proxies (P_s) into the application, which requires modifying the layer containing the client business logic to make it compatible with the service interfaces (I_s), EasySOC injects a layer of service adapters (see the right side of Fig. 1). To evidence the benefits of applying this approach, suppose a client application that uses a weather service to obtain information about the temperature. If the service provider shuts down the service permanently, it is necessary to change the service provider. This means that the client application components invoking the old service must adapt to the new service interface. On the other hand, for an EasySOC application only the adapter of the old service needs to be changed, leaving intact the application components.

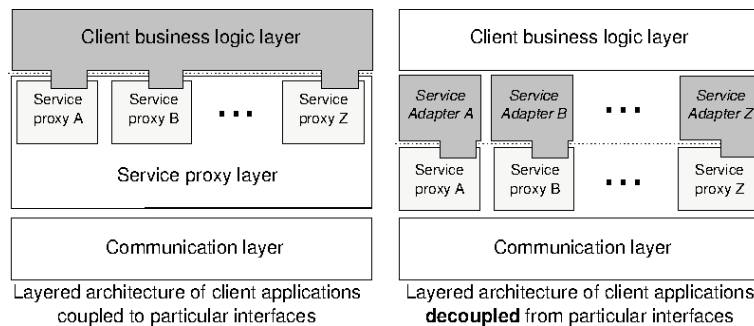


Fig. 1 Not using (left) and using (right) EasySOC in client applications: Architectural differences.

A service adapter is a specialized Web Service proxy, based on the Adapter pattern, which adapts the interface of a particular service according to the interface (specified by the developer at design time) expected by the in-house components. We refer to A_{sc} as an adapter that accommodates the actual interface of a service S to the interface expected by an in-house component C . In other words, an adapter carries the necessary logic to transform the operation signatures of the expected client interface to the actual interface of a selected Web Service. For instance, if a service operation returns a list of integers, but the application expects an array of floats, a service adapter would perform the type conversion.

2.2 Micro-operations

In the previous section we have described EasySOC, which isolates the client logic from the service interface exposed by providers. In this context, each adapter performs a set of operations (from now on referred as micro-operations) to transform the signature of the expected client interface to the interface expected by the provider.

As mentioned earlier, this work aims to create a model for estimating the energy consumption of Web Service interface adaptations. To achieve this, we followed three steps.

First, we determined which micro-operations are involved in the adapters commonly built. Second, we measured the energy consumption of each micro-

operation. Third, we applied the model explained in Section 2.3 to estimate the overall energy consumption. This section covers the first step.

Firstly, we created a catalog containing the most common operations involved in interface adaptations. Micro-operations were selected as a subset of the Java Grande6 benchmark suite, similarly to the studies presented in (3). Table 1 shows this catalog along with the Java equivalent code of each micro-operation. There are two groups of microoperations, one group containing primitive operations (i.e., cast int to float, cast int to double) and another group involving complex operations (i.e., create an Object, create an int array). In order to show the applicability of micro-operations for describing service adapters, let us suppose a reverse-geocoding service:

- Client (expected interface): String getAddress(int latitude, int longitude).
- Provider (real interface): String getAddress(GPSMessage_Type2 message).

Fig. 3 and Fig. 4 shows the structure of the GPSMessage_Type2 complex type and the Java code of the adapter responsible for transforming the client expected interface into the actual service interface, respectively. The resulting description of the necessary adaptation in terms of micro-operations is as follows:

Micro-operation name	Equivalent Java code
MOP1: cast int to double	<code>int a; double b; b = (double) a; o a = (int) b;</code>
MOP2: cast int to float	<code>int a; float b; b = (float) a; o a = (int) b;</code>
MOP3: cast long to double	<code>long a; double b; b = (double) a; o a = (long) b;</code>
MOP4: cast long to float	<code>long a; float b; b = (float) a; o a = (long) b;</code>
MOP5: create an Object instance	<code>Object a = new Object();</code>
MOP6: assign from other instance	<code>public class A { public int a; public int b; } public class B { public void foo (){ A aInstance = new A (); a Instance.a = aInstance.b; } }</code>
MOP7: assign from the same instance	<code>public class A{ private int a; private int b; public foo (){ a = b ; } }</code>
MOP8: create a float array[n]	<code>float [n] a = new float [n];</code>
MOP9: create a int array[n]	<code>int [n] a = new int [n];</code>
MOP10: create a long array[n]	<code>long [n] a = new long [n];</code>
MOP11: create a Object array[n]	<code>Object [n] a = new Object [n];</code>

Table 1 Micro-operations Catalog

- 1 occurrence of MOP5: GPSMessageType2, line 11.
- 2 occurrences of MOP1: (double) longitude and (double) latitude, lines 12 and 13.
- 2 occurrences of MOP6: setLongitude(longitude) and setLatitude(latitude), lines 12 and 13.

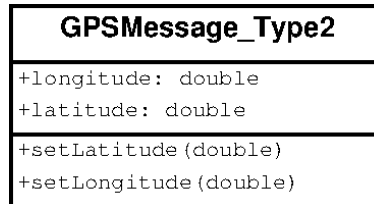


Fig. 2 A complex type of the expected interface

```

1. public class GetAddress_V2_MultipleAdapter implements
   IGetAddressAdapter {
2.     private ServiceMultipleConsumerProxy proxy;
3.     public GetAddress_V2_MultipleAdapter
   (ServiceMultipleConsumerProxy proxy) {
4.         this.proxy = proxy;
5.     }
6.     public String getAddress (int latitude, int longitude) {
7.         GPSMessage_Type2 msg = new GPSMessage_Type2 ( );
8.         msg.setLatitude((double) latitude);
9.         msg.setLongitude ((double) longitude);
10.        return proxy.getAddress (msg);
11.    }
12. }

```

Fig. 3 Example: Adapter Implementation

To conclude, in this section we have presented a catalog of micro-operations for describing data transformations and how to use this to represent an adapter in terms of such elemental adaptation operations. The estimation model covered in the next section uses these descriptions along with the energy consumption of each micro-operation as an input to estimate the overall energy consumption of an adapter. Also, we will expand this idea to estimate the energy consumption of mashups.

2.3 Estimation model

The model proposed in this section helps to estimate the least energy expensive combination of services when building a mashup. In order to achieve this, the model relies in the concept of micro-operation we have described earlier.

Formally, a Web Service mashup is a combination of one or more services $\{S_1..S_n\}$ from different sources. The EasySOC approach establishes that every service has associated an adapter that carries the necessary logic to transform the signatures of expected interfaces into the actual service interfaces. Formally, for each service $\{S_1..S_n\}$ there is an adapter $\{A_1..A_n\}$ each one having an energy cost, namely $\{C_1..C_n\}$. Additionally, an adapter can be described in terms of micro-operations $\{MO_1..MO_p\}$. Then, the estimated energy cost of a service adapter can be calculated as $C_k = \sum_{i=1}^p CMO_i$, where CMO_i is the energy cost of the corresponding MO_i micro-operation. Finally, the overall estimated energy cost of a mashup k (M_k) is $CME_k = \sum_{i=1}^n C_i$.

To illustrate the model, suppose a developer who wants to select the least energy expensive combination of services to get information about songs (e.g., SongDNA). The developer wants to get links to the song videos, and the lyrics of the song. There are four services available, two for either category, namely $\{S_{11}, S_{12}\}$ and $\{S_{21}, S_{22}\}$. Then, each service has an associated adapter, which can be described in terms of micro-operations:

- $A_{11} = \{MO_1 = 1, MO_2 = 5\}$,
- $A_{12} = \{MO_1 = 3, MO_2 = 2\}$,
- $A_{21} = \{MO_1 = 7\}$,
- $A_{22} = \{MO_1 = 1, MO_2 = 1, MO_3 = 1\}$.

To determine which alternative fits the energy requirements of the developer, we apply the model by adding the energy cost of the micro-operations in order to calculate the estimated energy cost of invoking each service:

- $\{S_{11} = 1 + 5 = 6\}$, $\{S_{12} = 3 + 2 = 5\}$ – $\{S_{21} = 7\}$, $\{S_{22} = 1 + 1 + 1 = 3\}$.

Finally, we calculate the estimated energy cost of each mashup alternative:

- $M_1 = \{S_{11} = 6, S_{21} = 7\} = 13$,
- $M_2 = \{S_{12} = 5, S_{21} = 7\} = 12$,
- $M_3 = \{S_{11} = 6, S_{22} = 3\} = 9$,
- $M_4 = \{S_{12} = 5, S_{22} = 3\} = 8$.

Then, the least energy consuming combination is $M_4 = \{S_{12}, S_{22}\}$.

In order to evaluate the feasibility of the proposed model, we considered a set $B_k = \{M_1..M_j\}$ composed by the alternatives that have lower energy consumption than M_k . We created two M_k groups, one corresponding to the estimated energy cost CME and another corresponding to the real energy cost CMR which we have measured from experiments. Then, we compared these groups to find errors between the estimated and the real cases. Next section covers these experiments, and uses this approach to evaluate the estimations provided by the model.

3 Experimental evaluation

This section describes the results of the measurements for energy consumption we made for micro-operations along with a case study in which we applied and evaluated the proposed estimation model.

Before introducing the experiments we will describe the measurement environment. We used a Power Monitor, a tool capable of measuring the power of any device that employs a lithium battery at a frequency of 5Khz (5,000 times per second). This tool comes along with a software application that shows the Voltage and Amperage measurements, and allows the user to export this information in different formats (e.g., CSV). Fig. 4 shows the device setup.



Fig. 4 Power Monitor Setup

Additionally, the mobile device used for testing was a Samsung I5500 with the following characteristics: 600 MHz CPU (model MSM7227-1 ARM11), 256 MB RAM, internal 170 MB uSD 4GB of storage and lithium ion battery of 1,200 mAh. The experiments were made after a device restart and executing only the essential operating system services.

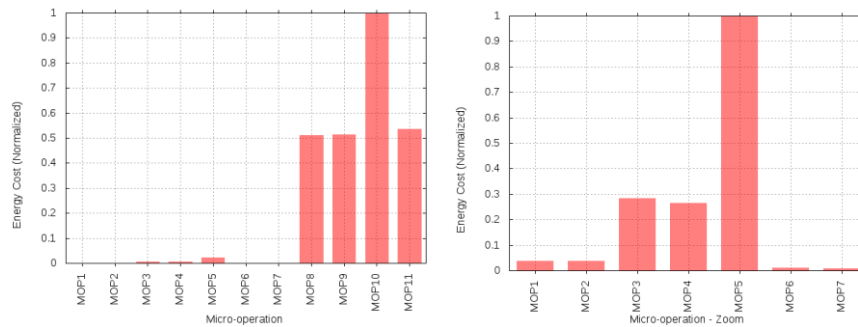
3.1 Micro-operations

Earlier, we have presented the concept of micro-operation and how use this to describe Web Service adapters. Also, we have mentioned that a necessary step before applying the model for estimating the overall consumption of a service adapter, is to measure the energy consumption of each micro-operation.

Table 2 shows the average energy consumption of several test runs for each micro-operation. Note that array-related micro-operations had a fixed size. To ensure representativeness, we simulated the creation of arrays of size N . In order to achieve this, we increased the array size in one element every run ($N = 1$ to $N=64,000$ for array-related micro-operations). Then, we calculated the average power and time.

Micro-operation	Time (s)	% Standard Deviation - Time	Power (Watts)	% Standard Deviation - Power	Power consumption (per unit) (μ Joules)
MOP1: cast int to double	0.9080	2.2764	0.6175	0.9717	0.0876
MOP2: cast int to float	0.8675	1.8100	0.6147	1.2931	0.0833
MOP3: cast long to double	6.4446	0.3918	0.6399	0.1818	0.6443
MOP4: cast long to float	5.9559	0.3915	0.6454	0.2372	0.6006
MOP5: create an Object instance	22.77	0.6202	0.6374	0.1080	2.2680
MOP6: assign from other instance	2.0355	0.9261	0.6779	2.1781	0.0216
MOP7: assign from the same instance	2.0451	1.0538	0.6236	3.0924	0.0199

MOP8: create a float array[n]	5.0734	0.8660	0.6270	0.1791	49.7047
MOP9: create a int array[n]	5.1084	0.9593	0.6256	0.2615	49.9425
MOP10: create a long array[n]	9.9019	1.0560	0.6290	0.1347	97.3300
MOP11: create a Object array[n]	5.3399	2.3442	0.6250	0.2291	52.1544

Table 2 Energy consumption results: Micro-operations**Fig. 5** Energy consumption comparison: Micro-operations

Error! No se encuentra el origen de la referencia. compares micro-operations energy consumption measured in Joules --the work required to produce one watt of power for one second--. The results show that array creation micro-operations consume more energy than any other micro-operation (about 4,800% in the worst case). Also, long array creation is about 100% more energy-hungry than any other array creation micro-operation. This is since the memory size of the long type (64-bits in 32-bit systems) is twice the size of any of the other array creation micro-operations. This means that long arrays micro-operations require twice as much memory for the same array size.

Service	Expected Interface
Reverse Geocoding	String getAddress(int latitude, int longitude)
Send an E-mail	sendMessage(String text, String subject, String[] recipients)
Send an SMS	sendSMSMessage(String text, String[] recipients)

Table 3 Client expected interfaces

3.2 Case Study

In this section we put in practice the above results to illustrate the utility and feasibility of the proposed estimation model. First, we will present a case study involving the development of a mashup, then we will apply the model to select services, and finally we will compare how accurate the estimations provided by the model are compared with the real energy consumption of the possible service combinations for building the mashup.

Service		Adapter micro-operations	Estimated cost (μ Joules)
Service name	Version		
1: Get a direction	1	<ul style="list-style-type: none"> • MOP5: 1 • MOP2: 2 • MOP6: 2 	2.4778
	2	<ul style="list-style-type: none"> • MOP1: 2 • MOP5: 1 • MOP6: 2 	2.4864
2: Send an E-mail	1	<ul style="list-style-type: none"> • MOP5: 1 • MOP6: 52 • MOP9: 1 	55.5456
	2	<ul style="list-style-type: none"> • MOP5: 1 • MOP9: 2 • MOP6: 50 	103.2330
3: Send an SMS	1	<ul style="list-style-type: none"> • MOP5: 1 • MOP6: 52 • MOP9: 1 	55.5456
	2	<ul style="list-style-type: none"> • MOP5: 1 • MOP10: 1 • MOP6: 52 	152.8756

Table 4 Case study: Adapter estimation results

Consider the three services and client-side expected interfaces from Table 3. We created eight mashup versions through the combination of these services, varying the type and number of parameters. Table 4 shows the micro-operations involved to transform the expected interfaces into the real service interfaces. However, in some cases, the micro-operations proposed cannot describe exactly the operations involved in the adapter. For example, the instantiation of user-defined types. In such cases we used the most similar micro-operations, considering creations of user-defined types as creations of Object types, for instance. Also, by basing on the results presented in Table 2, we applied the model and calculated the estimated cost for each service version. The size of the arrays in parameters were set to 50 for all services.

We measured the real energy consumption of each mashup combination. In order to achieve this, first, we implemented the services presented in Table 4 using Axis2 and the service clients using ksoap2 (SOAP services). Then, we used Apache Tomcat7 to host the services, running in a server with an Athlon2 X2@2.6Ghz processor, 4Gb RAM, and running Windows 7. The mobile device hosting the client application contacted the server through a Wi-Fi connection to a private access point. In addition, service invocation were made in 10 groups of 1,000 invocations each (10,000 invocations). Also, for each possible service combination, 5 runs were made (a deviation below 4.2% was obtained).

Table 5 presents the measurement results. We used the following nomenclature to describe a mashup: (GetDirectionVersion, SendAnEmailVersion, SendAnSMSVersion). According to the results presented in Section 3.1, these results show that array creation micro-operations have a high impact on energy consumption. Also, combinations that involve creating long arrays are significantly more energy-demanding than any other.

Mashup	Time (s)	% Standard Deviation - Time	Power (Watts)	% Standard Deviation - Power	Energy consumption (per unit) (μ Joules)
M_1 : (1, 1, 1)	0.0068	4.1311	0.7344	1.0593	5040.41
M_2 : (1, 1, 2)	0.0084	2.8573	0.7127	1.1205	6056.04
M_3 : (1, 2, 1)	0.0079	2.9796	0.7275	1.0218	5817.13
M_4 : (1, 2, 2)	0.0095	2.3428	0.7059	0.9156	6741.06
M_5 : (2, 1, 1)	0.0069	1.7495	0.7223	0.3350	4995.10
M_6 : (2, 1, 2)	0.0079	2.6022	0.7192	1.3780	5740.82
M_7 : (2, 2, 1)	0.0081	2.3029	0.7179	0.8427	5847.66
M_8 : (2, 2, 2)	0.0092	2.8843	0.7156	0.8468	6651.41

Table 5 Energy consumption results: Mashups

Finally, we used the formulas presented in Section 2.3 and the results of the estimations from Table 4 to estimate which mashup combinations have the lowest energy consumption. Fig. 6 shows a comparison between the estimated and the real energy consumption results. Our model was not intended to provide the exact amount of energy each alternative will actually consume, but to provide energy consumption estimations to EasySOC applications, allowing the developer to select the least expensive mashup combination. Thus, even if there is a large difference between the real and the estimated cost in some cases, this should not affect the results. For example, the mashup (1, 1, 1) combination is less energy-expensive compared to the (1, 1, 2) combination, both in the real and the estimated cases. To evidence this, we used the approach proposed in Section 2.3. First, we created the B_k sets:

– Real:

- $B_1 = \{M_5\}$
- $B_2 = \{M_1, M_3, M_5, M_6, M_7\}$
- $B_3 = \{M_1, M_5, M_6\}$
- $B_4 = \{M_1, M_2, M_3, M_5, M_6, M_7, M_8\}$
- $B_5 = \{\varepsilon\}$
- $B_6 = \{M_1, M_5\}$
- $B_7 = \{M_1, M_3, M_5, M_6\}$ • $B_8 = \{M_1, M_2, M_3, M_5, M_6, M_7\}$

– Estimated:

- $B_1 = \{\varepsilon\}$
- $B_2 = \{M_1, M_3, M_5, M_7\}$
- $B_3 = \{M_1, M_5\}$
- $B_4 = \{M_1, M_2, M_3, M_5, M_6, M_7\}$
- $B_5 = \{M_1\}$
- $B_6 = \{M_1, M_2, M_3, M_5, M_7\}$
- $B_7 = \{M_1, M_3, M_5\}$

- $B_8 = \{M_1, M_2, M_3, M_4, M_5, M_6, M_7\}$

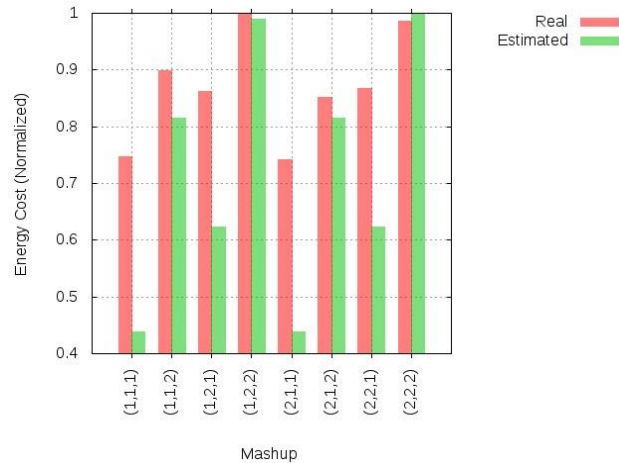


Fig. 6 Energy consumption comparison: Mashups

Then, we compared the differences in both cases. From 27 comparisons made, there were 5 errors (in red). This shows that the efficacy of the model, for this case study, is $\frac{22}{27} = 0.8148 \approx 81\%$.

4 Related work

Previous works in the area have studied 2G and 3G networks power consumption. The measurements described in (13) suggest that 2G networks reduce the energy consumption for SMS and voice services, while 3G networks should be the preferred option for network data transfers providing both speed and energy efficiency. Also, (14) analyzes energy consumption in 302.11 networks, particularly, during the association process. The authors claim that eliminating the ARP probe process from the dynamic addressing protocol could improve dynamic addressing power consumption by almost 400%.

Additionally, (15) proposes TailEnder, a network scheduler that lowers the energy consumption up to 40% in 3G networks by reordering network transfers in a mobile device. In the same direction, (16) improves Bit Torrent transfer scheduling, reducing power consumption by almost 50% with respect to conventional clients. Other similar studies have profiled applications network traffic and identified four major application energy leaks: misinterpretation of callback API semantics, poorly designed downloading schemes, repetitive downloads, and aggressive prefetching (17).

On the other hand, the implications of using SOC in mobile applications has been gained attention in the research community. (18) studies service invocation patterns claiming that in 3G networks, an effective service invocation scheduling can improve up to 50% the energy consumption. Other works analyzed proxy-based architectures for service invocation that reduce the intra-device CPU processing during the service invocation process, resulting in a reduction of energy consumption (19) (20) (21)

(22). Among the works in this line, however, there is no evidence of an approach for service-oriented mobile application development which considers energy consumption as a first-class citizen.

5 Conclusions

In this work we have presented a SOC development model for mobile devices that improves EasySOC by providing energy consumption estimations. This model helps to determine which services or mashup combinations are less energy-hungry, reducing the energy consumption during the service invocation process after the applications have been built and deployed.

First, we have described the EasySOC (11) (12) approach for developing SOC applications. Our power-aware enhanced model builds on the benefits of EasySOC and improves this by estimating the energy consumption due to interface adaptations. In order to achieve this, we have proposed to describe service adapters in terms of elemental operations or micro-operations. Then, to calculate the overall energy consumption involved in an adaptation, the energy cost of individual micro-operations resulting from describing the adapter are summed. This is in turn used to estimate which mashup combinations require less energy.

To assess our approach we measured the energy consumption of each individual micro-operation. We identified that array micro-operations are more energy-hungry than any others and that creation of larger data-types increment energy consumption. This suggests it could be a relationship between energy consumption and larger memory requirements. Second, using as an input the energy cost of the micro-operations provided by the measurements, we described a case study in which we applied the proposed model. Also, we have validated the model comparing the estimations against the real measurements. The results indicated that the model had an efficacy of 81% for the proposed case study.

This work can be extended in several directions. Firstly, we will generalize the proposed model by incorporating to the catalog micro-operations that help to describe more precisely service adapters, for example for considering array assignation and common programming language defined-types such as strings or char. Also we will study the trade offs between (partially) including vs. not including service adapters in applications, and how this impacts on energy consumption, maintainability, and performance. Indeed, a previous study in the context of EasySOC (1) shows the trade-off between performance and maintainability when not using adapters, but energy should be considered alongside these two attributes.

On the other hand, we will experiment with new device types (e.g., tablets) and smartphones to further generalize to the measurements results. On the other hand, we will extend WSQBE (23) by providing real energy estimations for invoking services. WSQBE is an approach that combines popular best practices for using external Web services with text-mining and machine learning techniques for facilitating service publication and discovery. Thus, when a client searches for a service by providing the expected interface, WSQBE will return a similar service in functional terms but at the same time the least energy hungry one according to the amount of energy that adaptations will consume. Finally, we will extend our model to support REST Web Services (24) (25), and we will analyze the trade-offs against SOAP services.

Acknowledgments

We acknowledge the financial support of ANPCyT through grants PAE-PICT-2007-2311 and PICT-2012-0045.

Bibliography

1. *EasySOC: Making web service outsourcing easier*. **Crasso, Marco, et al.** 0, 2014, Information Sciences. En prensa. ISSN: 0020-0255 DOI: 10.1016/j.ins.2010.01.013.
2. *Introducing Mobile Devices into Grid Systems: A Survey*. **Rodriguez, Juan Manuel, Zunino, Alejandro and Campo, Marcelo.** 1, s.l. : Inderscience Publishers, 2011, International Journal of Web and Grid Services, Vol. 7, pp. 1-40. ISSN: 1741-1106.
3. *Mobile device-aware refactorings for computational kernels*. **Rodriguez, Ana, Mateos, Cristian and Zunino., Alejandro.** 2012. 13th Argentine Symposium on Technology, 41st JAIIO.
4. *Are smartphones really useful for scientific computing?* **Rodriguez, Juan M., Mateos, Cristian and Zunino., Alejandro.** 2011, In F. V. Cipolla-Ficarra et al., editor, Advances in New Technologies, Interactive Interfaces and Communicability (ADNTIIC 2011), Lecture Notes in Computer Science, Vol. n/a, pp. 35-44.
5. *Environmental sizing of smartphone batteries*. **Flipsen, B., et al.** 2012. Electronics Goes Green 2012+ (EGG), 2012. pp. 1-9.
6. *RESTful Web Services improve the efficiency of data transfer of a whole-farm simulator accessed by Android smartphones*. **Arroqui, Mauricio, et al.** 2011, Computers and Electronics in Agriculture, Vol. 87, pp. 14-18.
7. *Who killed my battery?: analyzing mobile browser energy consumption*. **Thiagarajan, Narendran, et al.** 2012. Proceedings of the 21st international conference on World Wide Web.
8. *Dynamic power management using on demand paging for networked embedded systems*. **Agarwal, Yuvraj, Schurgers, Curt and Gupta, Rajesh.** 2005, ASP-DAC '05 Proceedings of the 2005 Asia and South Pacific Design Automation.
9. *Energy scavenging for mobile and wireless electronics*. **Paradiso, Joseph A. and Starner, Thad.** 2005, Pervasive Computing, IEEE.
10. *Experience with a low power wireless mobile computing platform*. **Raghunathan, Vijay, et al.** 2004, ISLPED '04 Proceedings of the 2004 international symposium on Low power electronics and design.
11. *Service-oriented Computing*. **Bichler, Martin and Lin, Kwei-Jay.** 3, s.l. : IEEE Computer Society, 2006, Computer, Vol. 39, pp. 99-101. ISSN: 0018-9162.
12. *Separation of concerns in service-oriented applications based on pervasive design patterns*. **Mateos, Cristian, et al.** 2010. SAC '10 Proceedings of the 2010 ACM Symposium on Applied Computing.
13. *On the Impact of 2G and 3G Network Usage for Mobile Phones Battery Life*. **Perrucci, Gian Paolo, et al.** 2009. Wireless Conference, 2009. EW 2009. European. pp. 255-259.
14. *Measuring mobile phone energy consumption for 802.11 wireless networking*. **Rice, Andrew and Hay, Simon.** 2010, Pervasive and Mobile Computing, Volume 6, Issue 6, Vol. 6, pp. 593-606.
15. *Energy Consumption in Mobile Phones: A Measurement Study and Implications for Network Applications*. **Balasubramanian, Niranjana,**

Balasubramanian, Aruna and Venkataramani, Arun. 2009. IMC '09 Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference. pp. 280-293.

16. *Bursty Content Sharing Mechanism for Energy-Limited Mobile Devices.* **Kelényi, Imre and Nurminen, Jukka K.** 2009. PM2HW2N '09 Proceedings of the 4th ACM workshop on Performance monitoring and measurement of heterogeneous wireless and wired networks.

17. *ADEL: an automatic detector of energy leaks for smartphone applications.* **Zhang, Lide, et al.** 2012. CODES+ISSS '12 Proceedings of the eighth IEEE/ACM/IFIP international conference on Hardware/software codesign and system synthesis.

18. *Invoking Web Services Based on Energy Consumption Models.* **Papageorgiou, Apostolos, et al.** 2012. Mobile Services (MS), 2012 IEEE First International Conference on.

19. *Mobile Service Platform: A Middleware for Nomadic Mobile Service Provisioning.* **van Halteren, Aart and Pawar, Pravin.** 2006. Wireless and Mobile Computing, Networking and Communications, 2006. (WiMob'2006). IEEE International Conference on.

20. *Architectures for Web services access from mobile devices.* **Rendon, et al.** 2005. Web Congress, 2005. LA-WEB 2005. Third Latin American.

21. *The Study on Mobile Phone-Oriented Application Integration Technology of Web Services.* **Luqun, Li, Minglu, Li and Xianguo, Cui.** 2004, Grid and Cooperative Computing, Vol. 3032, pp. 867-874.

22. *A Proxy-Based Architecture for Dynamic Discovery and Invocation of Web Services from Mobile Devices.* **Artail, Hassan, Fawaz, Kassem and Ghandour, Ali.** 2012, Services Computing, IEEE Transactions, Vol. 5, pp. 99-115.

23. *Combining Query-By-Example and Query Expansion for simplifying {W}eb {S}ervice discovery.* **Crasso, Marco, Zunino, Alejandro and Campo, Marcelo.** Amsterdam, The Netherlands : Springer Netherlands, 2011, Information Systems Frontiers, Vol. 13, pp. 407-428. ISSN: 1387-3326.

24. *Bridging the semantic Web and Web 2.0 with Representational State Transfer.* **Battle, Robert and Benson, Edward.** 2008, Web Semantics: Science, Services and Agents on the World Wide Web.

25. *Fielding, Roy Thomas.* *Architectural styles and the design of network-based software architectures.* University of California, Irvine. 2000. Ph.D. dissertation.