

# A Survey of Public Web Services

Su Myeon Kim  
KAIST, EECS Dept.  
Yusung-Gu Gusung-dong 373-1, Taejon  
KOREA  
+82-42-869-3586  
smkim@nclab.kaist.ac.kr

Marcel-Catalin Rosu  
IBM T.J. Watson Research Center  
19 Skyline Drive, NY  
USA  
+1-914-784-7242  
rosu@us.ibm.com

## ABSTRACT

This paper introduces a methodology to provide the first characterization of public Web Services in terms of their evolution, location, complexity, message size, and response time.

## Categories & Subject Descriptors: A.1

[Introductory and survey]: *Web Services, SOAP Traffic, Geographical Distribution*

**General Terms:** Measurement, Experimentation

**Keywords:** Web Services, SOAP, WSDL, UDDI Business Registry, Measurement, Web Services Traffic Characteristics.

## 1. Introduction

Enterprise IT infrastructures are currently migrating toward a service-oriented architecture, using Web Services (WS) as a de-facto implementation protocol. In spite of the wide acceptance of WS in computing infrastructures, there have been few studies on WS characteristics. In this paper, we analyze WS using publicly available information that we collected weekly, between August 8<sup>th</sup> 2003 and January 30<sup>th</sup> 2004 from a UDDI[3] Business Registry (UBR). First, we study the evolution of the WS population and its geographic distribution. Second, we determine a few characteristics of public WS, such as the frequency of elementary types. Third, we develop a methodology for estimating WS message sizes. Fourth, we examine the liveness and response times of public WS. Lastly, using our methodology, we analyze the Amazon WS site and compare the message sizes predicted by our methodology with the message sizes observed during invocations.

Our initial results contradict common intuition. First, the number of public WS has not increased dramatically, although the intensity of the standardization and development activities in the WS domain continues to be high. Second, the geographic distribution of public WS is largely skewed, with about three fifths of public WS located in USA. Lastly, the sizes of WS response messages and their variation are smaller than that of Web documents. We expect our results to benefit WS application and tool developers, and to improve our understanding of this emerging research area. This survey is part of an ongoing research and detailed analysis results are published on our web site [1].

## 2. Population, Distribution, and Structure

Currently, about 1200 WS are registered in a UBR. Figure 1 summarizes the data collected during a six-month period. The

number of ‘valid’ WS, i.e. with a retrievable WSDL file, is substantially smaller than the number of registered WS: approximately 67% of the registrations are not valid. Furthermore, many of the downloaded WSDL[4] files omit mandatory elements or contain other syntax errors. During the six month interval, the number of valid WS decreases a little, which is contrary to the slight increase in the number of registered WS.




Figure 1. Web Services in UBR

Figure 2 (a) shows the geographic distribution of public WS on November 7<sup>th</sup> 2003. 63% of the WS are hosted in United States. Figure 2 (b) shows the distribution of WS hosting sites on the same day. From the fact that the fraction of US sites is smaller, we can infer that a larger number of US-resident WS are co-hosted.




Figure 2. Geographical Distribution (a) WS, (b) Hosting Site




Figure 3. Frequency of Elementary Types

We also analyze the WSDL files to determine the frequency of elementary, array, and compound types in WS messages. We

Copyright is held by the author/owner(s).

WWW 2004, May 17–22, 2004, New York, New York, USA.  
ACM 1-58113-912-8/04/0005

found that responses use more arrays and compound variables than requests do. Figure 3 shows the distribution of elementary types, with array and compound types expanded into elementary types. As most WS definitions do not specify array lengths, we had to select lengths for these arrays. The figure shows type distributions for array lengths 2 and 16; in addition, it shows that the string and string array types are very common.

### 3. SOAP Message Size

A SOAP[4] message can be divided into three parts – HTTP header, essential tag, and payload. Below is the equation used to infer the size of a SOAP message:

$$\text{SOAP message size} = \text{HTTP header} + \text{essential tag (SOAP envelope tag + SOAP body tag + namespaces)} + \text{payload (payload tag + summation of (type tag + value) for each elementary type field in parameters)}$$

We determine the size of each message component by examining real SOAP messages. We investigate several messages and determine the default size for the HTTP header, essential tag, and payload tag. Next, we estimate the number of elementary type fields and determine the average size of the XML representations for the fields of each type, including type tag.

We compare the distributions of SOAP messages, estimated using the above equation, to that of existing Web objects (see Figure 4). For Web objects, we use the model presented in [2], which studies the population of unique files.




Figure 4. SOAP message vs. Web objects: Array Length is 2

Contrary to the common expectation that SOAP messages are larger than current Web objects due to XML formatting, most SOAP messages are smaller than existing Web objects. For instance, while about 92% of SOAP messages are smaller than 2KB, only 45% of the existing Web objects are smaller than 2KB.

### 4. Liveness and Invocation Delay




Figure 5. WS and HTTP response delay

We verify the liveness of the valid WS sites. Our weekly measurements from two locations, IBM Watson and KAIST, show

that approximately 16% of the valid WS are down and that 2~3% more Web servers are alive; in addition, 96% of the live WS respond in two seconds or less. Figure 5 shows the response time CDFs for WS as well as for Web servers, as measured on Nov. 13<sup>th</sup> 2003; measurements performed on other dates show similar results. Our attempts to measure ping delays do not show any meaningful results, as most sites block ICMP ping messages.

### 5. Case Study

Amazon provides their WS for associates, suppliers, and developers. The main Amazon WS site is located in US and it is operated by Amazon itself. Their WS operations use only string and string array types. HTTP and WS response delays are 327 and 502msec when measured from IBM Watson, and 501 and 510msec from KAIST.

Figure 6 (a) and (b) show the message sizes, both real and estimated, for requests and responses, respectively. Note that the browse operations have two kinds of responses – lite and heavy. A lite response delivers the summary of the selected items, while a heavy response delivers all the available information. Figure 6 (a) shows that our estimation of request sizes are accurate. Figure 6 (b) shows that our estimations of response message sizes are less accurate. However, it should be noted that the line patterns are almost identical and that the line for heavy response is between the estimated lines.

