Similarity-based SOAP Processing Performance Enhancement

Joe Tekli, Ernesto Damiani, Richard Chbeir, Gabriele Gianini

► To cite this version:

Joe Tekli, Ernesto Damiani, Richard Chbeir, Gabriele Gianini. Similarity-based SOAP Processing Performance Enhancement. IEEE Transactions on Services Computing, 2011, 99 (PrePrints), pp.1939-1374. 10.1109/TSC.2011.11 . hal-00650641

HAL Id: hal-00650641

<https://hal.science/hal-00650641>

Submitted on 13 Dec 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Similarity-based SOAP Processing Performance and Enhancement*

Joe Tekli, Ernesto Damiani, Richard Chbeir, and Gabriele Gianini

Abstract—The Web Services (WS) technology provides a comprehensive solution for representing, discovering and invoking services in a wide variety of environments, including SOA (Service Oriented Architectures) and grid computing systems. At the core of WS technology lie a number of XML-based standards, such as the Simple Object Access Protocol (SOAP), that have successfully ensured WS extensibility, transparency, and interoperability. Nonetheless, there is an increasing demand to enhance WS performance, which is severely impaired by XML's verbosity. SOAP communications produce considerable network traffic, making them unfit for distributed, loosely coupled and heterogeneous computing environments such as the open Internet. Also, they introduce higher latency and processing delays than other technologies, like Java RMI and CORBA. WS research has recently focused on SOAP performance enhancement. Many approaches build on the observation that SOAP message exchange usually involves highly similar messages (those created by the same implementation usually have the same structure, and those sent from a server to multiple clients tend to show similarities in structure and content). Similarity evaluation and differential encoding have thus emerged as SOAP performance enhancement techniques. The main idea is to identify the common parts of SOAP messages, to be processed only once, avoiding a large amount of overhead. Other approaches investigate non-traditional processor architectures, including micro- and macro-level parallel processing solutions, so as further increase the processing rates of SOAP/XML software toolkits. This survey paper provides a concise, yet comprehensive review of the research efforts aimed at similarity-based SOAP performance enhancement. A unified view of the SOAP performance enhancement problem is provided, covering almost every phase of SOAP processing, ranging over message parsing, serialization, de-serialization, compression, multicasting, security evaluation, and data/instruction-level processing.

Index Terms—H.3.5.e. Web-based Services, H.3.5.F.XML/XSL/RDF, D.2.8.b. Performance Measures, H.3.4.d Performance Evaluation, H.2.0.a.Security, Integrity and Protection.

1 INTRODUCTION

OVER the past decade, web services have transformed the web from a publishing medium used to simply disseminate information, into an ubiquitous infrastructure that supports transaction processing [48]. The Web Services (WS) technology differs from traditional software integration frameworks such as CORBA [54], DCOM [35] and Java RMI [66], in that WS utilize well-established and open Web protocols and formats, chiefly HTTP and XML [7], allowing smooth interoperability among heterogeneous systems. Nonetheless, the very feature that makes WS universally usable, namely the adoption of the ubiquitous XML standard [7], makes it difficult to reach the performance lever required by large-scale processes and applications [12]. In this paper, we survey a number of issues related to WS performance, particularly in the context of WS communications, discussing the main performance bottlenecks and possible improvements.

An individual web service generally comes down to a self-contained, modular application that can be described, published and invoked over the Internet, and executed on the remote system where it is hosted [61]. WS mainly rely on two standard XML schemata:

- WSDL (Web Service Description Language) [10] which supports the machine-readable description of a web service's interface. It allows the definition of XML grammar structures for describing WS as collections of communication endpoints capable of exchanging messages.
- SOAP (Simple Object Access Protocol) [82] is the protocol specification for message exchange among WS. It is based on the XML data model, and usually relies on existing application layer protocols (e.g., HTTP, FTP, SMTP...) for message negotiation and transmission.

While these basic building blocks of WS technology are now firmly in place, performance issues have prevented using WS to implement large-scale distributed processes over large corporate networks or on the global Net. A major performance bottleneck resides in SOAP message processing [68]. The reason for SOAP performance criticality is twofold:

- On one hand, SOAP communication produces considerable network traffic, and causes higher latency than competing technologies, like Java RMI and CORBA [38]. This is a central problem especially within wireless communication networks with their relatively low bandwidth and high latency [59], as well as the rising number of mobile computing devices (e.g., PDAs and mobile phones) increasing service demand, and consequently network bandwidth consumption [48].
- On the other hand, and perhaps more importantly, the generation and parsing of SOAP messages, and their conversion to-and-from in-memory application data can be computationally very expensive [1, 4]. In this paper we adopt the following terminology: the

*Work Supported in part by *Fondazione Cariplo*, and *Japan Society for the Promotion of Science (JSPS)*.

- Joe Tekli is with the Department of Science and Technology, Shizuoka University, Hamamatsu, 432-8011 Japan. Email: joe.tekli@inf.shizuoka.ac.jp
- Ernesto Damiani and Gabriele Gianini are with the Department of Information and Technology, Università degli Studi di Milano, Crema, 65 - 26013 Italy. E-mails: {ernesto.damiani, gabriele.gianini@unimi.it}
- Richard Chbeir is with the LE2I Laboratory UMR-CNRS, University of

process of translating a memory object according to a serialization format into an XML object is called *serialization*. The process of converting an XML structure into a memory object will be called *de-serialization*. For complex XML structures, both these processes are computationally expensive. In fact, the translation between in-memory numeric data of type double and the ASCII-based XML representation format has been shown to consume over 90% of the end-to-end SOAP message time [12], which proves critical for various kinds of WS applications, ranging over business transactions (e.g., online booking and stock quote services), and scientific data processing (e.g., grid computing).

Several techniques have been proposed to improve SOAP processing performance. Many of them exploit the well-known concepts of similarity and differential encoding to i) reduce processing time, in message parsing [45, 70, 71], serialization [4, 21], and de-serialization [1, 68], as well as to ii) reduce network traffic via SOAP message compression [81] and multicasting [6, 58, 59]. Similarity-based SOAP performance enhancement is based on the straightforward observation that SOAP message exchanges usually involve highly similar messages. Messages created by the same implementation usually have the same structure, and those sent from a server to multiple clients tend to show similarities in structure and content (e.g., stock quote services [59] involving a large number of similar transactions requesting the latest stock data, as well as online booking and meteorological broadcast services [6]).

Thus, various efforts have been undertaken to process SOAP messages taking into account their similarities. The main idea is to identify the common parts of SOAP messages, to be processed once, regardless of the number of messages. Processing is only repeated for those parts which are different, avoiding a large amount of unnecessary overhead.

Another source of overhead is checking SOAP messages against security policies. Recently, several research efforts have focused on the impact of WS-Security policy evaluation on SOAP messages. WS-Security policies [19] specify authorizations, signature and encryption schemes on SOAP elements and contents, and may introduce substantial processing overhead without (or despite) ad-hoc performance enhancement [6, 14, 71]. Indeed, evaluating WS-Security policies can introduce an overhead much larger than standard WS invocation processing (6.9 times in average, according to [37]). A major portion of this overhead is related to the requirement of providing message level security (as opposed to channel-level security such as with TLS [79]) and to the XML encoding of message content.

Other performance bottlenecks arise from the limited amount of parallelism available on a conventional processor. Efficient parsing of SOAP and XML streams, as well as processing variable length encoded character streams would require hardware support for longer processing pipelines than standard CPUs can support. Handling XML streams entirely in software (for instance, by mapping processing pipeline stages to software threads) prevents the execution speed to be improved beyond a best processing rate of tens of clock cycles per character, and that best case performance can result in rates on the order of hundreds of clock cycles

per character for many practical XML applications [78]. As a result, recent studies have addressed these performance bottlenecks by investigating non-traditional processors, namely parallel processing architectures and “XML machines”, e.g., [8, 23, 30].

The goal of this survey paper is to provide a unified view of the problem, connecting the different aspects and techniques related to similarity-based SOAP processing performance enhancement, including WS-Security policy evaluation and XML parallel processing architectures. The remainder of the paper is organized as follows. Section 2 presents a glimpse on SOAP message processing, introduces its performance metrics, and discusses its main bottlenecks. In Section 3, we categorize, discuss and compare some of the most prominent methods to SOAP performance enhancement. Section 4 discusses prominent ongoing challenges. Section 5 concludes the paper.

2 WS AND SOAP PROCESSING PERFORMANCE

Experience with Service Oriented Architectures (SOA) has shown that WS performance is a crucial success factor for large-scale business processes [48]. It becomes even more crucial when services are made available on the open Web, where (i) user requests to a certain service provider/company tend to increase with the amount of information and services the company makes available online [49], and (ii) the fidelization of service consumers is on average lower than on a SOA infrastructure. If service latency becomes too high, clients may become frustrated and simply switch to another site or service offering the same functionality. Hence, WS performance problems can bring all kinds of undesired consequences, including financial and sales losses, decreased productivity and a bad reputation for a company [48]. Moreover, as the web evolves, mobile computing devices (e.g., PDAs and mobile phones) add another challenge to web services performance: wireless communication networks with their relatively low bandwidth and high latency [59]. Finally, current web systems and services are usually characterized by integration with databases, scheduling and tracking systems (e.g., Google Maps), requiring altogether high performance levels [27].

In the following, we first briefly present the key metrics which characterize WS performance levels. We subsequently discuss the various aspects of SOAP processing, and the corresponding performance bottlenecks.

2.1 Evaluation Metrics

Service-oriented infrastructures share some properties with component-based [26, 60] and web-based [47] applications, hence to some extent it is possible to apply existing resource metrics from the component-based software engineering and web applications domains in the context of SOA [60]. Namely, it is possible to classify performance metrics in three main categories: delay, bandwidth and usage, with *response time*, *throughput* and *network traffic* [48, 59] as the most relevant metrics normally used to assess the performance of WS for each category respectively. Summary values of those metrics are normally obtained by aggregation in time and/or aggregation in space, or concatenation in space. A taxonomy of the relevant metrics can be found in [72] and references therein.

Fig. 1. Outline of a classic SOAP remote service call processing chain.

2.1.1 Response Time

Response time (also called *latency* or *end-to-end response time*) is the time perceived by a client to obtain a reply for a request for a web service. It includes the network time (latency and transmission delays on the communication link), as well as the processing delays at the server end-point (service execution) and at intermediary nodes (switching time introduced by hubs, routers and modems)[48]. The process with the longest processing delay in the processing chain is usually the key determinant of response time, and is identified as *bottleneck* (or time-sink). Response time is measured in time units.

2.1.2 Throughput

While response time is a performance metric typically of interest to end-users, throughput, which is defined as the number of requests executed per unit of time (e.g., I/O operations per second), is of more interest to administrators. It is usually evaluated on the server side [48]. There are many possible throughput metrics depending on the definition of unit of work. It is common to distinguish between point-to-point (or link) throughput (to quantify transport performance), node throughput (to quantify processing performance) and overall throughput in the system (a.k.a. consistent throughput in the system) [60]. The overall system throughput is bound by the local throughput (link throughput and nodal throughput) of the least performing components in the transport and processing chain. Its basic unit of measure is *byte/sec*, however, for web service providers, it can be measured in *req/sec* – requests per seconds, *HTTPOps/sec* – HTTP operations per seconds for web servers, or *tps* – transactions per seconds [71].

2.1.3 Network Traffic

The total network traffic for a communication scheme or session (e.g., *conversation*, i.e. a SOAP message exchange among two service end-points) consists of the total size of all session-related messages sent over the network for the duration of the communication [59]. In other words, it encompasses the total number of bytes (corresponding to all messages exchanged during the communication session being evaluated) that are transmitted over the network [81]. Other related performance metrics exist, including: average utilization of a node, incoming/outgoing message rates, incoming/outgoing traffic for a node or the overall message rate in the system, which can also be measured in bytes or number of messages.

Over the past few years, several works have studied web service performance, e.g., [3, 22, 45, 68]¹. Most of them focus on SOAP processing and message exchange as the major players affecting web service performance levels. In the remainder of this section, we present a glimpse on SOAP processing, so as to pinpoint SOAP performance bottlenecks.

2.2 A Glimpse on SOAP Processing

SOAP (Simple Object Access Protocol) [29] was specifically conceived as a messaging protocol to support interdependent interactions between otherwise independent entities, namely WS [12]. It is based on XML [4] and can support a variety of message exchange patterns, including request-response, one way messages, remote procedure calls, and peer-to-peer interactions [28].

Fig. 1 depicts a simplified activity diagram describing a typical SOAP remote service call processing scenario. Given two end-point services, usually identified as *client* and *application server*, an outgoing client SOAP message consists of a method invocation, a.k.a. (also known as) client SOAP request, underlining a client call for method destined to the application server. An outgoing server SOAP message consists of a method response, a.k.a. server SOAP response, carrying the result of the action performed at the application server, following the corresponding method invocation. SOAP request and response messages are usually similar in structure. They both follow the same schema defined in the WSDL interface definitions of the services involved in the communication process. In general, a SOAP request/response message consists of a root node entitled *Envelope*, encompassing two elements: *Header* and *Body*. Consider for instance the sample SOAP messages in Fig. 2.

- *Envelope* provides the serialization context and namespace information for elements and parameters utilized in the message.
- *Header* contains auxiliary information which is not related to the method invocation (or response) itself, such as transaction management and client/server information (e.g., client/server addresses, URL of final message destination).
- *Body* contains the actual data carried in the SOAP message. It usually starts with a sub-element entitled with the method (or method response) name. The latter would encompass a child node for every parameter required to perform the local invocation.

¹Most references in this paper address, in one way or another, web ser-

vice performance. We only give a few here for clearness of presentation.

Fig. 2. Sample SOAP request and response messages.

As shown in Fig. 1, a common SOAP message exchange scenario consists of the following steps. First, a SOAP request message is created at the client side. Message creation requires serialization which consists in converting between in-memory application data representations and XML-based messages (Step 1). The request message is sent to the server application, usually via classic IP unicast routing (Step 2). At the server side, the message is first parsed, i.e., processed for lexical analysis (identifying characters and extracting tokens such as tags and contents) and validation (verifying the message's structural integrity w.r.t. the corresponding WSDL definition) (Step 3). The application server consequently evaluates its security policy rules on the received message, so as to identify and process those parts of the message which were assigned security constraints (authorization rules, signature verification...) (Step 4), followed by message de-serialization (converting between XML and the in-memory data representation) in order to be processed via the service executor (Step 5). As for the SOAP response message, the same procedure is undertaken, but this time in the inverse direction. The response message is created, i.e., serialized (Step 6), sent back to the client service via unicast routing (Step 7), parsed (Step 8), evaluated w.r.t. the client security policy rules (Step 9), and de-serialized so as to transfer the processed data to the client service component (Step 10).

2.3 SOAP Performance Bottlenecks

SOAP's XML-based nature, which makes the SOAP protocol universally usable, tends unfortunately to work against achieving high performance [12]. The impact of XML message encoding on overall SOAP performance is omnipresent in almost every step of SOAP processing, underlining: i) high response time and low throughput in SOAP serialization [2, 4], parsing [45, 70, 71], security evaluation [6, 14], and de-serialization [1, 68], mainly due to XML processing and the conversion between in-memory data and the ASCII-based XML format, as well as ii) high network traffic and bandwidth consumption during message transmission and routing [58, 59, 81], due to XML's verbosity and redundant textual characteristics.

To give an idea of the problem size at hand, we discuss the results of three studies, [17, 37, 81], evaluating the performance levels of SOAP in comparison with existing integration technologies, namely CORBA [54] and Java RMI [66]. Fig. 3 depicts the response time for a SOAP service call

processing, i.e., the time required to generate and send a service request message and to receive its corresponding service response message, using two SOAP implementations (Java-based, Microsoft VB 6.0 toolkits) [17], in comparison with similar procedures to remote method invocations using CORBA [54] and Java RMI [66]. Timing results in both Fig. 3.a and Fig. 3.b show that SOAP performs very poorly in comparison with competing technologies. The time performance gap increases significantly when exchanging numeric data (e.g., integer arrays in [17]), which is due to the expensive process of converting in-memory numeric data to-and-from ASCII-based XML [12]. Fig. 4 depicts network traffic created by SOAP (two Java-based and Microsoft .Net based toolkits were considered) [81], CORBA [54] and Java RMI [66], when varying the number of method invitations between two client and application server end-points. Results show that SOAP produces significantly more network traffic than existing technologies. It requires almost three times more bandwidth than Java-RMI and CORBA, the latter using dedicated binary encodings for message exchange, in comparison with SOAP's XML-based textual format [81].

Fig. 3. Comparing SOAP response time, with CORBA [54] and Java RMI [66].

Fig. 4. Comparing SOAP service call network traffic, with CORBA [54] and Java RMI [66].

The need of encrypting and signing SOAP messages, which is of paramount importance especially when accessing services available on the open Net, has introduced additional delays. The WS-Security standard [19] is now widely used to express (in XML) the service providers' policies regarding what parts of the SOAP XML tree need to be encrypted and signed. In a recent study [37], the authors evaluate the additional overhead introduced by WS-Security policy evaluation w.r.t. standard processing of SOAP invocations. Their results show that WS-Security increases SOAP response time by a factor of 3 on average, while SOAP messages when using WS-Security are 6.9 times larger than unsecured SOAP messages (affecting network traffic accordingly).

In addition to evaluating the performance bottlenecks of SOAP itself, related works in [8, 39, 78] (among others) have addressed the shortcomings of conventional hardware computing architectures in handling XML-based data for large scale data sets and WS computing environments. They highlight the limited amount of parallelism in XML processing: both at the data level [8, 78] (i.e., in processing multiple pieces of data with one instruction), and at the instruction level [39, 78] (i.e., executing multiple instructions concurrently, a.k.a. multi-processing). This family of hardware-based studies usually underlines the limitations of conventional processors in providing an efficient enough solution to evaluate multiple conditions of various types in parallel, which is central in XML string and character processing (e.g., verifying character integrity, whether an end tag matches a previously processed start tag, whether an attribute name is unique for a given element, and so on).

Some works [12, 28] address transport protocol bindings, namely the shortcomings of HTTP [24] as the application layer protocol used with SOAP for message negotiation and transmission. The authors in [12, 28] conclude that HTTP (specifically the earlier HTTP 1.0 version) negatively affects SOAP processing, and that it induces higher SOAP response time due connection and message transmission overheads.

All relevant aspects of SOAP processing, the impact of the XML-based parallelism on SOAP performance, as well as the various solutions to SOAP performance enhancement to-date, are detailed in the following sections.

3 IMPROVING SOAP PROCESSING PERFORMANCE

As mentioned previously, SOAP processing performance enhancement has been widely researched [6, 45, 58, 59, 70, 71]. Many approaches build on the simple observation that SOAP message exchange usually involves a number of highly similar messages. Invocations sent from the same client often reflect similar information needs, and thus similar SOAP message requests [21]. Likewise, messages sent from the same server to a single and/or multiple clients usually share strong similarities. Typical examples are various [6] such as stock quote services [59] (involving a large number of transactions requesting the latest stock data, hence similar stock quote request and response messages are processed), as well as online booking systems, and meteorological broadcast services [6], etc.

Several proposals addressing SOAP performance enhancement exploit, in one way or another, the similarity between SOAP messages, in order to gain in performance, e.g., reducing execution time, increasing throughput, and saving on network traffic. The main idea is to identify the common parts of SOAP messages, to be processed once, regardless of the number of messages.

We classify these solutions based on the performance metrics they target, and on the specific SOAP processing operations they address.

3.1 Methods for Improving Service Execution Time

Improving service execution time (i.e., attaining lower response time and higher throughput), has been investigated in various aspects of SOAP processing, addressing serialization, parsing and de-serialization operations.

3.1.1 SOAP Serialization

As mentioned previously, the serialization of SOAP messages consists in converting in-memory data types into XML. In this context, the main bottleneck consists in transforming in-memory data of numeric types into the ASCII-based XML representation format [12]. Consequently, the authors in [4], building upon the findings in [12], introduce a method for differential SOAP serialization, called bSOAP. The main idea consists in storing the SOAP messages in a dedicated buffer, to be used as templates for future outcalls, instead of discarding them after they have been sent over the wire. The message is normally serialized and saved during the first invocation of the SOAP call. Subsequent calls which share identical or similar message structures, as the message in the buffer, would avoid a significant amount of processing by only serializing the changes to the previously sent message. The authors address the problem of change tracking between in-memory data, and their serialized representations. Dedicated indexed tables, i.e., DUTs (Data Update Tracking), are associated with each serialized message, keeping track of the in-memory location of each field in the original structure to be serialized, and its position in the serialized message. A *dirty bit* is associated with each field, to keep track of those fields whose values have changed since the last send, in order to check which parts of the last message could be reused. Experimental results in [4] confirm the approach's better time performance, in comparison with regular serialization, and show that serialization time is linearly dependent on the percentage of in-memory values that must be re-serialized (reflected by the number of dirty bits that are changed). When the whole message has to be serialized, bSOAP's serialization time is almost equivalent to that of existing SOAP toolkits, e.g., gSOAP [77] and XSOAP [63] (cf. Fig. 5.a). Nonetheless, when the exact message is to be sent again (i.e., when none of the dirty bits are changed), time performance gain is maximal (almost 1000%, cf. Fig. 5.b).

a. Comparing bSOAP, to alternative approaches, i.e., gSOAP [77] and XSOAP [63].

b. Serialization time, when various percent-tages of stored values are re-serialized.

Fig. 5. Time performance of bSOAP differential serialization (reported from [4]).

In subsequent studies [2, 3], the authors address bSOAP's buffer management, mainly padding, which consists in stuffing the serialized message with white spaces to reduce the cost of message expansion when the latter is to be updated. Padding is useful when the new serialized form of some value does not fit in the current space allocation (e.g., the value of an integer variable $i=3$ which holds a single character space, is to be updated to $i=1003$ in the new serialized message, which requires four character spaces). Hence, padding allows on-the-fly message expansion, DUT table entries being updated accordingly.

Various other SOAP buffer optimization techniques have been proposed [2, 3, 12, 77], namely chunking (dividing the SOAP message into chunks stored in different memory locations, to be processed separately) and streaming (pipelined-send, each message chunk being sent as soon as it is serialized, thus allowing an overlap of computation and communication). However, even after these optimizations, the conversion from in-memory data to the ASCII representation (over 90% of the end-to-end time) remains the most critical bottleneck [12], which emphasizes the relevance of differential serialization [4].

An approach comparable to differential serialization [4] is introduced in [21]. It addresses client-side SOAP message caching and allows entire request messages to be cached and sent as is. It also allows partial caching by reusing cached messages with identical structures, updating element values for subsequent sends. Similarly to [4], it relies on dedicated indexed structures in detecting correspondences between cached and outgoing messages. Nonetheless, the approach in [21] does not address partial structural matches (i.e., caching messages with partially different structures) as in [4], but only caches messages with identical structures. In addition, the authors in [21] do not discuss how to handle mismatched data sizes that require message resizing and expansion.

3.1.2 SOAP Parsing

As mentioned previously, SOAP parsing consists in analyzing the contents of the incoming SOAP message, to be consequently transformed into their in-memory application format via the de-serialization component. In general, SOAP parsing consists in analyzing the characters in the SOAP message, extracting tokens such as tags and text, and then extracting and validating the underlying XML structure (cf. Fig. 6.a). These tasks can be achieved using functions of existing XML parsers such as DOM [84] and SAX [47].

In this context, a few studies have proposed using special-purpose parsers, considering the particularities of XML and SOAP messages in order to amend performance. One of the earlier XML-based approaches promotes partial parsing [53], by i) extracting the XML document structure (node references and hierarchical relations) in a pre-processing phase, and then ii) parsing only those parts of the document required by the application program, by looking up the document structure. The authors in [53] show that performance improves only when document (application) coverage is less than 80%, and that it otherwise declines due to pre-processing overhead. In [11, 74], the authors investigate the optimization of SOAP lexical analysis, using schema (WSDL) information, to more efficiently identify lexical tokens (e.g., tag names, attributes...). Yet, such methods only target lexical analysis, disregarding byte-level character encoding and validation optimizations [69]. On the other hand, XSOAP [63] targets validation optimization and attempts to improve SOAP message validation performance by only executing the validation process on those elements specific to SOAP, namely *Envelope*, *Header* and *Body*. Remaining parts, which usually consist of classic XML tagging, are disregarded in order to gain in parsing time. However, when the corresponding service requires complete message validation, the invalidated SOAP message parts have to be processed via a dedicated validation function to be added by the programmer in the service program [70], thus minimizing performance enhancement. A recent work[87]introduces a Table Driven XML

(TDX) parser, that combines the lexical analysis and validation of SOAP XML messages in a single pass. The idea is to pre-record the states of an XML parser produced from the corresponding (Schema) WSDL service description, as grammar productions rules in tabular form, and then to utilize a runtime streaming parsing engine to break up the SOAP message into a token stream, to be processed for well-formedness verification and validation at once. The authors in [87] show that their approach is more efficient than existing XML and SOAP toolkits where validation is enforced separately [5, 65, 77](e.g., it runs six times faster than gSOAP [77]). Yet, TDX's performance is shown to be comparable (and even lower) when evaluated against a non-validating schema-specific SOAP parsing approach [74].

a. Traditional SOAP (XML-based) parsing.

b. Differential SOAP parsing.

Fig. 6. SOAP parsing.

Instead of focusing on a specific phase of SOAP parsing, such as lexical analysis, or limiting the range of SOAP elements validation, more recent proposals in[45, 70, 71] focus on differential parsing, exploiting the similarities between SOAP messages, in order to skip unnecessary parsing altogether (including character encoding, lexical analysis, and validation) as depicted in Fig. 6.b. In the following, we discuss the main approaches to differential SOAP parsing.

Template-based: T-SOAP [70] makes use of a predefined template, modeled via a finite state automaton (FSA), memorizing the basic structure of the SOAP messages, extracted from the corresponding WSDL definition schema¹. It

allows the identification of invariant and variable tag parts in the SOAP messages. Consequently, each incoming SOAP message is matched to the predefined template, and only those parts of the message, which correspond to variable parts in the template, are parsed (the invariant parts being already parsed in advance). While it induces a significant gain in processing time, in comparison with classic SAX [47] and DOM [84] parsers, a major limitation of T-SOAP [70] is its restriction to messages conforming to the same basic structure. In other words, a SOAP message with a structure different than that underlined in the predefined template would not benefit from T-SOAP [70] and would have to be parsed from scratch.

Multiple Templates: In [45], the authors propose a more dynamic approach by managing multiple templates based on actual SOAP message structures, instead of using a single predefined schema structure. Incoming messages are first matched against the automaton, describing multiple message templates merged together. If the message matches any of the templates, then parsing is undertaken w.r.t. the variable parts of the corresponding template, similarly to [70]. Otherwise, parsing is undertaken via an ordinary DOM-based processor [84], and a new template corresponding to the unmatched message is created and appended into the automaton, to be exploited in upcoming parsing operations. While this technique provides more flexibility than T-SOAP [70], the authors in [45] underline that their method requires more memory for storing the combined automaton, and additional processing time for updating the latter with new message templates. Experimental results in [45] show however that the proposed approach performs better, in time and memory usage, than classic SAX [47] and DOM [84] parsers.

Detecting Repeatable Structures: An extension to the approach in [45] is provided in [71]. The authors in [71] introduce an improved automaton, able to consider repeatable structures in SOAP messages, which are not considered in [45]. That is because the automaton in [45] is string-based and processes SOAP messages as a series of invariant and variable sections of string characters (i.e., byte sequences), whereas the new automaton in [71] considers the XML syntax (e.g., XML tagging) in its definition of states and state transitions. Detecting repeatable structures allows reducing the number of templates to be appended to the automaton, the latter becoming more expressive. Consequently this allows reducing memory and processing time needed for storing and updating the automaton respectively, thus further enhancing parsing performance. Experimental results in [71] show improved memory usage and time performance w.r.t. the approach in [45], as well as a classic DOM parser [84].

Note that both methods described in [45, 71] have been developed in the context of WS-Security processing. Their main objective is therefore to improve security policy evaluation performance, by repetitively applying security rules only on those parts of SOAP messages which are different,

processing the common parts only once. Yet, other methods aimed at improving security policy evaluation performance have been proposed in the context of SOAP message multicasting [6, 14] (which is discussed subsequently). Thus, for clearness of presentation, we disregard security aspects in this section, and provide a unified view of SOAP security policy evaluation performance, covering all related methods, in Section 3.3.

3.1.3 SOAP De-serialization

De-serialization is the process of converting XML messages to in-memory application objects, to be processed by the service executor. It can be viewed as the symmetric function of serialization. Recall that with serialization, the SOAP message is the target for recycling, whereas with de-serialization, the target is an application object.

Approaches to improving SOAP de-serialization performance build on the observation that memory object creation, based on SOAP XML messages, is an expensive task (mainly due to data-type transformation – conversion from ASCII-based textual representation to in-memory numeric types, and the processing of the XML tree hierarchy [68]). Hence, the main idea is to avoid fully de-serializing each incoming message, by exploiting already constructed objects which were de-serialized previously. In other words, de-serialization is differential and is only applied to those portions of the SOAP messages which have not been de-serialized previously. To our knowledge, two studies have been developed in this direction, which we identify as *automaton-based* [68] and *checksum-based* [1]. We also stumbled on a more recent approach, *XML Screamer* [39], which promotes tight integration between software layers to avoid unnecessary de-serialization processing.

Automaton-based: The authors in [68] propose an automaton-based approach, consisting of two main functions. The first consists in generating an automaton based on incoming SOAP messages (similarly to SOAP parsing approaches in [45, 70]), and then conducting de-serialization in the usual way, creating a link between the defined automaton and the application object. The second function is to match an incoming message with the existing automaton, and if matched, return the linked application object to the SOAP engine after partially de-serializing only the portions that differ from previous messages. The de-serialization approach described in [68] could exploit the methods in [45, 70, 71] in building the de-serialization automaton. Recall that SOAP parsing and de-serialization are complementary operations, and allow SOAP message analysis (Fig. 1).

Checksum-based: In [1], the authors propose to periodically checkpoint the state of the de-serializer and to compute checksums⁴ for portions of the incoming SOAP messages. In short, the de-serializer runs in one of two modes: *regular* and *fast*. In regular mode, the de-serializer processes SOAP message tags and contents as a normal SOAP de-serializer, creating checkpoints and corresponding message portion checksums along the way. It switches to fast mode once it recognizes that the parser state is the same as one that has been saved in a checkpoint. In fast mode, the de-serializer compares the sequence of checksums against those associated to the most recently received message. If the checksums match,

⁴A FSA is usually modeled as $(P, \Sigma, p_s, F, \delta)$ where: P is a set of states, Σ the set of labels, $p_s \in P$ is the start state, $F \subset P$ is a set of final states, and $\delta: e \times R \rightarrow P$ is a transition function where $e \in \Sigma$, R is an expression over P , and $p \in P$ [34]. Standard procedures for producing automata and testing the membership of data instances w.r.t. automata have been thoroughly studied in language theory [34].

then the already de-serialized objects corresponding to the portions of the SOAP message at hand are exploited in a straightforward manner, without additional processing. Otherwise, when a checksum mismatch occurs, the system switches from fast to regular mode, where it processes SOAP tags and contents as a normal de-serializer.

The authors discuss and experimentally validate the performance of their approach, considering the relation between i) the amount of similarity between incoming messages, which otherwise determines the percentage of time the de-serializer spends in fast mode, ii) how quickly the system can recognize the need to switch modes (from fast to regular, and vice-versa), and iii) the overhead of creating checkpoints, and comparing checksums.

Fig. 7. Comparing regular de-serialization and full differential de-serialization time[1].

On one hand, if the new message is completely different from the previous one (which is the worst case scenario), the differential de-serializer runs slightly slower than a normal de-serializer since it does the same work, plus the added work of calculating and comparing checksums. On the other hand, when all checksums match, i.e., when the new message is identical to the previous one (which is the best case scenario), the cost of de-serialization is replaced by that of computing and comparing checksums, which is significantly faster (speedups up to 41 times have been recorded by the authors, cf. Fig. 7). The authors also mention that using checksums to match portions of SOAP messages can be error-prone, (since checksums themselves are not perfect by definition), but the possibility of changes going undetected is extremely low, in comparison with the substantial gain in performance.

Note that both methods in [1, 68] have not been evaluated w.r.t. each other, so as to compare their relative improvements in SOAP de-serialization performance.

XML Screamer: In a more recent study, the authors introduce XML Screamer [39], an optimized system providing tight integration across levels of software, combining: i) schema-based XML parsing (character encoding, token extraction, and validation) and ii) de-serialization, in one single processing layer (as opposed to separate layers -Fig. 6.a), in order to avoid unnecessary data processing, copying (to/from memory), and data-type transformations. The authors adopt a design principle requiring that each character and/or string in the input document be ‘visited’ only once (if

possible), so as to reduce repeatable scans of the same data and corresponding unnecessary overhead (e.g., tests to verify whether a character is an angle bracket ‘>’, or an expected element name character, are performed only once following [39], whereas such tests are repeated multiple times - during parsing, and de-serialization - in traditional XML/SOAP toolkits). Experimental results in [39] show that XML Screamer delivers from 2.3 to 5.3 times the throughput of traditional SOAP toolkits[5, 65] (cf. Fig. 8).

Note that the combination of software layer integration optimization [39], with similarity-based SOAP parsing [45, 70, 71] and de-serialisation [1, 68], has not been investigated to date. We believe this to be a very interesting research topic which could yield promising performance improvements in the near future.

3.2 Methods for Reducing Network Traffic

Another major drawback of using SOAP is its voracity for bandwidth, compared to competing solutions such as CORBA [54] and Java RMI [66]. Even though today’s networks can be powerful enough to provide sufficient bandwidth, the latter remains crucial in several applications, namely in mobile computing [59] (e.g., wireless and cellular platforms), as well as sensor networks [81]. In this context, the problem of SOAP bandwidth reduction has been investigated on two levels: i) SOAP compression [81] in order to reduce message size prior to transmission, and ii) SOAP multicasting [58, 59] so as to optimize SOAP traffic travelling on the wire.

3.2.1 SOAP Compression

Various methods have been proposed for classic text and XML compression, namely gzip [20], WBXML [46], XMILL [42], and ESAX [9]. Text compression techniques (e.g., gzip) could be exploited with XML-based data (e.g., SOAP), since the latter are usually stored as ASCII-based text files. Nonetheless, a comparative study conducted in [81] showed that existing compression methods for classic XML documents might not always be appropriate in the context of SOAP. That is due to the fact that SOAP messages are of relatively smaller sizes (a few kilobytes), in comparison with other kinds of XML-based documents (e.g., SVG [85], MPEG-7 [52]..., usually in the order of hundreds of kilobytes). Hence, existing compression methods might yield coding tables (i.e., tables mapping symbols to their bit codes) which require more space than the original SOAP messages themselves [81] (cf. Fig. 9.a). In other words, compression results for large files are not necessarily transferable to small files, which is the case of SOAP messages. Following this observation, the authors in [81] propose a differential compression framework specifically aimed toward SOAP messages, exploiting the similarities between SOAP messages sent or received by the same service. The approach is based on XML differential encoding, which basically means that only the differences between SOAP messages should be sent over the wire. In brief, the authors exploit the WSDL schema definition to generate a SOAP message skeleton (the same would be available at the sender/receiver sides) describing the structure and tagging of corresponding SOAP messages (i.e., SOAP element/attribute names and corresponding parent/child relations, disregarding values). Consequently, only the differences between the SOAP message and the prede-

¹A checksum is a fixed size datum computed from a block of digital data (of fixed and/or variable size) to detect accidental errors that may occur during transmission or storage [50].

finer skeleton are transmitted, along with corresponding SOAP message element/attribute values. The differences in structure and tagging, as well as element/attribute values, are consequently patched to the same skeleton at the receiver side in order to reconstruct the original message.

The authors argue that the effectiveness of their method depends on the degree of resemblance between the generated skeleton and the actual SOAP messages, which strictly influences compression rate: a higher resemblance yields smaller difference files, which in turn underlines a higher compression rate. They test two existing implementations of XML diff encoding tools (XUdate [41] and DUL [51]) in their experimental evaluation, proving that their approach yields better compression rates than existing XML-based compression techniques (Fig. 9).

Fig. 9. Comparing the effectiveness of differential SOAP compression, in comparison with alternative text-based (gzip [20]) and XML-based (XMILL [42]) techniques.

The authors evaluate the execution speed of their approach, and show that it is slower than gzip [20], which introduces a major computational burden w.r.t. service execution time. In fact, gzip itself has been shown to be computationally expensive, exceeding the combined cost of XML serialization and data transport over LANs [28, 73]. Thus, while SOAP compression seems central in reducing network traffic, particularly when network bandwidth is very limited, its execution time underlines an equally serious drawback, which (to our knowledge) remains an open problem.

3.2.2 SOAP Multicasting

Another approach to reduce SOAP network bandwidth consumption would be to perform multicasting, a well-known technique that allows to conserve network bandwidth in applications where the same data is to be transmitted to multiple clients [86]. The main idea is to avoid sending replicated unicast messages over the wire by simultaneously delivering identical messages to a group of destinations, in a single aggregate message, only creating copies when the network links to the multiple destinations split [59, 86]. In general, multicasting would be effective when the number of receivers for a given service is sufficiently large and there is sufficient commonality in their interests, which happens to be the usual case with SOAP [59].

In this context, the authors in [59] put forward SMP, a Similarity-based SOAP Multicasting Protocol. It is built on top of SOAP unicast, and does not rely on low level (IP) multicast, in order to avoid complex network configurations at intermediate nodes (hubs and routers). In addition, SMP's main contribution and originality consists in grouping and transmitting together similar SOAP messages, and not only

identical messages such as with traditional (IP) multicasting. An SMP message consists of two parts: SMP header and SMP body. The SMP header stores the addresses of destinations to which the messages should be sent. The SMP body is composed, in turn, of two parts: the *common* part section containing common values of the messages, and *distinctive* part section containing the different parts of each message. The aggregate SMP message is consequently encapsulated within the body of a classic SOAP message, which header encompasses the address of the next router along the path to all intended recipients. Each midway router would parse the SMP header and examine its routing table to decide the next hops for each client address. The router then separates client addresses into groups, splits the SMP message accordingly, and forwards the appropriate information to the next hop. The SMP message is split so that only relevant information (i.e., information destined to the designated clients) is sent down the stream path. During splitting, multiple copies of the input message are first produced, one for each downstream link that the router connects to. The client list in each newly generated message header includes only those destinations that will be routed through that hop. *Distinctive* items in the original SMP message are analyzed and removed if they are not intended for clients beyond the next hop. The *common* part is obviously replicated in all outgoing messages. If the next hop connects directly to an end-point service, a standard SOAP unicast message is extracted from SMP and sent to the client service component.

The authors exploit an XML-based similarity measure [44] to quantify the resemblance between SOAP messages, so as to only aggregate the most similar ones. In addition, a dedicated indexing technique is also introduced to reduce SOAP message size by omitting full tag names and leveraging the organization of common and distinct parts in the SMP message.

In a subsequent study [58], the authors propose an enhanced routing protocol to further improve the performance of their SMP multicasting approach. In their original proposal [59], they used Dijkstra's Open Shortest Path First (OSPF) routing algorithm, which routes the message using the shortest path from a source to a destination. In their later study [58], the authors introduce tc-SMP (traffic constrained SMP) exploiting a similarity-based routing algorithm for transmitting messages following paths which maximize shared links between highly similar messages. This allows optimizing SMP network traffic distribution and thus further reducing overall network traffic (cf. Fig. 10.a).

Fig. 10. Comparing network traffic and average response time with tc-SMP [58], SMP [59], traditional multicasting and unicast (reported from [58]).

The authors also evaluate the performance penalty, in response time, of tc-SMP and SMP over traditional multicasting (simply multicasting identical messages) and unicast transmissions (cf. Fig. 10.b). It is mainly due to the processing overhead required to measure the similarity between messages and aggregate similar ones (for both tc-SMP and SMP), as well as setting up the routing tree (in the case of tc-SMP). In short, results show that tc-SMP induces an average 3.5 to 5 times reduction in network traffic, compared to an average 2.5 times increase in average response time, which is considered acceptable by the authors, particularly in scenarios where bandwidth is limited such as with wireless and sensor networks.

In addition to network traffic optimization with classic SOAP message communications, differential SOAP multicasting (SMP) has been recently investigated in the context of secure SOAP message exchange [6, 14], in order to improve SOAP security policy evaluation performance.

3.3 Improving SOAP Security Policy Evaluation Performance

In the past few years, the growing demand on mission-critical WS applications (e.g., financial transactions, stock market...), has underlined an urgent need to provide trustworthy and secure services [48]. Nonetheless, security provision may introduce a substantial additional overhead, which has motivated researchers to start investigating the impact of security policy evaluation on WS performance.

WS-Security policy evaluation [19] consists in checking and verifying the access and usage security constraints defined on SOAP messages. It is performed both at the client and server application end-points, each w.r.t. its own policy rules (cf. Fig. 1). A WS-Security policy usually underlines a set of rules (actions), specifying security constraints (e.g., authorizations, signatures, encryption...) on particular SOAP elements and contents [6, 15]. A security policy rule can be characterized in a 3-tuple entity: (*subject*, *object*, *rule*), where *subject* identifies the users to whom the rule applies, *object* identifies to which messages, or portions of messages, the corresponding policy rule applies, and *rule* specifies the actions (e.g., access, signature or encryption [6]) authorized for the policy *subject* (user), on the policy *object*. Consider for instance the XML-based security rules in Fig. 11. The first rule allows service points with role 'booking agency' to access encrypted credit card numbers of client requests, whereas the second rule denies subjects with role 'customer' from accessing credit card numbers of other clients.

```

1 <subject><role>BookingAgency</role></subject>
  <object>//BookingConfirmation/CreditCardNb</object>
  <rule>
    <Access>Allowed</Access>
    <Encryption>AES</Encryption>
  </rule>

2 <subject><role>Customer</role></subject>
  <object>//BookingConfirmation/CreditCardNb</object>
  <rule>
    <Access>Denied</Access>
  </rule>

```

Fig. 11. Sample SOAP security policy rules (expressed in XML).

The need for evaluating WS-Security policies may introduce additional overhead, which in some cases dwarfs the latency of standard SOAP message processing. The results

of [37] show that WS-Security policy evaluation can cause: i) an increase in SOAP response time by a factor of 3 on average, ii) a substantial increase in network traffic (SOAP messages size) by a factor 6.9 in overall (regardless of the type of data, e.g., integer, double, string..., being exchanged). In this context, a few proposals have addressed the issue of improving SOAP security policy evaluation performance through improving other underlying techniques, namely parsing [45, 71], caching [76] and multicasting [6, 14]. Methods for improving SOAP parsing performance, e.g. [45, 71], consist in parsing and simultaneously processing the SOAP message for security evaluation, providing the de-serializer module with the parsed output message (or parts of the message) the destination client is allowed to access. Simultaneous parsing and security policy evaluation is undertaken via automata (cf. Section 3.1.2) which consider both the parser context and security context, at the same time, for each incoming SOAP message. In other words, security-enabled parser automata identify SOAP events (e.g., opening element tag, element text...) which correspond to classic parsing events, as well as their corresponding policy rules (e.g., authorization, signature or encryption schemes, allowing security processing), so as to process SOAP messages accordingly. These methods have been discussed in Section 3.1.2.

In [76], the authors investigate various techniques for WS-Security performance optimization, including digest-based caching, pre-hashing, and on-demand canonicalization. They propose to store the de-serialized objects of digitally signed XML messages in cache, and then match the IDs and digest hash values of inbound elements to the objects in the cache, to be retrieved and utilized in case of a cache hit. Similarly, the digest hash value for each signed element in the outbound message is stored in the cache, along with its serialized content, so as to re-serialize and re-hash (in subsequent message exchanges) only those objects which are different. The authors show that the digest-caching and pre-hashing methods reduce overhead by a factor of 3 to 4 [76], at the expense of increased memory use (which they do not experimentally quantify). The authors also investigate on-demand canonicalization [75] (i.e., re-canonicalizing contents only when the signature verification fails), and show that it effectively improves performance when more than 88% of the WS-Security messages need not be re-canonicalized (otherwise, it might introduce additional overhead) [76].

Approaches in [6, 14] discuss and compare different scenarios where SOAP multicasting, namely SMP [59], could improve policy evaluation performance. In [14], the authors focus on a single sender/receiver SOAP message exchange scenario. They discuss how policy evaluation could be performed on an aggregate SMP message so as to only repeat policy evaluation processing on the SMP common part section once. Following the authors, security policy evaluation would be only repeated on those parts of the SOAP messages which are distinctive, inducing a substantial gain in processing time. In a subsequent study [6], the authors extend their discussion to multiple scenarios, with multiple senders/receivers, and investigate different approaches to improve SOAP signing/encryption through multicasting. They discuss different strategies for achieving optimal ordering of signing and multicasting operations, such as *Sign-Join-Split-Verify* and *Join-Sign-Split-Verify*. Fig. 12 depicts the classic approach, and the one ultimately adopted by the authors. They conclude that the best strategy, minimizing processing

time and thus maximizing the gain in performance, would be to i) first aggregate the SOAP messages (*Join*), ii) process the aggregate SMP message for signing/encryption (*Sign*), iii) transmit the signed/encrypted aggregate message to the receiver where it is first checked w.r.t. the latter's policy rules and processed for signature recognition and decryption (*Verify*), and then iv) decompose the SMP message to reconstruct the original SOAP messages (*Split*, cf. Fig. 12.b).

Fig. 12. Different scenarios to security policy evaluation.

Experimental results to quantify the actual gain in performance are not provided in [6], the corresponding prototypical implementation being under development. Indeed, research on the interplay between WS-Security policy evaluation and SOAP multicasting is still at a preliminary stage.

3.4 Parallelization and Hardware Approaches

Despite of the various kinds of software optimizations to improve SOAP and XML processing performance, no parser software can process input faster than its supporting hardware accesses data. With most current XML software toolkits, the maximum processing rate usually attains a best of tens of clock cycles per character [39] (a simple character-scanning loop runs at about 100 Mbytes/second on a 1 GHz Pentium processor, which amounts to 10 cycles/byte [39]), and that for many XML applications can result in processing rates of the order of hundreds of clock cycles per character (traditional parsers, e.g., [5, 65], perform in the range of 2.5–6 Mbytes of input per second or 160–400 cycles/byte, with a penalty of between 16x and 40x [39]). Recent benchmarking works in [32, 33] demonstrate that most existing implementations of WS do not scale well when the size of the SOAP/XML document being processed is increased. The authors in [32, 33] argue that most existing software toolkits are typically designed to process small-sized XML datasets, and thus are not suited for large-scale computing applications, e.g., [25, 62]. Hence, recent studies have attempted to alleviate the limitations of XML software performance bottlenecks by applying non-traditional parallel processor architectures, e.g., [8, 23, 30, 36, 55, 78]. On one hand, general-purpose (scalar) processors are characterized by the sequential nature of instruction execution, where instructions are selected based on their sequential memory addresses, conditions being evaluated one at a time. On the other hand, XML processing usually requires the evaluation of multiple conditions of various types that can occur simultaneously, namely during XML string and character parsing (e.g., verifying character integrity, whether an end tag matches a previously processed start tag, whether an attribute name is unique for a given element, and so on). Hence, the nature and frequency at which XML processing conditions occur result in a less predictable instruction flow,

which calls for higher processing parallelism to improve performance [8, 78].

Parallel processing solutions can be roughly classified according to the level at which the hardware supports parallelism [13], namely: bit-level, data-level, and instruction-level. In addition to single-node parallelism, a.k.a. micro-parallelism (achieved on a single computer system, with multiple processing units connected via the same bus and sharing the same memory), recent XML-related studies [23, 30, 31] have addressed cluster computing, a.k.a. macro-parallelism (i.e., distributed computing on large datasets of computer clusters). In the following, we provide a concise overview of the most prominent XML and SOAP parallel processing methods in the literature, roughly organized following the type of parallelism they achieve.

Bit-Level Parallelism: It consists in increasing the processor word size (i.e., the amount of bits the processor can manipulate per cycle) and optimizing the inner-processor architecture so as to reduce the number of instructions the processor must execute to perform operations on variables whose sizes are greater than the length of the word, and thus gain in processing rate. In this context, the authors in [78] introduce ZUXA, an XML accelerator engine which provides a processing model optimized for conditional execution in combination with dedicated instructions for XML character and string-processing functions. It is based on a programmable XML Finite State Machine technology, B-FSM, specifically tailored to provide high XML processing performance (a processing rate of one state transition per clock cycle), wide input and output vectors (with words of at least 64 bits for each transition), storage efficiency (to allow cost-efficient use of fast on-chip memory technologies), as well as full programmability (supporting fast incremental updates, allowing dynamic addition/removal of states and transitions), and scalability to tens of thousands of states and state transition rules. Related hardware solutions have been developed in the industrial arena, e.g., Datapower [16], which exploits Just-In-Time virtual machine technology [40] and ASICs customized for XML processing.

Data-Level Parallelism: Also known as SIMD (Simple Instruction Multiple Data), data-level parallelism describes computer systems with multiple processing elements that perform the same operation on multiple data simultaneously. An application that may take advantage of data-level parallelism is one where the same operation is being executed on a large number of data points, which is a common operation in many multimedia applications (e.g., image/video rendering and filtering), as well as in XML parsing and lexical analysis (e.g., reading input characters, and identifying string tokens). Parabix [8] is an XML parser designed to exploit the data-level parallelism capabilities of modern processors to deliver performance improvements over traditional byte-at-a-time parsing technology. Byte-oriented character data is first transformed to a set of 8 parallel bit streams, each stream comprising one bit per character code unit. Character validation, transcoding, and lexical item stream formation are all then carried out in parallel using bitwise logic and shifting operations. Byte-at-a-time scanning loops in the parser are replaced by bit scan loops that can advance by as many as 64 positions with a single instruction. Experimental results in [8] show that Parabix performs substantially better

than traditional XML parsers: ranging from twice as fast as Expat [65], to an order of magnitude faster than Xerces [5].

Instruction-Level Parallelism: It is a processing paradigm which underlines the re-ordering and combination of instructions into instruction sets, which are then executed in parallel without affecting the result of the program. Instruction-level parallelism could be achieved in a number of ways to improve XML parsing performance, namely through i) pipelining, and/or ii) multi-processing (a.k.a. superscalar computing) [13]. On one hand, pipelining allows splitting the processing of an instruction into a series of independent steps, executed in parallel by different threads. On the other hand, multi-processing allows the execution of more than one instruction during a clock cycle, by simultaneously dispatching multiple instructions to redundant execution units on the processor. Superscalar processors are identified as *multi-core* when their constituent processing units are embedded in the same processor chip. While pipelining may provide significant speedup, XML software pipelining is often hard to implement due to synchronization and memory access bottlenecks, and to the difficulties of balancing the pipeline stages [55]. Hence, most studies in the context of XML and WS have focused on multi-processing solutions. One prominent approach is the Meta-DFA project [43, 56], introducing a parallelization method that uses a two-stage DOM parser. The main idea is to divide the XML document into chunks, such as multiple threads would work on the chunks independently. The first stage consists in *pre-parsing* the XML document, to determine its logical tree structure (made of start and end tag node references). This structure is then used in a subsequent stage to divide the XML document such that the divisions between the chunks occur at well-defined points in the XML grammar. As the chunks are parsed, the results are then merged. In a following study [55], the authors investigate static partitioning and load-balancing in order to minimize thread synchronization overhead. The authors in [43, 55, 56] show that their technique, while effective, does not scale to large numbers of cores (from 1 to 4 cores). In addition, while DOM-style parsing can be intuitive and convenient with applications requiring random access/manipulation of XML-based data, nonetheless, it can also be memory-intensive, both in the amount of memory used (to store the DOM structure), and in the high overhead of memory management [55].

In a related project by Head *et al.*, the Piximal toolkit [23, 30, 31] presents a parallelized SAX parsing solution, focusing on a different class of applications than the DOM-based Meta-DFA project, tailored around event-streams and fast sequential access of XML-based data. Piximal conducts parsing work dynamically, and generates as output a sequence of SAX events. This results in a larger number of parser states and state transitions, underlining more opportunities for parallelization optimization, and scaling well with increasing numbers of processing cores. Experimental results demonstrate that the level of speedup obtainable using Piximal's micro-level parallelization techniques can be limited due to: i) memory bandwidth, which could become a bottleneck [31], and ii) the amount of computation required to parse the input, which would induce little performance gain if the computation required is small in comparison to the time required to access the bytes of the input in memory [23]. Hence, the authors in [23, 30, 31] also address macro-level

parallelism. They investigate the distributed processing of large-scale XML data stored in a cluster, by applying Google's MapReduce processing paradigm [18]. The simplicity and robustness of the MapReduce model, as well as its relaxed synchronization constraints, tend to work favorably for large-scale XML data sets and WS computing environments [23]. Experimental results on Piximal's macro-level parallelization technique show that securing additional resources for each thread by distributing the workload to a cluster of machines using MapReduce can increase performance [23, 30, 31]. Nonetheless, the authors also show that if not enough processing is taking place on each cluster, the latter would be burdened with redundancy checks and network traffic for just small chunks of input. The authors conclude that when computation is not sufficient enough to offset communication latencies due to the number of running computers, a single node, which minimally suffers from the same condition, would perform better than a cluster of computers.

4 ONGOING CHALLENGES

Despite the wide array of techniques proposed to enhance SOAP processing performance, yet various challenges and limitations remain unaddressed. Three major hurdles remain to the wide adoption of similarity-based techniques.

First, while similarity-based methods have been shown in many cases to produce a significant gain in speed-up when many similar messages are involved [69], as well as a noticeable reduction in network traffic [58], nonetheless, similarity computations can sometimes introduce additional overhead on their own (as shown with SOAP compression [81] and multicasting [58, 59]), especially when the SOAP messages being processed are fairly different (i.e., not similar to the documents processed before). Hence, a comprehensive empirical analysis addressing the trade-off between: i) the amount of additional processing overhead, and ii) the amount of processing time and network traffic reduction, induced by similarity-based approaches, is required in order to identify and better understand each method's optimum usage constraints (e.g., percentage of similar SOAP messages, amount of inner-message similarities, number of messages, and so on).

Secondly, interference and synergy between different similarity-based techniques is not yet completely understood. One can realize that the various techniques covered in the paper are not mutually exclusive, but are rather complementary. For instance, similarity-based methods to SOAP serialization, parsing, and de-serialization could very well exploit XML parallel processing architectures so as to better improve their clock cycle character processing rates. In addition, software-based methods could make use of tight integration architectures, such as in [39], so as to avoid repeated/unnecessary data processing, copying to/from memory buffers, and expensive data-type transformations (ASCII/UTF to in-memory types, and vice-versa). In this context, recent efforts have been made toward combining efficient SOAP multicasting, on one hand, with fast security policy evaluation on the other hand (as discussed in Section 3.3). Nonetheless, corresponding techniques are still in their preliminary stages. Comparative theoretical and experimental studies are required to better understand the interplay and actual gain in performance between WS-Security policy

evaluation and SOAP multicasting.

TABLE 1.
Characteristics of Existing (Similarity-based) SOAP Performance Enhancement Approaches.

Performance	SOAP Processing	Approach	Features
Reducing Response time and increasing Throughput	Serialization	Abu-Ghazaleh <i>et al.</i> [4]	bSOAP, differential serializer: - DUTs (Data Update Tracking), tracking between in-memory data, and their serialized representations. - <i>Dirty bits</i> to identify fields whose values changed, recognizing parts to be reused.
		Abu-Ghazaleh <i>et al.</i> [2, 3]	bSOAP buffer management: - Padding and chunk overlaying to allow on-the-fly message expansion.
		Devaram and Andersen[21]	Client-side SOAP message caching: - Indexing structures to detect correspondences between cached and outgoing messages. - Does not address partial structural matches (only caches identical structures).
	Parsing	Zhang and Van Engelen[87]	TDX: Table Driven XML parsing - Combining the lexical analysis and validation - Pre-recording parser states as grammar productions in tabular form, and breaking up the SOAP message into a token stream
		Takeuchiet <i>al.</i> [70]	T-SOAP, template-based differential parser: - Predefined template, modeled via a finite state automaton (FSA). - Identification of invariant/variable tag parts in the SOAP messages. - Variable parts are only parsed.
		Makino <i>et al.</i> [45]	Multi-template differential parser: - Appending new templates to the FSA, - More flexible than T-SOAP [70] (bound to one single template), - Requires more memory than T-SOAP.
		Teraguchi <i>et al.</i> [71]	Detecting repeatable structures: - Improved XML-based automaton, to consider repeatable structures in SOAP messages, in comparison with string-based ones in [45, 70], - More expressive automaton, reducing memory and time consumption.
	De-Serialisation	Kostoulas <i>et al.</i> [39]	XML Screamer: - Tight integration across software levels, - Combines parsing and de-serialization in one layer, so as to avoid unnecessary data processing, copying (to/from memory), and data-type transformation.
		Suzumura <i>et al.</i> [68]	Automaton-based approach: - Classic de-serialisation and automaton creation, - Matching messages to automaton and only de-serialising those different portions (could complement parsers in [45, 70, 71])
		Abu-Ghazaleh and Lewis[1]	Checksum-based approach: - Regular mode, periodically checkpointing de-serialiser state, - Compare checkpoints, and switches to fast mode, when parser state is similar to state saved in previous checkpoint, - Checksumming is fast, yet error prone.
		Makino <i>et al.</i> [45], Teraguchiet <i>al.</i> [71]	Security-based SOAP message parsing: - Automata to consider both the parser context and security context, - Identifying SOAP events (tags, text...) and their corresponding policy rules (authorizations, signatures...)
Reducing Network traffic	Security Policy Evaluation	Damiani and Marrara[14]	Security-based SOAP multicasting: - Single sender-receiver scenario, - Policy evaluation on aggregate SMP message [59], - Policy evaluation repeated only on those parts of SOAP messages which are different.
		Azzini <i>et al.</i> [6]	Security-based SOAP multicasting: - Multiple senders/receivers scenario - Different approaches to improve SOAP signature/encryption (<i>Sign-Join-Split-Verify</i> , <i>Join-Sign-Split-Verify</i> ...), - Best strategy is <i>join-sign-verify-split</i> .
		Van Engelen and Zhang[76]	WS-Security performance optimization: - Digest-based caching, storing and using de-serialized digitally signed objects, - Pre-hashing, storing and using digest values of digitally signed objects, - On-demand canonicalization, re-canonicalizing contents only when the signature verification fails.
	Compression	Werneret <i>al.</i> [81]	Differential compression: - XML differential encoding (tree edit distance), - Identifying differences between SOAP messages and predefined WSDL-based SOAP templates, - Only differences are transmitted, - Patching differences with the same skeleton at the receiver side, to reconstruct the original message.
	Multicasting	Phanet <i>al.</i> [59]	SMP, Similarity-based SOAP Multicasting Protocol: - Built on top of IP unicast (avoiding complex network configurations), - Grouping and transmitting together similar SOAP messages (not only identical ones such as with classic multicasting), - SMP message encapsulated in classic SOAP message, with common and distinct parts.
		Phan <i>et al.</i> [58]	tc-SMP, traffic constrained SMP: - Enhanced routing protocol for transmitting messages following paths which maximize shared links between highly similar messages, - Reducing traffic in comparison with the OSPF-based SMP [59].

TABLE 2.
Characteristics of SOAP and XML-based Parallellization and Hardware related approaches.

Performance	SOAP Processing	Approach	Features
Micro-Parallelism	Bit-level	Van Lunteren <i>et al.</i> [78]	ZUXA XML Accelerator Engine: <ul style="list-style-type: none"> - Increasing processor word size, i.e., the amount of bits the processor can manipulate per cycle, - Optimized for conditional execution with dedicated instructions for XML character processing, - Based on a programmable State Machine technology, B-FSM, tailored to provide high XML processing performance, wide input/output vectors, storage efficiency, as well as full programmability.
	Data-level	Cameron <i>et al.</i> [8]	PARABIX: <ul style="list-style-type: none"> - Designed to exploit the data-level parallelism, - Byte-oriented character data is first transformed to a set of 8 parallel bit streams, each stream comprising one bit per character code unit, - Character validation, transcoding, and lexical item stream formation are all then carried out in parallel using bitwise logic and shifting operations.
	Instruction-level	Pan <i>et al.</i> [43, 56]	Meta-DFA: <ul style="list-style-type: none"> - Two-stage DOM parser : i) <i>pre-parsing</i> to determine its logical XML tree structure, and then ii) dividing the XML document such that the divisions between the chunks occur at well-defined points in the XML grammar, - Merges results as the chunks are parsed, - Exploits static partitioning and load-balancing to minimize thread synchronization overhead, - Scalable to a maximum of 4 cores.
		Head <i>et al.</i> [23, 30, 31]	Piximal: <ul style="list-style-type: none"> - Introduces a parallelized SAX parser, tailored around event-stream XML data (different class of applications than the DOM-based Meta-DFA), - Larger number of parser states, thus more opportunity for parallelization and scalability with increasing numbers of cores (in comparison with Meta-DFA), - Speed-up could be limited due to: i) memory bandwidth, and ii) the amount of computation required to parse the input (if the computation required is small in comparison to the time required to access the bytes of the input in memory).
Macro-Parallelism		Head <i>et al.</i> [23, 30, 31]	Piximal, with cluster computing: <ul style="list-style-type: none"> - Exploits distributed processing of large-scale XML data stored in a cluster, by applying Google's MapReduce processing paradigm [18], - Introduces relaxed synchronization constraints, which tend to work favorably for large-scale XML data sets and WS computing environments, - Experiments show that macro-parallelism can increase performance (in comparison with micro-parallelism). Yet, if not enough processing is taking place on each cluster, the latter would be burdened with redundancy checks and network traffic for just small chunks of input, and could perform worst than a single node, - Examining computation costs to determine the best computation strategy.

Thirdly, and perhaps more importantly, interference may arise between SOAP similarity-based multicasting described in this paper and attempts at boosting SOAP performance via custom protocol bindings.

Several commercial SOAP engines, including Noemax and Sun Metro, are based on custom protocol bindings that exploit information on the XML stream data to improve the performance of transport layer protocols. In these implementations of SOAP, HTTP binding has been dropped altogether in favor of an integrated SOAP/TCP transport where each message sent during a communication session is accompanied only by new entries (if any) to the XML Infoset vocabulary[67]. The vocabulary is a table that associates string values with identifiers. In this context, the technique used to reduce the size of the XML text encoding is to enter string values (such as XML markup) in the vocabulary and substitute all occurrences of these string values in the document with their corresponding identifier. This vocabulary-based technique is sometime coupled with GZIP compression [20] of messages, and is a major competitor of similarity-based multicasting when non-standard protocol bindings are acceptable - e.g., on clusters or grids [80] when no firewall traversal is required. However, the effect of using similarity-based SOAP multicasting in the context of custom SOAP/TCP bindings is still largely unexplored, but, great potential have been shown by enhancements in the underlying HTTP transport protocol (particularly in the context of HTTP 1.1) to reduce the overhead of creating a new connec-

tion for every SOAP message (with persistent connections and message chunking [12, 28]), as well as by ongoing investigations in XML-based binary encodings for SOAP [57, 64, 83]. In short, techniques to SOAP performance enhancement are yet to be further improved and perfected, promising further performance improvements in the near future, which presents an overwhelming motivation to do research in this field.

5 CONCLUSION

In this survey paper, we have given an overview of current research related to SOAP processing performance enhancement, focusing on similarity-based approaches, as well as WS-Security optimizations, and XML parallel processing architectures. We provide a concise, yet comprehensive review of how different techniques have been exploited to enhance SOAP performance in almost every phase of SOAP processing, ranging over message parsing [45, 70, 71], serialization [4, 21], de-serialization [1, 68], compression [81], multicasting [6, 58, 59], security evaluation [6, 14], and data/instruction-level processing [8, 55, 78](cf. Tables 1 and 2). Most methods build on the observation that SOAP message exchange usually involves highly similar messages (messages created by the same implementation usually have the same structure, and those sent from a server to multiple clients tend to show similarities in structure and content).

The main idea is then to identify the common parts of SOAP messages, to be processed once, only repeating the processing for parts which are different, and substantially reducing SOAP processing overhead. Other approaches investigate non-traditional processor architectures, including micro- and macro-level parallel processing solutions, so as further increase the processing rates of SOAP/XML software toolkits. In addition, we have also discussed some of the main challenges and possible future research directions, covering SOAP software and parallel architecture integration, as well as custom protocol bindings.

We hope that the unified presentation of SOAP-related performance enhancement techniques in this paper will foster further research on the subject.

ACKNOWLEDGMENT

This work was supported in part by *Fondazione Cariplo2007Capitale Umano di Eccellenza* research grant, and *Japan Society for the Promotion of Science (JSPS)2010* research fellowship n. PE10006.

References

- [1] Abu-Ghazaleh N. and Lewis M.J., *Differential Deserialization for Optimized SOAP Performance*. Proceedings of the ACM/IEEE Conference on Supercomputing, 2005. pp. 21-31, Seattle, USA.
- [2] Abu-Ghazaleh N., M.J.L., and M. Govindaraju. , *Performance of Dynamic Resizing of Message Fields for Differential Serialization of SOAP Messages*. Proceedings of the International Symposium on Web Services and Applications, 2004. pp. 783-789.
- [3] Abu-Ghazaleh N.; Govindaraju M. and Lewis M.J., *Optimizing Performance of Web Services with Chunk-Overlaying and Pipelined-Send*. . Proceedings of the International Conference on Internet Computing (ICIC), 2004. pp. 482-485.
- [4] Abu-Ghazaleh N.; Lewis M.J. and Govindaraju M., *Differential Serialization for Optimized SOAP Performance*. Proceedings of the 13th International Symposium on High Performance Distributed Computing (HPDC'04), 2004. pp. 55-64.
- [5] Apache Foundation. *Xerces XML Parser*. Available from <http://xerces.apache.org/>, [cited Nov 2010].
- [6] Azzini A.; Marrara S.; Jensen M. and Schwenk J., *Extending the Similarity-Based XML Multicast Approach with Digital Signatures*. Proceedings of the 2009 ACM Workshop on Secure Web Services (SWS'09), 2009. pp. 45-52, Chicago.
- [7] Bray T.; Paoli J.; Sperberg-McQueen C.; Mäler Y.; and Yergeau F. *Extensible Markup Language (XML) 1.0 - 5th Edition*. W3C recommendation, 26 Novembre 2008. <http://www.w3.org/TR/REC-xml/> [cited November 2008].
- [8] Cameron R.D.; Herdy K.S. and Lin D., *PARABIX: High Performance XML Parsing using Parallel Bit Stream Technology*. . In Proceedings of the 2008 Conference of the Center for Advanced Studies on Collaborative Research: Meeting of Minds (CASCON '08), 2008. 17:222-235, ACM, New York, NY, USA.
- [9] Cheney J., *Compressing XML with Multiplexed Hierarchical PPM Models*. In Proceedings of the Data Compression Conference, 2001. pp. 163-173.
- [10] Chinnici R.; Moreau J.J.; Ryman A. and Weerawarana S. *Web Services Description Language (WSDL) Version 2.0 Part 1: Core Language*, W3C Recommendation 26 June 2007, <http://www.w3.org/TR/wsdl20/>. [cited 25 August 2009].
- [11] Chiu K. and Lu W., *A Compiler-based Approach to Schema-specific XML Parsing*. Proceedings of the Workshop on High Performance XML Processing, New York., 2004.
- [12] Chiu K.; Govindaraju M. and Bramley R., *Investigating the Limits of SOAP Performance for Scientific Computing*. Proceedings of ACM International Symposium on High Performance Distributed Computing (HPDC), 2002. pp. 246-254, Edinburgh, Scotland.
- [13] Culler D.E.; Singh J.P. and Anoop Gupta, *Parallel Computer Architecture - A Hardware/Software Approach*. 1999. Morgan Kaufmann Publishers, pp. 1100, ISBN 1-55860-343-3.
- [14] Damiani E. and Marrara S., *Efficient SOAP Message Exchange and Evaluation Through XML Similarity*. Proceedings of the 2008 ACM workshop on Secure Web Services (SWS'08), 2008, 29-36.
- [15] Damiani E.; De Capitani di Vimercati; Paraboschi S. and Samarati P., *Securing SOAP E-Services*. International Journal of Information Security (IJIS), 2001. 1:100-115.
- [16] Datapower. <http://www.datapower.com/>. [cited Nov 2010].
- [17] Davis D. and Parashar M., *Latency Performance of SOAP Implementations*. Proceedings of the 2nd IEEE/ACM International Symposium on Cluster Computing and the Grid, 2002. pp. 407-412.
- [18] Dean J. and Ghemawat S., *MapReduce: Simplified Data Processing on Large Clusters*. Communications of the ACM, 2008. 51(1):107-113.
- [19] Della-Libera G. et al., *Web Services Security Policy Language (WS-SecurityPolicy)*. V1.1 Specification, July 2005. <http://download.boulder.ibm.com/ibmdl/pub/software/dw/specs/ws-secpol/ws-secpol.pdf>.
- [20] Deutsch L.P., RFC 1952: *GZIP file format specification version 4.3*. 1996.
- [21] Devaram K. and Andersen D., *SOAP Optimization via Parameterized Client-Side Caching* Proceedings of the IEEE/ACM 2nd International Symposium on Cluster Computing and the Grid (CCGRID'02), 2002. pp.439-312.
- [22] Elfving R.; Paulsson U. and Lundberg L., *Performance of SOAP in Web Service Environment Compared to CORBA*. Proceedings of the 9th Asia-Pacific Software Engineering Conference (APSEC'02), 2002. pp. 84-94.
- [23] Fadika Z.; Head M.R. and Govindaraju M., *Parallel and Distributed Approach for Processing Large-Scale XML Datasets*. In Proc.of 10th IEEE/ACM International Conference on Grid Computing (GRID 2009), 2009, 105-112, Banff, Canada.
- [24] Fielding R.; Gettys J.; Mogul J.; Frystyk H.; Masinter L.; Leach P.; Berners-Lee T., N.W.G. *Hypertext Transfer Protocol -- HTTP/1.1*, <http://www.ietf.org/rfc/rfc2616.txt>, 1999, [cited May 2010].
- [25] Gannon D.; Krishnan S.; Fang L.; Kandaswamy G.; Simmhan Y. and Slominski A., *On Building Parallel and Grid Applications: Component Technology and Distributed Services*. In Proc. of the 2nd Inter. Workshop on Challenges of Large

- Applications in Distributed Environments (CLADE '04), 2004. IEEE Computer Society, p. 44, Washington DC, USA.
- [26] Gao J.Z.; Tsao H.S.J. and Wu, Y., *Testing and Quality Assurance for Component-based Software*. Artech House, 2003, 439.
- [27] Ginige A. and Murugesan S., *Web Engineering: An Introduction*. IEEE Multimedia, 2001. 8(1):14-17.
- [28] Govindaraju M.; Slominski A.; Chiu K.; Liu P.; Van Engelen R.; Lewis M.J., *Toward Characterizing the Performance of SOAP Toolkits*. Proceedings of 5th IEEE/ACM International Workshop on Grid Computing (GRID'04), Pittsburgh, 2004. pp. 365-372.
- [29] Gudgin M.; Hadley M.; Mendelsohn N.; Moreau J.-J.; Canon and Nielsen H.F. *Simple Object Access Protocol 1.1* <http://www.w3.org/TR/SOAP>. June 2003 [cited April 2010].
- [30] Head M.R. and Govindaraju M., *Parallel Processing of Large-Scale XML-Based Application Documents on Multi-core Architectures with PiXiMaL*. In Proceedings of the 4th IEEE International Conference on e-Science, 2008. pp. 261-268, Indianapolis, Indiana, USA.
- [31] Head M.R. and Govindaraju M., *Performance Enhancement with Speculative Execution Based Parallelism for Processing Large-scale XML-based Application Data*. In Proceedings of International Symposium on High Performance Distributed Computing (HPDC 2009), 2009, 21-30, Munich, Germany.
- [32] Head M.R.; Govindaraju M.; Slominski A.; Liu P.; Abughazaleh N.; Van Engelen R.; Chiu K. and Lewis M.J., *A Benchmark Suite for SOAP-based Communication in Grid Web Services*. In Proceedings of the ACM/IEEE Conference on Supercomputing (SC'05), 2005. pp. 19.
- [33] Head M.R.; Govindaraju M.; Van Engelen R. and Zhang W., *Benchmarking XML Processors for Applications in Grid Web Services*. In Proceedings of the ACM/IEEE Conference on Supercomputing (SC'06), 2006. pp. 30.
- [34] Hopcroft J. E.; Motwani R. and Ullman J. D., *Introduction to Automata Theory, Languages, and Computation*. 2001. Addison Wesley, 2nd edition.
- [35] Horstmann M. and Kirtland M. *DCOM Architecture*. Microsoft MSDN, <http://msdn.microsoft.com/en-us/library/ms809311.aspx> 1997 [cited January 2010].
- [36] Intel Corporation. *Intel Core i7-800 Processor Series and the Intel Core i5-700 Processor Series*. [cited Nov 2010], <http://download.intel.com/products/processor/corei7/>
- [37] Juric M.B.; Rozman I.; Brumen B.; Colnaric M. and HerickoM., *Comparison of Performance of Web Services, WS-Security, RMI, and RMI-SSL*. Journal of Systems and Software, 2006. Volume 79, Issue 5, 689-700.
- [38] Kohlhoff C. and Steele R., *Evaluating SOAP for High Performance Business Applications: Real-Time Trading Systems*. Proceedings of the World Wide Web (WWW) Conference, 2003. Budapest, Hungary.
- [39] Kostoulas M. G.; Matsa M.; Mendelsohn N.; Perkins E.; Heifets A. and Mercaldi M., *XML Screamer: An Integrated Approach to High Performance XML Parsing, Validation and Deserialization*. In Proceedings of the 15th International Conference on World Wide Web (WWW '06), 2006. pp. 93-102.
- [40] Kuznetsov E., *Method and Apparatus of Data Exchange Using Runtime Code Generator and Translator*. US Patent 6,772,413 B2, 2004.
- [41] Laux A. and Martin L., *XUUpdate Working Draft*. XML:DB Initiative, 2000.
- [42] Liefke H. and Suciu D., *XMill: An Efficient Compressor for XML Data*. University of Pennsylvania Technical Report MSCIS-99-26., 2000.
- [43] Lu W.; Chiu K. and Y. Pan, *A Parallel Approach to XML Parsing*. In Proceedings of the 7th IEEE/ACM International Conference on Grid Computing (Grid'06), 2006. pp. 223-230.
- [44] Ma Y. and Chbeir R., *Content and Structure Based Approach for XML Similarity*. Proceedings of the International Conference on Computer and Information Technology (ICCIT), 2005. pp. 136-140.
- [45] Makino S.; Tatsubori M.; Tamura K. and Nakamura Y., *Improving WS-Security Performance with a Template-Based Approach*. Proceedings of the IEEE International Conference on Web Services (ICWS'05), 2005. pp. 581-588.
- [46] Martin B. and Jano B. *WAP Binary XML Content Format*. W3C Note 24 June 1999 1999 [cited February 2010].
- [47] Megginson D. et al. *The Simple API for XML* <http://www.megginson.com/SAX/> [cited February 2010].
- [48] Menascé D.A. and Almeida V.A.F., *Capacity Planning for Web Services – Metrics, Models and Methods*. 2002. p.556, Prentice Hall.
- [49] Menascé D.A.; Almeida V.A.F. and Dowdy L.WL., *Capacity Planning and Performance Modeling: From Mainframes to Client-Server Systems*. Prentice Hall, Upper Saddle River, New Jersey 1994.
- [50] Moon T.K., *Error Correction Coding: Mathematical Methods and Algorithms*. New Jersey: John Wiley & Sons, 2005, 756.
- [51] Mouat A., *XML Diff and Patch Utilities*. CS4 Dissertation., 2002. Edinburgh Scotland: Heriot-Watt University.
- [52] Moving Pictures Experts Group. *MPEG-7*. <http://www.chiariglione.org/mpeg/standards/mpeg-7/> [cited 2 June 2009].
- [53] Noga M. L.; Schott S. and Lowe W., *Lazy XML Processing*. In Proceedings of the 2002 ACM Symposium on Document Engineering (DocEng '02), 2002. Virginia, USA.
- [54] Object Management Group. *The Common Object Request Broker: Architecture and Specification*. Version 3.0.3, http://www.omg.org/technology/documents/formal/corba_2.htm, 2004 [cited January 2010].
- [55] Pan Y.; Lu W.; Zhang Y. and Chiu K., *A Static Load-Balancing Scheme for Parallel XML Parsing on Multicore CPUs*. In Proceedings of the 7th IEEE International Symposium on Cluster Computing and the Grid (CCGrid '07), 2007. pp.351-362.
- [56] Pan Y.; Zhang Y.; Chiu K. and Lu W., *Parallel XML Parsing Using Meta-DFAs*. In Proceedings of the IEEE Third International Conference on eScience and Grid Computing (eScience'07), 2007 pp. 237-244.
- [57] Paul Sandoz et al. *Fast Web Services*, 2003 [cited May 2010], <http://java.sun.com/developer/technicalArticles/WebServices/>
- [58] Phan K.A.; Bertok P.; Fry A. and Ryan C., *Minimal Traffic-Constrained Similarity-Based SOAP Multicast Routing Protocol*. OTM Confederated International Conferences, 2009. LNCS 4803, pp. 558-576.
- [59] Phan K.A.; Tari Z.; and Bertok P., *Similarity-Based SOAP Multicast Protocol to Reduce Bandwidth and Latency in Web*

- Services. IEEE Transactions on Services Computing, 2008. Vol 1, No 2, pp. 88-103.
- [60] Rud D.; Schmietendorf A. and Dumke, R., *Product Metrics for Service-Oriented Infrastructures*. In Abran A., Bundschuh M., Buren G., Dumke, R., eds.: *Applied Software Measurement*. . Proceedings of the International Workshop on Software Metrics and DASMA Software Metrik Kongress (IWSM/MetriKon'06). 2006. pp.161–174.
- [61] Sahai A. and Machiraju V., *Enabling fo the Ubiquitous e-services Vision on the Internet* Hewlett-Packard Laboratories, HPL-2001-5, 2001.
- [62] Singh G.; Bharathi S.; Chervenak A.; Deelman E.; Kesselman C.; Manohar M.; Patil S. and Pearlman L., *A Metadata Catalog Service for Data Intensive Applications*. In proceedings of the 2003 ACM/IEEE conference on Supercomputing., 2003. IEEE Computer Society, 2003, p. 33, Washington DC, USA.
- [63] Slominski A. XSOAP. 2004, [cited February 2010]<http://www.extreme.indiana.edu/xgws/xsoap/>
- [64] SourceForge.NET. XML Binary Information Set (XBIS). [cited May 2010], Available from: <http://xbis.sourceforge.net/>.
- [65] SourceForge.NET. *The Expat XML Parser*. Available from <http://expat.sourceforge.net/> [cited Oct 2010].
- [66] Sun. *Java Remote Message Invocation (RMI)*, [cited January 2010], <http://java.sun.com/j2se/1.5.0/docs/guide/rmi/> 2005
- [67] Sun Microsystems, *SOAP/TCP Specification v.1.0*. <http://java.sun.com/webservices/reference/apis-docs/soap-tcp-v1.0.pdf>, May 2007.
- [68] Suzumura T.; Takase T. and Tatsubori M., *Optimizing Web Services Performance by Differential Deserialization* Proceedings of the IEEE International Conference on Web Services (ICWS'05), 2005. Vol. 1, pp.185- 192.
- [69] Takase T.; Miyashita H.; Tatsubori M. and Suzumura T., *An Adaptative, Fast and Safe XML Parser Based on Byte Sequence Memorization*. Proceedings of the World Wide Web (WWW) Conference, 2005. pp. 692 - 701.
- [70] Takeuchi Y.; Okamoto T.; Yokoyama K. and Matsuda S., *A Differential-Analysis Approach for Improving SOAP Processing Performance*. Proceedings of the IEEE International Conference on e-Technology, e-Commerce and e-Service (EEE'05), 2005, 472-479.
- [71] Teraguchi M.; Makino S.; Ueno K. and Chung H.V., *Optimized Web Services Security Performance with Differential Parsing*. Proceedings of the 4th International Conference on Service-Oriented Computing (ICSOC'06), 2006. pp. 277-288.
- [72] Truong H.L; Dustdar S. and Fahringer T., *Performance Metrics and Ontologies for Grid Workflows*. Future Generation Computer Systems, 2007. 23:760–772.
- [73] Van Engelen R., *Pushing the SOAP Envelope with Web Services for Scientic Computing*. Proceedings of the International Conference on Web Services (ICWS), 2003. pp. 346-352.
- [74] Van Engelen R., *Constructing Finite State Automata for High Performance XML Web Services*. Proceedings of the International Conference on Internet Computing (ICIC), 2004. pp. 975-981.
- [75] Van Engelen R., *A framework for service-oriented computing with C and C++ Web service components*. ACM Transactions on Internet Technology (ACM TOIT), 2008. 8(3):1-25, New York, NY, USA.
- [76] Van Engelen R. and Zhang W., *An Overview and Evaluation of Web Services Security Performance Optimizations*. In Proceedings of IEEE International Conference on Web Services (ICWS), 2008. pp. 137-144.
- [77] Van Engelen R.A. and K. Gallivan K., *The gSOAP Toolkit for Web Services and Peer-To-Peer Computing Networks*. In Proceedings of the 2nd IEEE International Symposium on Cluster Computing and the Grid (CCGrid2002), 2002. pp. 128-135, Berlin, Germany.
- [78] Van Lunteren J.; Bostian J.; Carey B.; Engbersen T. and Larsson C., *XML Accelerator Engine The First International Workshop on High Performance XML Processing*, 2004. New-York, NY, USA.
- [79] Viega J.; Messier M. and Chandra P., *Network Security with OpenSSL*. O'Reilly, 2002.
- [80] Wang N.; Welzl M. and Zhang L., *A High Performance SOAP Engine for Grid Computing*. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, 2009. Volume 2, pp. 1-8.
- [81] Werner C.; Buschmann C. and Fischer S., *WSDL-Driven SOAP Compression*. International Journal of Web Services Research, 2005. Vol. 2, Issue 1, pp. 18-35.
- [82] Word Wide Web Consortium. *SOAP Version 1.2*. W3C Recommendation (Second Edition), 2007, [cited February 2010]. <http://www.w3.org/TR/soap/>
- [83] World Wide Web Consortium. *XML Binary Characterization Working Group*, [cited May 2010], Available from: <http://www.w3.org/XML/Binary/>.
- [84] World Wide Web Consortium. *The Document Object Model*. <http://www.w3.org/DOM/>, [cited 28 May 2009].
- [85] World Wide Web Consortium. *Scalable Vector Graphics (SVG)*. <http://www.w3.org/Graphics/SVG/>, [cited May 2009].
- [86] Zhang B.; Jamin S. and Zhang L., *Host Multicast: A Framework for Delivering Multicast to End Users*. Proceedings of the IEEE Conference on Computer Communications (INFOCOM'02), 2002. pp. 1366-1375.
- [87] Zhang W. and Van Engelen R. A., *A Table-Driven Streaming XML Parsing Methodology for High-Performance Web Services*. In Proceedings of the IEEE International Conference on Web Services (ICWS'06), 2006. pp. 197-204.

Joe M. Tekli is a visiting researcher at the Department of Science and Technology, University of Shizuoka, Japan (since May 2010), and is a former post-doc of the University of Milan, Italy (2009). He holds a PhD in CS from the University of Bourgogne, LE2I CNRS, France, acquired (in Oct. 2009) with Highest Honors. He also holds a Research Masters in CS from the University of Bourgogne (July 2006), and a Masters of Engineering in Telecommunications from the Antonine Fathers University, Lebanon (July 2005), both acquired with Honors (ranked top of his class in both programs). He has been awarded various prestigious postdoctoral fellowships, of

the FAPESP (Brazil), JSPS (Japan), and Fondazione IRI (Italy). He was also awarded a PhD Fellowship of the Ministry of Education (France), and a Masters Scholarship of the AUF (France). His research activities cover XML processing, web services, data semantics and taxonomies, data clustering and classification, RSS integration, and multimedia fragmentation. He is a member of IEEE and ACM SIGAPP French Chapter. He is an organizing member of various international conferences such as SITIS, ICDIM, MEDES and ACM SAC'06. His research results have been published in various international journals and conferences (e.g., Computer Science Review, WWW Journal, ER, SBB, WISE, ADBIS, COMAD, etc.).

Ernesto Damiani is a professor at Università degli Studi di Milano and the director of the same University PhD program in computer science. He has held visiting positions at a number of international institutions. He has done extensive research on advanced network infrastructure and protocols, taking part in the design and deployment of secure high-performance networking environments. His areas of interest include business process representation, Web services security, processing of semi and unstructured information, and semantics-aware content engineering for multimedia. He is interested in models and platforms supporting open source development. He has served and is serving in all capacities on many congress, conference, and workshop committees. He is a senior member of the IEEE. In 2008 he was nominated ACM distinguished scientist and he received the Chester Hall Award from the IEEE Society on Consumer Electronics. Web page www.dti.unimi.it/~damiani.

Dr. Richard Chbeir received his PhD in Computer Science from the University of INSA-FRANCE in 2001. The author became a member of IEEE since 1999. He is currently an Associate Professor in the Computer Science Department of the Bourgogne University, Dijon-France. His research interests are in the areas of distributed multimedia database management, XML similarity and rewriting, spatio-temporal applications, indexing methods, multimedia access control models, security and watermarking. Dr. CHBEIR has published (more than 80 peer-reviewed publications) in international journals and books (IEEE Transactions on SMC, Information Systems, Journal on Data Semantics, Journal of Systems Architecture, etc.), conferences (ER, WISE, SOFSEM, EDBT, ACM SAC, Visual, IEEE CIT, FLAIRS, PDCS, etc.), and has served on the program committees of several international conferences (ICDIM, IEEE SITIS, ACM SAC, IEEE ISSPIT, EuroPar, SBB, etc.). He has been organizing many international conferences and workshops (ICDIM, CSTST, SITIS, etc.). He is currently the Chair of the French Chapter ACM SIGAPP and the vice-chair of ACM SIGAPP.

Gabriele Gianini, Ph.D., is Assistant Professor at the Department of Information Technology of the University of Milan where he is lecturer of Probability and Statistics, and since 2005 Visiting Professor at the Free University of Bolzano. He has been working between 1990 and 2000 at the Fermi National Accelerator Laboratory (Fermilab) in Chicago and at the CERN in Geneva. He is involved

in several research projects funded by the Italian Ministry of Research and by the European Union.