

Evaluating SOAP for High Performance Business Applications: Real-Time Trading Systems

Christopher Kohlhoff
Tenermerx Pty Ltd
6/50 Ben Boyd Rd, Neutral Bay
NSW 2089 Australia
chris@tenermerx.com

Robert Steele
University of Technology, Sydney
PO Box 123 Broadway
NSW 2007 Australia
rsteele@it.uts.edu.au

ABSTRACT

Web services, with an emphasis on open standards and flexibility, may provide benefits over existing capital markets integration practices. However, web services must first meet certain technical requirements including performance, security and fault-tolerance. This paper presents an experimental evaluation of SOAP performance using realistic business application message content. To get some indication of whether SOAP is appropriate for high performance capital markets systems, the results are compared with a widely used existing protocol. The study finds that, although SOAP performs relatively poorly, the difference is less than in scientific computing environments. Furthermore, we find that in realistic business applications it is possible for text-based wire formats to have comparable performance to binary, and that the text-based nature of XML is not sufficient to explain SOAP's inefficiency. This suggests that further work may enable SOAP to become a viable wire format for high performance business applications.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols—*applications*; C.2.4 [Computer-Communication Networks]: Distributed Systems—*distributed applications*; K.4.4 [Computers and Society]: Electronic Commerce—*electronic data interchange*

General Terms

Performance

Keywords

Performance study, SOAP, Web services, FIX

1. INTRODUCTION

Over the last decade, rapid advances in computing technology have driven dramatic changes in the financial sector. The provision of online services like real-time share trading has prompted a move to new business models such as global markets, 24-hour trading, and straight-through processing [21]. Tremendous growth in the number of third party financial networks has increased competition and accelerated the race to develop advanced, automated trading systems [7].

To compete effectively, existing organisations need to form alliances and provide integrated services. However, business-to-business integration is not a new concept for capital markets. The

financial domain has utilised industry standard protocols for the integration of distributed applications since the 1970s [26].

Web services are now emerging as a technology for systematic and flexible application-to-application integration [5]. Web services differ from most existing integration practices in that they utilise established, proven web protocols and open XML standards.

Capital markets systems may benefit from the introduction of web services. Existing integration practice is characterised by competing industry standards and many proprietary protocols, but web services' emphasis on open standards may be an advantage in getting industry players to agree. The use of XML, and its formal definition language XML Schema, can help improve integration protocol implementations by eliminating ambiguities, as well as providing support for automatic validation of messages. Finally, the extensibility of web services and XML can allow integration mechanisms to evolve as markets require new functionality, without causing further fragmentation of protocols.

However, before web services can be used for capital markets systems, various technical requirements must be met. These requirements include performance, security and fault-tolerance.

Existing research into SOAP performance has considered its application in scientific areas such as grid computing, and has focused on the transmission of numerical data. With this emphasis, the predominant cost and weakness of XML-based messages was identified as the encoding and decoding of floating point values [4].

On the other hand, this study examines the performance of SOAP in realistic business computing scenarios. More specifically, the primary goal of this research is to consider the feasibility of using SOAP in capital markets systems, and particularly in real-time trading systems.

The approach adopted by this study is to evaluate the performance of SOAP against existing practice. The study compares the performance of SOAP with the established, widely used, domain-specific protocol, FIX. The relative performance of SOAP and FIX may be useful in determining whether SOAP can meet the performance requirements of capital markets.

The emphasis of this study was on the inherent performance limitations of the wire formats. Both SOAP and FIX use a text-based wire representation, and therefore it may seem reasonable to conclude, based on the existing research, that both would be impacted by the same inherent inefficiencies. For this reason, a binary wire format, CDR, was included in the comparison to gauge the costs associated with text encoding.

The study finds, firstly, that in business applications SOAP does indeed perform poorly compared to the binary wire format, CDR. SOAP messages are some 2-4 times the size of the equivalent CDR messages. Latency over local networks is substantially increased,

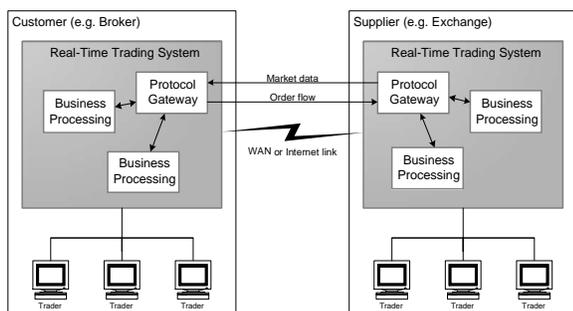


Figure 1: Integration between real-time trading systems

with encoding 8–10 times and decoding some 5 times more expensive. These results are similar to the conclusions of earlier studies, although the results show a less marked difference than when the focus is on transmission of numerical data.

When compared to FIX, SOAP again exhibits poorer performance. SOAP messages are 3.5–4.5 larger than FIX, latency is 2–3 times worse, and encoding/decoding costs are increased by up to nearly 9 times.

Given that FIX, like SOAP, is text-based, the surprising result is that FIX performed comparably to CDR. From this we have been led to conclude that, in realistic business application scenarios, SOAP's poor performance cannot be adequately explained simply by the disadvantages of text-based over binary wire formats. This also suggests that improvements in the efficiency of SOAP encoders and decoders may enable its use in high performance business applications.

2. BACKGROUND

Software systems used in capital markets can be classified according to their position in the trading lifecycle [21]:

- **Pre-trade services.** The delivery of real-time and historical market data, analysis of this data, and routing of an order to the best trading entity for a transaction.
- **Trading.** The execution of a trade itself at a trading entity, such as an exchange like the Australian Stock Exchange (ASX).
- **Post-trade services.** Operations performed prior to finalisation of a trade, such as surveillance, compliance checking, or risk management.
- **Settlement.** Finalisation of a trade by transferring money between the buyer and seller.
- **Registry.** Transfer of ownership of the securities from seller to buyer.

This paper will focus on the pre-trade part of the lifecycle, and in particular the integration needs of real-time trading systems. Integration between real-time trading systems typically involves the communication of live market data as well as the flow of buy and sell orders, as shown in Figure 1. Given the potentially large volumes of data and the need for timely delivery, integration between real-time trading systems has, in the authors' experience, the highest performance requirements in the domain.

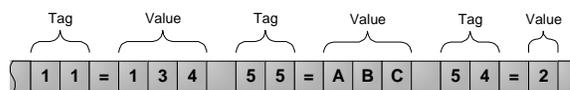


Figure 2: An example FIX protocol message fragment

2.1 FIX Protocol

The Financial Information eXchange (FIX) protocol [10] is a messaging standard developed specifically for the real-time electronic exchange of securities transactions.

FIX messages are text-based, and consist of tag-value pairs separated by a special delimiter character (SOH, which is ASCII value 0x01) as illustrated by Figure 2. The tags are short strings of digits, and types of values include strings, integers, floating point values, timestamps and arbitrary binary data. Although the content of a message is represented by complex application structures, the layout of an encoded message is flat with flexible ordering of fields. The protocol specification describes, in natural language, the set of available tags, their corresponding business meanings, and the required message structure.

Recent versions of the FIX protocol have introduced an XML-based message format, called FIXML [9]. This provides FIX messages with a rich on-the-wire structure, enabling automated validation and reducing the inherent ambiguities of the tag-based approach. XML also allows the FIX standard to evolve to include new functionality without causing further version fragmentation.

As this paper evaluates the suitability of SOAP for capital markets systems, the FIX protocol will be used as the basis for some comparisons. FIX has been selected for this purpose over other industry protocols due to its wide usage. A 1999 survey of market participants, referenced in [12], found that 82% of surveyed brokers used FIX. The influence of FIX also extends to many organisations that use variants of the standard protocol, or use protocol message definitions that may be classed as FIX-like, such as the ASX's SEATS Open Interface [1].

3. RELATED WORK

Several studies have evaluated the performance of SOAP and XML [6, 13, 3]. These studies all agreed that SOAP and XML incur a substantial performance penalty compared to binary protocols.

[6] conducted an experimental evaluation of the latency performance of various SOAP implementations, comparing with other protocols such as Java RMI and CORBA/IIOP. A conclusion drawn from these results was that SOAP is orders of magnitude slower, although for some of the slowest SOAP systems this can be partly explained by poor implementation.

[13] evaluated the performance of SOAP for high performance scientific computing. Their experiments compared Java RMI with SOAP by sending large arrays of doubles (i.e. floating point values with 18 decimal digits of precision). The results showed that SOAP is much slower than Java RMI, typically by about a factor of ten. They concluded that SOAP's XML messages were inherently unsuitable for use in transferring bulk data, but due to the format's flexibility and accessibility, may be useful as part of a multi-protocol system with SOAP as a "lingua franca".

[3] presented the results of experiments that compared the encoding, decoding and network performance of various message formats, including XML. They found that the marshalling and communications costs of XML are staggeringly high in comparison to more traditional approaches, with XML some 2 to 4 orders of mag-

nitude slower in encoding and decoding than CORBA/IIOP and similar binary wire formats. They concluded that XML wire formats are inappropriate for high performance systems, as the baseline performance of all systems is strongly determined by their wire format.

These studies identified some factors that can affect the performance of web services and SOAP, which can be broadly grouped into three main categories.

3.1 Quality of Implementation

Design and implementation decisions made by SOAP infrastructure vendors can have a considerable impact of performance. These factors include:

- **Choice of XML parsing method.** Different XML parsing models have different trade-offs with respect to memory efficiency, computational speed, and ease of use. Of the generic XML parsing models, pull parsing offers the highest performance, as well as high memory efficiency [14]. However, [4] found that schema-specific parsers can greatly enhance performance compared to general purpose XML parsers, particularly where large data structures are involved.
- **Message length calculation.** When using SOAP with an HTTP/1.0 network binding, the length of the body must be specified in the “Content-Length” header field [2]. Setting the HTTP Content-Length field is difficult for dynamic data, because the message must be buffered as it is constructed to determine how long it is, and the message itself is not sent until encoding is finished [6].
- **Connection establishment costs.** HTTP/1.0 [2] mandates the establishment of a new connection for each operation. Establishing a new connection for each transaction can have a negative impact on performance due to interaction with certain TCP features [23], such as the three-way handshake [24] and slow start algorithm [16]. When looking at HTTP/1.0 performance over the Internet, at least one quarter of the transaction time may be taken up by connection establishment [17]. HTTP/1.1 provides a connection keep-alive feature that allows a client to perform multiple operations over a single connection [8].
- **Pipelining** HTTP/1.1 adds support for *pipelining* [8], which is the ability to send multiple requests on the connection before waiting for a response. This allows the connection to be used more efficiently. As an alternative to pipelining, an HTTP/1.0 implementation can make multiple single-operation connections in parallel, typically by using multithreading. [19] found that an HTTP/1.1 implementation using buffered pipelining will use less than 10% of the number of TCP packets that HTTP/1.0 does, and execute in less elapsed time. However HTTP/1.1 without pipelining has a higher elapsed time than HTTP/1.0 using multiple connections. This indicates that connection keep-alive alone is insufficient to improve performance.
- **Inappropriate TCP options.** [6] found that some SOAP implementations suffered from significantly worse performance due to interaction between the Nagle algorithm and the TCP delayed acknowledgement algorithm in the operating systems for the client and server [24]. Clearly this is not an inherent problem with either SOAP or HTTP.

3.2 Network Protocol Binding

The FIX protocol defines a session as a “bi-directional stream of ordered messages between two parties” [10], and so there are no request-response semantics imposed by its specification. Consequently, when seeking to apply SOAP to a real-time trading system we would prefer to use messaging-style rather than RPC-style communication. Since HTTP is a request-response protocol [8] with strict client and server roles, it may be ill-suited to use in message-style communication.

Fortunately, SOAP does not specify a particular network transport binding, and using SOAP with alternative network protocols may offer performance advantages. This clearly implies that the inefficiencies attributed to HTTP are not inherent to SOAP.

3.3 Inherent Limitations of SOAP

Open metadata technologies such as XML can provide a large gain in usability, but the success of these technologies requires that their use does not unreasonably degrade performance [28].

XML is extremely robust with respect to changes in the format of the incoming record [3]. However, the use of XML can negatively impact the performance of SOAP in the following areas:

- **Speed of encoding and decoding.** The conversion of data from binary to ASCII and vice-versa is the major performance cost of XML [3]. This is particularly the case for floating point values, which were found constitute the major cost of XML encoding and decoding [4]. The use of a text-based protocol also precludes the application of optimisations available to binary protocols when communication occurs between homogeneous systems [3].
- **Message size.** For XML, an expansion factor of 6–8 times over the original binary data is not unusual [3]. [4] measured expansion at 4 to 10 times, and [13] found that SOAP’s data representation size is typically about 10 times the size of the equivalent binary representation. This substantially greater size may result in higher network transmission costs and increased latency.

One suggested strategy for overcoming these inherent performance inefficiencies is the use of binary XML representations [28, 11, 25, 18].

4. EXPERIMENTAL DESIGN

The FIX protocol, like XML and SOAP, is text-based [10]. This means that FIX has the same performance issues with regard to the encoding and decoding of numerical data. Similarly, FIX messages may be larger than their equivalent binary representation, although overhead is lower than for XML due to FIX’s compact tag-value format.

The focus of this study is on the inherent performance issues of the SOAP and FIX wire formats. With this in mind, the experiments were designed to eliminate quality of implementation and network protocol factors from consideration. This was done by sending messages encoded in the various wire formats over “raw” TCP sockets, using a consistent network programming model in each case. SOAP bindings such as HTTP were not used. Furthermore, initial transmissions were excluded from the results to eliminate effects from the TCP slow start algorithm [16], and the TCP_NODELAY option was turned on to disable the Nagle algorithm [24].

To aid in the identification of performance issues associated with text-based wire formats, comparisons were also made with a binary

```

msg.header.SenderCompID = "ABC"
msg.header.TargetCompID = "XYZ"
msg.header.SendingTime = "20021116-10:15:28"

msg.body.MarketDataInc.MDReqID = "MYREQ"

msg.body.MarketDataInc.MDEntry[0].MDUpdateAction = Change
msg.body.MarketDataInc.MDEntry[0].MDEntryID = "FOO.last"
msg.body.MarketDataInc.MDEntry[0].MDEntryPx = 13.42
msg.body.MarketDataInc.MDEntry[0].MDEntrySize = 1200
...

```

Figure 3: An example application data structure for the “market data incremental refresh” message

wire format. The Common Data Representation (CDR) [20], which is used as the basis of CORBA communication, was selected for this purpose.

Three types of experiment were conducted:

- **Message Size.** The encoded representation of the application structure in the three wire formats was compared.
- **Latency.** Test programs were written to measure round trip times for sending a single message using each of the wire formats. High resolution timers enabled the separation of the times spent on encoding, decoding and network transmission.
- **Throughput.** These tests attempted to measure the amount of time required to transmit, in a single direction, a large volume of messages. The first set of tests measured throughput from application data structure at one end to application data structure at the other. Comparing with the throughput of pre-encoded messages allows a determination as to whether encoding/decoding costs or the network was the limiting factor.

4.1 Data

The “market data incremental refresh” FIX Message was selected as the business data for the experiment. This message is used for sending market data updates, such as the latest stock prices, throughout a trading day and would typically be high volume and time critical. Randomly generated instances of this message type, similar to that shown in Figure 3, were transmitted. The number of MDEntry items in each message is varied from 1–10 to alter the message size and allow estimation of the fixed and incremental performance costs for each message.

4.2 Software

Application data structures were translated to and from the wire formats using schema-specific encoders and decoders. The software tools used to accomplish this were as follows:

- SOAP messages were produced using gSOAP [27], a free, high performance toolkit. Studies [27, 4] showed that gSOAP had substantially higher performance than some commonly used SOAP implementations, although not as high as another special-purpose research implementation. The XML schema for the message was based on the FIXML DTD [9], with encoders and decoders generated in C++.
- An implementation of a subset of the FIX protocol was developed specifically for the research. The application structure was defined in CORBA IDL, and the encoders and decoders were in C++.

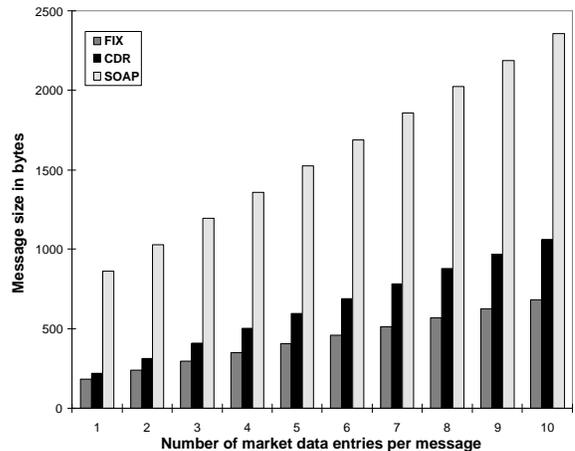


Figure 4: Message size comparison of the wire formats

Table 1: Summary of message size costs

	Fixed Cost	Per-Entry Cost
FIX	130.0 bytes	55.05 bytes
CDR	129.2 bytes	93.28 bytes
SOAP	695.8 bytes	166.0 bytes

- The same CORBA IDL definition as used with the FIX implementation was compiled using the TAO CORBA ORB’s [22] IDL compiler to generate C++ encoders and decoders for the CDR wire format.

4.3 Environment

The latency and throughput tests were run over both 10 Mbps and 100 Mbps Ethernet. Communication between real-time trading systems often occurs across leased lines or the public Internet, and so 10 Mbps may be more comparable to the available bandwidth in these cases.

The client system was a uniprocessor 900 MHz Pentium 3 with 256 MB of RAM, 256 KB level 2 cache and running Windows 2000. The test software for this system was compiled using the Borland C++ 5.6.1 compiler.

The server was a uniprocessor 500 MHz Pentium 3 with 256 MB of RAM, 512 KB level 2 cache, running Redhat Linux 7.3 with the 2.4.18-3 kernel. The test software for this system was compiled using the g++ 3.2.1 compiler.

5. EXPERIMENTAL RESULTS

5.1 Message Size

Prior to running experiments to measure latency and throughput, data were collected to compare the size of the application messages when encoded in each of the SOAP, FIX and CDR wire formats.

Figure 4 and Table 1 show the message sizes of the application data structure when encoded into each format. The figure shows that SOAP messages are substantially larger, being some 3.5–4.5 times larger than the equivalent FIX message, and 2–4 times larger than one in encoded using CDR. This is on the low side of the relative size results presented by existing SOAP and XML performance studies [3, 4, 13].

Figure 5 shows the application structure from Figure 3 encoded using FIX, and Figure 6 shows the same data encoded into SOAP.

```

<?xml version="1.0" encoding="UTF-8"?>
<SOAP-ENV:Envelope
xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"
xmlns:SOAP-ENC="http://schemas.xmlsoap.org/soap/encoding/"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <SOAP-ENV:Body
SOAP-ENV:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
    <sendMessage>
      <FIXMLMessage>
        <Header>
          <Sender><ComplID>ABC</ComplID></Sender>
          <Target><ComplID>XYZ</ComplID></Target>
          <SendingTime>2002-11-16T10:15:28.000</SendingTime>
        </Header>
        <ApplicationMessage>
          <MarketDataInc>

```

```

<MDReqID>MYREQ</MDReqID>
<MDInclList>
  <MDInclGroup>
    <MDUpdateAction>1</MDUpdateAction>
    <MDEntryID>FOO.last</MDEntryID>
    <MDEntryPx>13.42</MDEntryPx>
    <MDEntrySize>1200</MDEntrySize>
  </MDInclGroup>
</MDInclList>
</MarketDataInc>
</ApplicationMessage>
</FIXMLMessage>
</sendMessage>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>

```

Figure 6: SOAP representation

```

8=FIX.4.3●9=00000098●35=X●49=ABC●56=XYZ●34=1●52=20
021116-10:15:28●262=MYREQ●268=1●279=1●278=FOO.last●
270=13.42●271=1200●10=185●

```

Figure 5: FIX representation

```

union Optional_String switch(boolean) {
case TRUE:
  string value;
};

```

Figure 7: Optional field idiom used for application data structure definition in CORBA IDL

Here we see that the XML namespaces, the more verbose tag names and syntax contribute to the SOAP message being substantially larger.

Figure 4 also shows that FIX has a more compact wire representation than CDR, which runs counter to what is expected. CDR is approximately 50% larger due to CORBA IDL's lack of built-in support for optional fields. We have instead used a common idiom for defining optional fields in CORBA [15], as illustrated by Figure 7. This means that each optional field uses a single-byte indicator to show whether it is present or not. As the message header contains approximately 20 optional fields, and as each market data entry contains more than 40, then with most of these fields unset there is considerable overhead. Alternative binary wire formats with true support for optional fields, such as ASN.1/BER, may offer more compact messages.

5.2 Latency

Figure 8 and Table 2 present the measurements for round-trip times over a 10 Mbps network. This shows that FIX has the lowest time with CDR not much greater, especially when compared to SOAP which has a round trip time of slightly more than twice the other two.

The breakdown of costs in Figure 9 shows that for all three wire formats, over a 10 Mbps network, the largest cost is the time spent on the network. This would suggest that in this environment the size of the message on the wire is the major limiting factor. Over the slower network, FIX's more compact message representation contributes to its lower round-trip times than CDR.

Over a 100 Mbps network, time spent on the network is less significant in overall round-trip times. Figure 10 and Table 3 show the encoding costs for the wire formats, and Figure 11 and Table 4 show the relative decoding costs. For 100 Mbps Ethernet, the substantially higher encoding and decoding costs for SOAP contribute

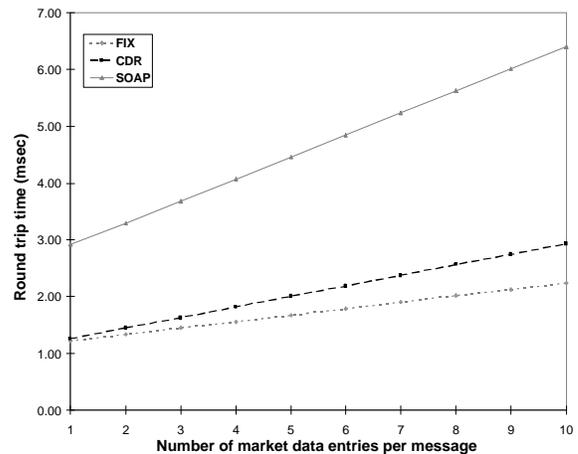


Figure 8: Round-trip times over 10 Mbps network

Table 2: Summary of round-trip costs over 10 Mbps network

	Fixed Cost	Per-Entry Cost
FIX	1.103 msec	0.1130 msec
CDR	1.076 msec	0.1851 msec
SOAP	2.525 msec	0.3875 msec

most to its poorer performance, with round-trips some 2–3 times more expensive than FIX or CDR.

An interesting result shown in Figure 11 is that FIX, a text-based wire format, has lower decoding costs than CDR, a binary format. This is particularly significant given the greater complexity involved in decoding FIX, with the presence of the tags in the wire format, flexible field ordering, and the fact that many fields may or may not be present at all on the wire. With CDR, on the other hand, all fields would be decoded in a fixed order as determined by their definition in the CORBA IDL. This result suggests two things:

- With this realistic application message content there is a mix of string, integer and floating point values. The cost of converting numerical data from text to binary, identified as major by other studies, does not have a predominant role.
- The cost of handling the complexity of the FIX message structure is minor compared to the cost of decoding a field, and that greater benefit is derived from not having to process the large number of optional fields that are not present on the

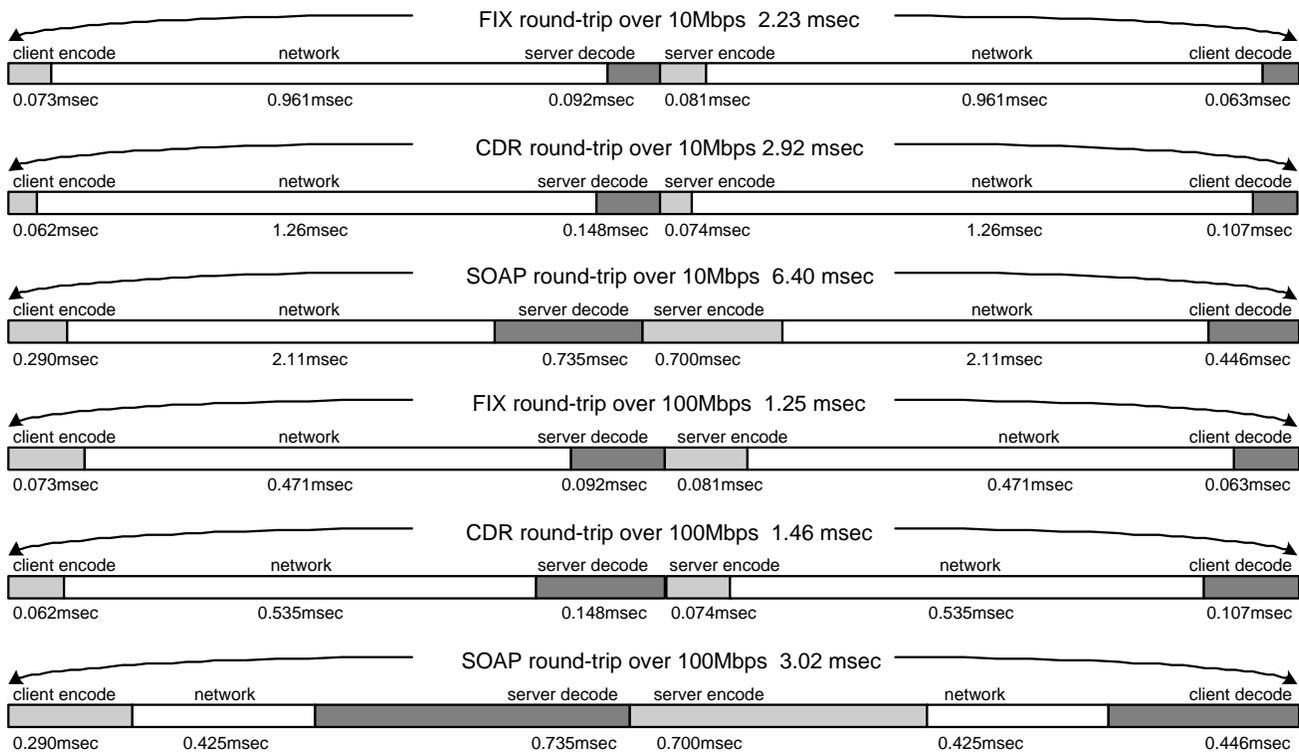


Figure 9: Cost breakdown for round-trip times for messages with 10 market data entries

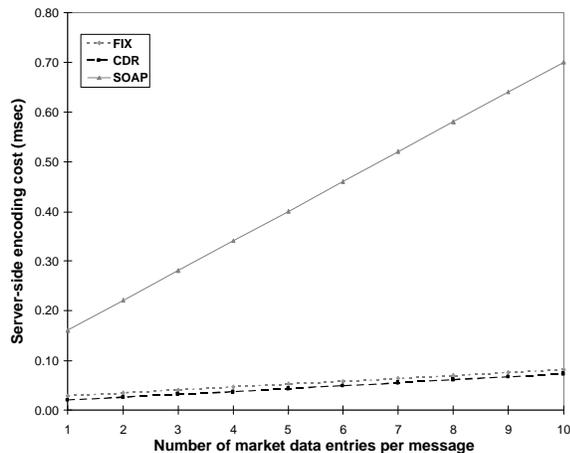


Figure 10: Server-side encoding costs

wire. With the message encoding, encoders for all wire formats must test the presence of every field, and the advantage is lost.

5.3 Throughput

Figure 12 displays the measurements for throughput over a 10 Mbps network. For this slower network configuration, the network itself was observed to be the bottleneck for all three wire formats. As with latency, this result suggests that in an environment with lower bandwidth, the size of the message is the major factor affecting performance. This allows FIX, with the most compact mes-

Table 3: Summary of server-side encoding costs

	Fixed Cost	Per-Entry Cost
FIX	0.0232 msec	0.0058 msec
CDR	0.0141 msec	0.0059 msec
SOAP	0.1012 msec	0.0599 msec

Table 4: Summary of server-side decoding costs

	Fixed Cost	Per-Entry Cost
FIX	0.0358 msec	0.0057 msec
CDR	0.0358 msec	0.0112 msec
SOAP	0.1878 msec	0.0547 msec

sages, to achieve the highest throughput values.

For 100 Mbps networks, the CPU on the slower server machine was observed to be the bottleneck, and consequently the network was under utilised. The results in Figure 13 show that FIX again achieved the highest throughput, although CDR has lower encoding costs. This is a result of the decoding, which had a lower cost for FIX than CDR, being performed on the slower machine. Reversing the roles of the machines changes the relative throughput of the wire formats.

Interestingly, the throughput performance of SOAP relative to the other two wire formats is worse for the 100 Mbps network. This is due to the ratio of SOAP decoding cost to FIX or CDR being greater than the equivalent ratio for message size.

5.4 SOAP Message Compression

Over lower network bandwidth the size of the message on the wire is the limiting factor for performance. As a result, it may be

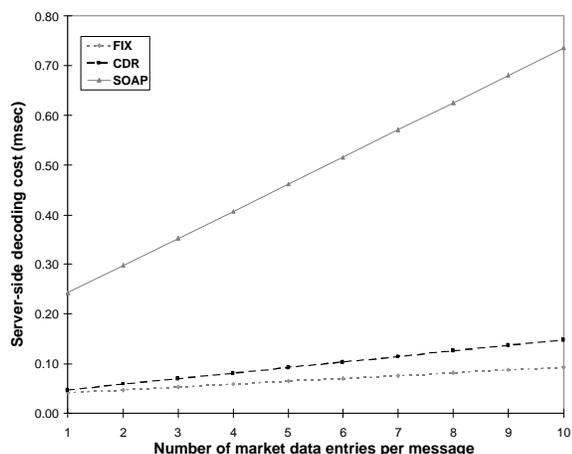


Figure 11: Server-side decoding costs

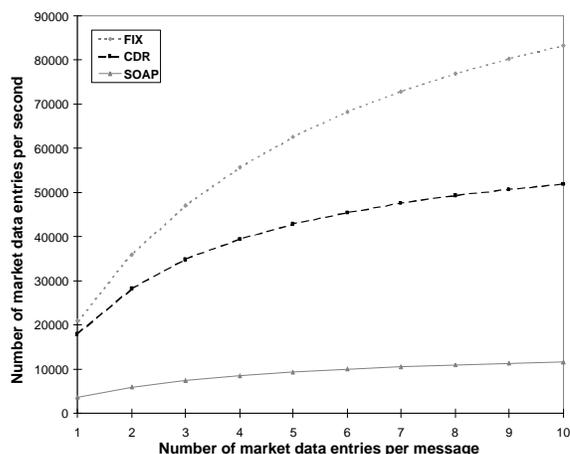


Figure 13: Throughput over 100 Mbps network

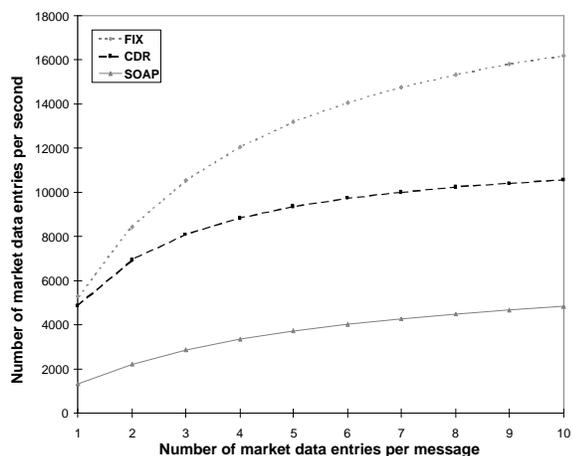


Figure 12: Throughput over 10 Mbps network

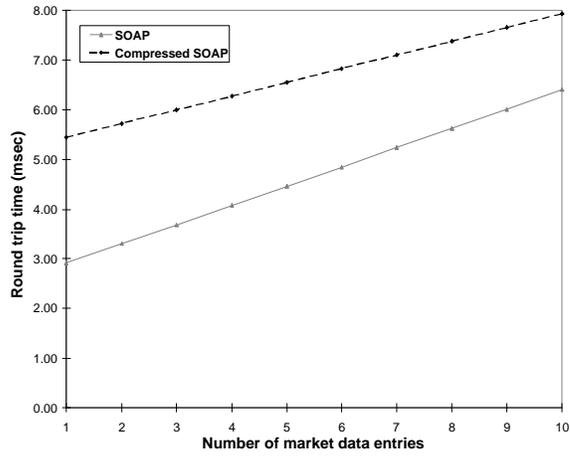


Figure 14: Round-trip times over 10 Mbps network for compressed and uncompressed SOAP messages

possible that compression of the SOAP message data would confer some advantage. To determine if this is the case an additional latency test was run where the SOAP messages were compressed immediately before being transmitted. For this purpose, the `zlib` compression library was used on the lowest (and fastest) compression level. This achieved compression savings of 50–70%.

The results, as shown in Figure 14, indicate that compression is in fact detrimental, substantially increasing the round-trip time. The increased CPU time spent compressing and decompressing the messages outweighs any benefits. Compression may only be useful for considerably slower networks.

5.5 Compact XML Tags

An alternative method for reducing the size of the SOAP messages investigated in this study was to reduce the length of the XML tag names. This was done by replacing the FIXML names with short 2–4 character strings based on the numeric FIX tags. This reduced the size of the SOAP messages by approximately 25–35% as shown in Figure 15, but clearly sacrifices message readability in favour of the potential performance gains.

Figure 16 shows that the more compact SOAP messages do provide gains in performance over 10 Mbps Ethernet, where the time

on the network is the major cost. However, the performance improvement is not in the same proportion to the reduction in message size. When considering the relative decoding costs shown in Figure 17, we see that there is not a commensurate improvement in decoding performance. Furthermore, the use of compact SOAP has a negligible effect on encoding efficiency. This suggests that the major cost of the XML encoding and decoding is in the structural complexity and syntactic elements, rather than the data contained in the message or the tag names.

6. DISCUSSION

Earlier studies into SOAP and XML performance [3, 4] found that the conversion from text to binary and vice versa was the major cost, and particularly the costs associated with encoding and decoding floating point values. However, these studies were oriented towards the application of SOAP and XML to scientific computing, with message data consisting, for the large part, of numerical values.

In this study we have attempted to study the performance of SOAP using realistic business application messages, with capital markets trading systems as the context. The results comparing

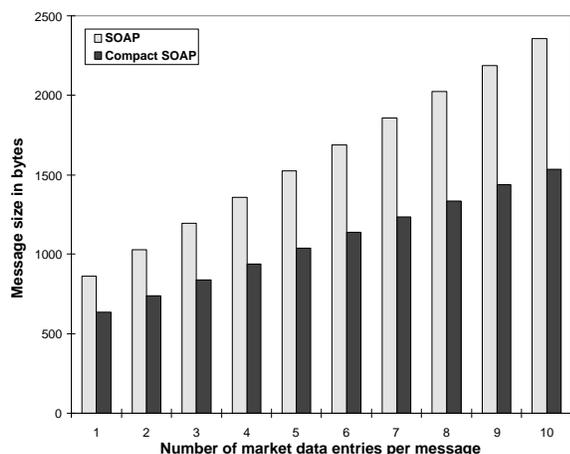


Figure 15: Message size reduction when compact XML tags are used

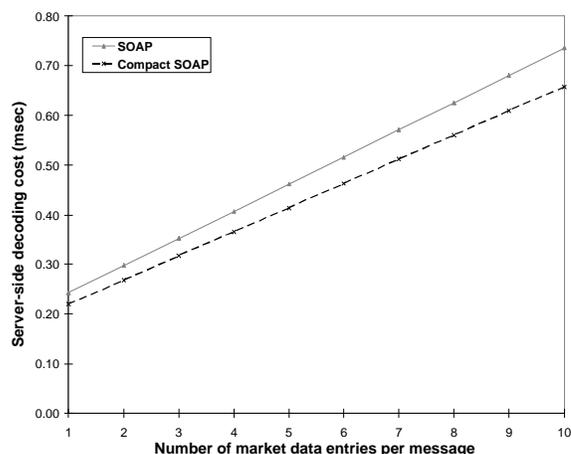


Figure 17: Decoding costs for compact SOAP messages

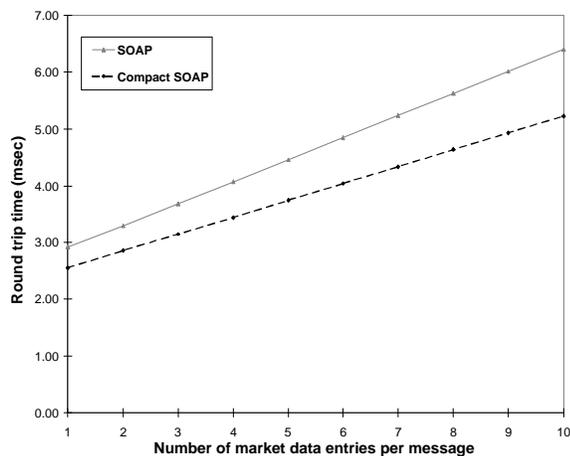


Figure 16: Round-trip times for compact SOAP messages over 10 Mbps network

SOAP to the binary wire format, CDR, do display poor performance for SOAP, although the difference is not as large as for the numerical data used in earlier studies. Given that the overall performance of FIX, with its text-based wire format, was comparable to CDR — and in fact outperforming it for decoding — it is clear that conversion of text-to-binary and back is not a major factor affecting performance in this case.

Two important results of this study with respect to the performance of SOAP are:

- In business computing scenarios, it is possible for text-based formats to have comparable performance to binary wire formats. The poorer performance of SOAP implementations compared to other wire formats cannot be fully explained simply by the fact that SOAP is text-based.
- Simply reducing the size of the encoded SOAP messages by using shorter tags did not proportionally improve the speed of encoding and decoding.

Together, these results mean that a likely cause of the poor performance of SOAP as a wire format is the complexity of the XML

syntax and the richness of its on-the-wire structure. The SOAP message definition used in this study, based closely on FIXML [9], is complex with a high degree of nesting. It may be useful to conduct further research to gauge the effect on performance of alternative XML message representations. Results from such research could provide some guidance to developers on how to effectively design SOAP message layouts for high performance applications.

The results of this study suggest some areas where SOAP implementors, in focusing any efforts to improve performance for business applications, may find the most benefit. Further study would be valuable in clarifying the causes of SOAP's poor performance, and what approaches may be used to address them.

Furthermore, in this study we have considered only the inherent performance characteristics of the SOAP wire format. The other requirements for using SOAP in capital markets systems, such as security and fault tolerance, may have an additional impact on SOAP performance.

Finally, the results show that it is important to consider the environment in which a system will be deployed when identifying the performance issues related to SOAP most relevant to that application. Although for fast networks the speed of encoding and decoding is the predominant determining factor, for slower networks it is the size of the encoded message that determines both latency and throughput performance. This is important for business-to-business integration which, in capital markets as in most other domains, often occurs over wide area networks.

7. CONCLUSIONS

In this paper we have presented the results of a performance evaluation of SOAP in a business application context. Our results indicate that, while SOAP did fare poorly when compared to both binary CDR and the established industry protocol FIX, the difference is less than that measured for scientific computing applications. Furthermore, in realistic business environments it is possible for text-based wire formats to have comparable performance to binary. This indicates that the text-based nature of XML is not in itself the major contributing factor to inefficiency in SOAP encoding and decoding. This finding suggests that further work in improving the performance of SOAP encoders and decoders may make it viable for use in high performance business applications. In spite of this, when designing performance-conscious systems for integration across wide area networks, bandwidth is generally the limiting

factor, and it is worth considering the size of an encoded message when selecting an appropriate wire format.

8. REFERENCES

- [1] Australian Stock Exchange. The SEATS computer system, 2000. http://www.asx.com.au/markets/14/SEATS_AM4.shtm, accessed 1 June 2002.
- [2] T. Berners-Lee, R. Fielding, and H. Frystyk. Hypertext transfer protocol - HTTP/1.0, 1996. IETF RFC 1945, <http://www.ietf.org/rfc/rfc1945.txt>.
- [3] F. E. Bustamante, G. Eisenhauer, K. Schwan, and P. Widener. Efficient wire formats for high performance computing. In *Proceedings of the 2000 Conference on Supercomputing*, 2000.
- [4] K. Chiu, M. Govindaraju, and R. Bramley. Investigating the limits of SOAP performance for scientific computing. In *Proceedings of the 11th IEEE International Symposium on High Performance Distributed Computing*, pages 246–254, 2002.
- [5] F. Curbera, M. Duftler, R. Khalaf, W. Nagy, N. Mukhi, and S. Weerawarana. Unraveling the web services web: An introduction to SOAP, WSDL, UDDI. *IEEE Internet Computing*, 6(2):86–93, March–April 2002.
- [6] D. Davis and M. Parashar. Latency performance of SOAP implementations. In *Proceedings of the 2nd IEEE/ACM International Symposium on Cluster Computing and the Grid*, pages 407–412, 2002.
- [7] M. Fan, J. Stallaert, and A. B. Whinston. The internet and the future of financial markets. *Communications of the ACM*, 43(11):83–88, November 2000.
- [8] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee. Hypertext transfer protocol - HTTP/1.1, 1999. IETF RFC 2616, <http://www.ietf.org/rfc/rfc2616.txt>.
- [9] FIX Protocol Ltd. FIXML: A markup language for the FIX application message layer. <http://www.fixprotocol.org/WORKGROUPS/928951581/wpaper.html>, accessed 8 June 2002.
- [10] FIX Protocol Ltd. The Financial Information Exchange Protocol (FIX), version 4.3, August 2001. <http://www.fixprotocol.org/specification/fix-43-pdf.zip>, accessed 8 June 2002.
- [11] M. Girardot and N. Sundaresan. Millau: An encoding format for efficient representation and exchange of XML over the web. In *Proceedings of the 9th International World Wide Web Conference*, pages 747–765, 2000.
- [12] J. Goeller. FIXML and STP related efforts, 2000. http://www.fixprotocol.org/WORKGROUPS/928951581/XML_STP_John6.ppt, powerpoint presentation, accessed 8 June 2002.
- [13] M. Govindaraju, A. Slominski, V. Choppella, R. Bramley, and D. Gannon. Requirements for and evaluation of RMI protocols for scientific computing. In *Proceedings of the 2000 Conference on Supercomputing*, 2000.
- [14] S. Graham, S. Simeonov, T. Boubez, D. Davis, G. Daniels, Y. Nakamura, and R. Neyama. *Building Web Services with Java: Making Sense of XML, SOAP, WSDL, and UDDI*. Sams Publishing, Indianapolis, 2002.
- [15] M. Henning and S. Vinoski. *Advanced CORBA Programming with C++*. Addison-Wesley, Reading, Massachusetts, 1999.
- [16] V. Jacobson. Congestion avoidance and control. In *Symposium proceedings on Communications architectures and protocols*, pages 314–329. ACM Press, 1988.
- [17] B. Liu and E. A. Fox. Web traffic latency: Characteristics and implications. *J.UCS: Journal of Universal Computer Science*, 4(9):763–778, 1998.
- [18] B. Martin and B. Jano. WAP binary XML content format, June 1999. <http://www.w3.org/TR/wbxm1/>, accessed 1 June 2002.
- [19] H. F. Nielsen, J. Gettys, A. Baird-Smith, E. Prud’hommeaux, H. W. Lie, and C. Lilley. Network performance effects of HTTP/1.1, CSS1, and PNG. In *Proceedings of the ACM SIGCOMM ’97 conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, pages 155–166, 1997.
- [20] Object Management Group. The Common Object Request Broker Architecture: Core Specification, version 3.0, November 2002.
- [21] F. A. Rabhi and B. Benatallah. An integrated service architecture for managing capital market systems. *IEEE Network*, 16(1):15–19, 2002.
- [22] D. C. Schmidt, D. L. Levine, and S. Mungee. The design of the TAO real-time object request broker. *Computer Communications*, 21(4):294–324, April 1998.
- [23] S. E. Spero. Analysis of HTTP performance problems, 1994. <http://www.w3.org/Protocols/HTTP/1.0/HTTPPerformance.html>, accessed 15 June 2002.
- [24] W. R. Stevens. *TCP/IP Illustrated, Volume 1: The Protocols*. Addison-Wesley, Reading, Massachusetts, 1994.
- [25] N. Sundaresan and R. Moussa. Algorithms and programming models for efficient representation of XML for internet applications. In *Proceedings of the 10th International World Wide Web Conference*, pages 366–375, 2001.
- [26] SWIFT. About SWIFT - History. <http://www.swift.com>, accessed 3 June 2002.
- [27] R. A. van Engelen and K. A. Gallivan. The gSOAP toolkit for web services and peer-to-peer computing networks. In *Proceedings of the 2nd IEEE/ACM International Symposium on Cluster Computing and the Grid*, pages 128–135, 2002.
- [28] P. Widener, G. Eisenhauer, and K. Schwan. Open metadata formats: Efficient XML-based communication for high performance computing. In *Proceedings of the 10th IEEE International Symposium on High Performance Distributed Computing*, pages 371–380, 2001.