COMPOSITION, PERFORMANCE ANALYSIS AND SIMULATION OF

WEB SERVICES

by

SENTHILANAND CHANDRASEKARAN

(Under the direction of Dr. John. A. Miller)

ABSTRACT

The new paradigm for distributed computing over the Internet is that of Web services. They are Web accessible software components which can be combined and linked together to create new functionality in the form of Web processes. This also creates the need to compose services into processes that are efficient in terms of their performance. In this work, we describe our Service Composition and Execution Tool (SCET) and also describe the various methodologies that could be adopted for evaluating the performance of a Web process. SCET allows for composing services statically using its designer and storing them as Web Service Flow Language (WSFL) based specifications. For executing a process, SCET automatically generates the Perl execution code for a composed process. As processes involving real world services are difficult to be analyzed for their performance, we have used simulation as an alternate technique for analyzing the efficiency of a process.

INDEX WORDS: Web services, Web processes, Web service composition,
             Simulation, Performance evaluation

COMPOSITION, PERFORMANCE ANALYSIS AND SIMULATION OF

WEB SERVICES

by

SENTHILANAND CHANDRASEKARAN

B.E., Anna University, India, 2000

A Thesis Submitted to the Graduate Faculty of The University of Georgia in Partial

Fulfillment of the Requirements for the Degree

MASTER OF SCIENCE

ATHENS, GEORGIA

2002

COMPOSITION, PERFORMANCE ANALYSIS AND SIMULATION OF

WEB SERVICES

by

SENTHILANAND CHANDRASEKARAN

Approved.

Major Professor.   John A. Miller

Committee.   Budak Arpinar
                Walter D. Potter

Electronic Version Approved.

Maureen Grasso
Dean of the Graduate School
The University of Georgia
December 2002

DEDICATION

To my father Chandrasekaran, mother Jeeva and brother Premanand

ACKNOWLEDGEMENTS

TABLE OF CONTENTS

CHAPTER 1

INTRODUCTION

A new wave of development based upon the eXtensible Markup Language (XML) has started. One of the interesting aspects of this development influencing the Web is the "Web services" technology, which is a new distributed computing paradigm based on XML. Web services are universally accessible software components deployed on the Web. These software components which are available as services, create an interesting scenario in which efficient Web processes can be created with existing services.

Web services are "self contained, self-describing modular applications that can be published, located, and invoked across the Web" (Tidwell 2001). The Web services architecture (Austin 2002) models software as individual components available on the Web. Such software components are described by an interface listing the collection of operations that are network accessible through standard XML messaging (Colan 2001). Web services are suitable for integrating e-business applications as they allow for creating loosely coupled distributed systems based on XML messaging protocols (e.g., the Simple Object Access Protocol (SOAP) (Kulchenko 2002)).

As individual services are limited in their capability, we need to compose existing Web services to create new functionality in the form of Web processes. Web service composition is, as explained in (Piccinelli 1999), the ability to take existing services (or building blocks) and combine them to form new services. In carrying out this composition task, one has to be concerned about the efficiency and the Quality of Service (QoS) that the composed process will

exhibit upon its execution. This task of composing services to create efficient Web processes raises the following issues:

- *Composition of a Web Process:* There are two types of process composition: static and dynamic. In a static composition, the services to be composed are chosen at design time, while in a dynamic composition, they are chosen at run-time. Both types of composition involve searching for Web services. Process composition involves three phases (Yang et al. 2001): planning, definition and implementation. In the planning phase, issues such as the type of composition to be employed and the Web service search techniques to be applied are decided based on the process requirements. The composer also has to deal with the issue of representing the process structure in a form and language that are appropriate to the problem being solved.

- *Execution of a Composed Web Process:* Executing a process enables one to realize its functionality and also analyze its performance. The central authority pattern and peer-to-peer enactment patterns (Benatallah et al. 2002) are two major execution techniques applied for service compositions. In the central authority pattern, the Web process execution engine acts as a scheduler, interpreting the process specification and executing the services involved accordingly. In the peer-to-peer enactment patterns, the responsibility for coordinating the execution of a Web process is distributed across several service providers.

- *Efficiency of a Composed Web Process:* A composed process should be efficient in terms of its service time, its ability to handle higher loads. For static compositions, as the services in the process are bound at the design time, the designer can search for services that have operational metrics (such as service time, load capacity, cost, and reliability) satisfying the requirements of the problem being solved. Mathematical methods have

been used by Cardoso, Miller et al. (2002) to analyze and estimate the overall QoS of a process. Another alternative for estimating the QoS of a process is to utilize simulation analysis (Miller et al. 2002). Simulation can play an important role in evaluating the quality of a Web process, before its actual execution. For dynamic compositions, the actual efficiency of the process cannot be determined until it is invoked.

In this paper, we address the above issues related to composing efficient processes and executing them, using our Service Composition and Execution Tool (SCET). We also describe the methodologies that could be followed in analyzing the performance of individual Web services involved in a process. Processes involving world-altering services and external services are difficult to be analyzed for their performance, as it might be an expensive and impractical task to test them under various conditions. To address this problem, simulation has been used as an alternate means to estimate the efficiency of a process.

As part of this work, being done at the University of Georgia, we have developed a Service Composition and Execution Tool (SCET). As the problem of composing services involves representing the process, executing the process and analyzing the executed process, we have worked on the above three related issues using our service composition tool. SCET allows static composition of services, by modeling the process as a digraph in a graphical designer. The process descriptions are stored as a Web Service Flow Language (WSFL) based specification (Leymann 2001). As WSFL supports dynamic selection of services, our system can be modified to support dynamic discovery of services at execution time from a Web Services registry (such as the Universal Description, Discovery and Integration (UDDI) registry).

SCET can automatically generate Perl execution code from the WSFL based process specification, enabling easier execution of the composed process. For performance analysis, we

list different techniques that can be used to evaluate the individual Web services involved in a process.

Analyzing the performance of a Web process by executing it, is not always feasible, as it requires control over the Web services and the hosts involved. In cases where we do not have control over the Web services (like live Internet Web services) to do performance analysis tests, SCET users can use simulation as a tool in evaluating the efficiency of the composed process. The JSIM simulator (Nair et al. 1996; Miller et al. 1997) integrated with SCET can simulate the execution of a Web process under various hypothetical conditions and generate statistical results. These results approximate the actual invocation, allowing decisions to be made on the behavior of the process without actual execution. This also allows us to find bottlenecks and performance problems in the service components, suggesting reordering or replacement of those components.

CHAPTER 2

COMPOSITION, PERFORMANCE ANALYSIS AND SIMULATION OF

WEB SERVICES[1]

---

1. Introduction

A new wave of development based upon the eXtensible Markup Language (XML) has started. One of the interesting aspects of this development influencing the Web is the "Web services" technology, which is a new distributed computing paradigm based on XML. Web services are universally accessible software components deployed on the Web. These software components which are available as services, create an interesting scenario in which efficient Web processes can be created with existing services.

Web services are "self contained, self-describing modular applications that can be published, located, and invoked across the Web" (Tidwell 2001). The Web services architecture (Austin 2002) models software as individual components available on the Web. Such software components are described by an interface listing the collection of operations that are network accessible through standard XML messaging (Colan 2001). Web services are suitable for integrating e-business applications as they allow for creating loosely coupled distributed systems based on XML messaging protocols (e.g., the Simple Object Access Protocol (SOAP) (Kulchenko 2002)).

As individual services are limited in their capability, we need to compose existing Web services to create new functionality in the form of Web processes. Web service composition is, as explained in (Piccinelli 1999), the ability to take existing services (or building blocks) and combine them to form new services. In carrying out this composition task, one has to be concerned about the efficiency and the Quality of Service (QoS) that the composed process will exhibit upon its execution. This task of composing services to create efficient Web processes raises the following issues:

- *Composition of a Web Process:* There are two types of process composition: static and dynamic. In a static composition, the services to be composed are chosen at design time,

while in a dynamic composition, they are chosen at run-time. Both types of composition involve searching for Web services. Process composition involves three phases (Yang et al. 2001): planning, definition and implementation. In the planning phase, issues such as the type of composition to be employed and the Web service search techniques to be applied are decided based on the process requirements. The composer also has to deal with the issue of representing the process structure in a form and language that are appropriate to the problem being solved.

- *Execution of a Composed Web Process:* Executing a process enables one to realize its functionality and also analyze its performance. The central authority pattern and peer-to-peer enactment patterns (Benatallah et al. 2002) are two major execution techniques applied for service compositions. In the central authority pattern, the Web process execution engine acts as a scheduler, interpreting the process specification and executing the services involved accordingly. In the peer-to-peer enactment patterns, the responsibility for coordinating the execution of a Web process is distributed across several service providers.

- *Efficiency of a Composed Web Process:* A composed process should be efficient in terms of its service time, its ability to handle higher loads. For static compositions, as the services in the process are bound at the design time, the designer can search for services that have operational metrics (such as service time, load capacity, cost, and reliability) satisfying the requirements of the problem being solved. The operational metrics of services can be described using a suitable Quality of Service (QoS) model (Cardoso, Miller et al. 2002; Cardoso, Sheth et al. 2002; Miller et al. 2002). Mathematical methods have been used by Cardoso, Miller et al. (2002) to analyze and estimate the overall QoS of a process. Another alternative for estimating the QoS of a process is to utilize

simulation analysis (Miller et al. 2002). Simulation can play an important role in

evaluating the quality of a Web process, before its actual execution. For dynamic

compositions, the actual efficiency of a dynamically composed process cannot be

determined until it is invoked. In both static and dynamic compositions, the performance

data from the executed process could be analyzed to provide feedback on the efficiency

of the composed process.

As discussed above, the problem of composing a Web process involves the tasks shown in

Figure 2.1.
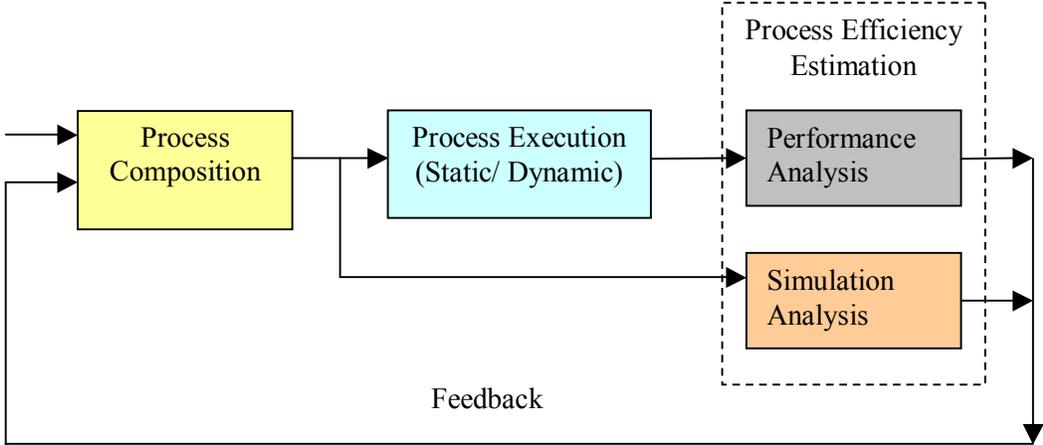
Figure 2.1. Process Composition Tasks

In this paper, we address the above issues related to composing efficient processes and

executing them, using our Service Composition and Execution Tool (SCET). We also describe

the methodologies that could be followed in analyzing the performance of individual Web

services involved in a process. Processes involving world-altering services and external services

are difficult to be analyzed for their performance, as it might be an expensive and impractical task

to test them under various conditions. To address this problem, simulation has been used as an

alternate means to estimate the efficiency of a process. As a further research item, we consider

providing useful feedback from simulation analysis for improving the composition. Apart from the above mentioned concerns, process composition also involves other issues related to payment mechanisms, trust, reliability, security, geographical location, transaction management, coordination, exception handling and service guarantees between different service providers.

As part of this work, being done at the University of Georgia, we have developed a Service Composition and Execution Tool (SCET). As the problem of composing services involves representing the process, executing the process and analyzing the executed process, we have worked on the above three related issues using our service composition tool. SCET allows static composition of services, by modeling the process as a digraph in a graphical designer. The process descriptions are stored as a Web Service Flow Language (WSFL) based specification (Leymann 2001). As WSFL supports dynamic selection of services, our system can be modified to support dynamic discovery of services at execution time from a Web Services registry (such as the Universal Description, Discovery and Integration (UDDI) registry).

SCET can automatically generate Perl execution code from the WSFL based process specification, enabling easier execution of the composed process. For performance analysis, we list different techniques that can be used to evaluate the individual Web services involved in a process. These techniques allow the composer to evaluate the services involved in the process by providing information about

- the Web services involved in the process that are overloaded;
- the Web services which have high (SOAP) message communication overhead;
- the load capacity of individual services involved in the process.

Analyzing the performance of a Web process by executing it, is not always feasible, as it requires control over the Web services and the hosts involved. In cases where we do not have control over the Web services (like live Internet Web services) to do performance analysis tests,

SCET users can use simulation as a tool in evaluating the efficiency of the composed process. The JSIM simulator (Nair et al. 1996; Miller et al. 1997) integrated with SCET can simulate the execution of a Web process under various hypothetical conditions and generate statistical results. These results approximate the actual invocation, allowing decisions to be made on the behavior of the process without actual execution. This also allows us to find bottlenecks and performance problems in the service components, suggesting reordering or replacement of those components.

The rest of the paper is organized as follows. In Section 2, we present the existing and emerging Web service composition languages. Section 3 discusses the related work in this area, while Section 4 analyzes the issues related to Web services and their composition. Section 5 covers our system architecture, Service Composition and Execution Tool (SCET), and our process execution technique. Section 6 explains our approach for analyzing the performance efficiency of a composed Web process and evaluating/comparing the invoked Web services. Simulation and its application to Web process composition task are discussed in Section 7. Conclusion and future work are presented in Section 8.

2. Service Composition Languages

Web service composition is currently an active area of research, with many languages being proposed by academic and industrial research groups.

IBM's Web Service Flow Language (WSFL) (Leymann 2001) and Microsoft's XLANG (Thatte 2001) were two of the earliest languages to define standards for Web services composition. Both extended W3C's Web Service Description Language (WSDL) (Christensen 2002), which is the standard language used for describing the syntactic aspects of a Web service. WSFL was proposed as an XML based language developed to describe complex service compositions. WSFL supported both static configuration of services as well as run time searching

of services in Web service registries. Microsoft's service composition language, XLANG extended WSDL with behavioral specifications to provide a model for orchestration of services. XLANG required services to be configured statically. Business Process Execution Language for Web Services (BPEL4WS) (Curbera et al. 2002) is a recently proposed specification that represents the merging of WSFL and XLANG. BPEL4WS combines the graph oriented process representation of WSFL and the structural construct based processes of XLANG into a unified standard for Web services composition. ebXML (Waldt 2002) is also an effort in this area, that enables enterprises to conduct business over the Internet using an open XML-based infrastructure. XL (Florescu et al. 2002) is another portable W3C compliant XML programming language designed for implementing Web services, with support for composition.

In contrast to these commercial XML based standards, researchers are developing a unique Web service Markup language called DAML-S (Ankolekar et al. 2001). "DAML-S supplies Web service providers with a core set of markup language constructs for describing the properties and capabilities of their Web services in unambiguous, computer-interpretable form. DAML-S markup of Web services will facilitate the automation of Web service tasks including automated Web service discovery, execution, interoperation, composition and execution monitoring" (Ankolekar et al. 2001). Table 2.1 highlights the differences in some of the main service composition standards.

Table 2.1. Comparison of Composition Standards

	WSFL	XLANG	BPEL4WS	DAML-S
Process Modeling	Supports graph based process modeling	Supports construct based process modeling	Supports a blend of graph/construct based process modeling	Supports construct based process modeling
Support for Semantics	Does not support semantics	Does not support semantics	Does not support semantics	Supports semantic description and discovery of services using ontologies
Support for QoS Specification	Uses WSEL extensibility elements to specify QoS	Does not support QoS specification	Does not support QoS specification yet	Supports QoS specification through its Service Profile class
Relationship to WSDL	Layered on top of WSDL providing composition	Layered on top of WSDL providing composition	Layered on top of WSDL providing composition	Models a language which provides more features than that supported by WSDL

3. Related Work

The problem of dynamic service composition has recently gained more importance with the emergence of Web services. As the number of service providers and services is always on the rise, dynamic service composition systems such as eFlow (Casati et al. 2000), can take advantage of the availability of a wide variety of services. eFlow allows nodes (activities) to have service selection rules. When the eFlow engine tries to execute an activity it calls a "service broker" which executes the service selection rules and returns a list of services (with ranking information). Service selection rules are defined using a service broker-specific language.

The default service broker of the eFlow system, processes service selection rules expressed in XML Query Language (XQL) (Jonathan et al. 1998).

Dynamic discovery of Web services and composition of Web processes closely benefit from each other. UDDI, the widely used XML-based registry for advertising businesses and services, utilizes only keyword and classification based searching of services and businesses. Paolucci et al. (2002) stress the importance of discovering Web services based on semantic matching between the declarative description of the service being sought and the description of the service being offered. In this direction, Cardoso and Sheth (2002) present a methodology and a set of algorithms for Web service discovery based on three dimensions: syntax, operational metrics, and semantics. This work on composing eWorkflows, emphasizes the discovery of Web services based on operational metrics. The algorithms employ a feature-based model to find similarities across tasks (activities) and Web service interfaces.

Run-time adaptability of a composed process is another research issue in this area. Earlier work on the METEOR (Kochut et al. 1999) project, at the LSDIS lab of the University of Georgia, has addressed this issue for workflows. The distributed task schedulers of METEOR perform the work of scheduling the execution of tasks. This allows for dynamic information to be provided to these task schedulers at run-time, creating an adaptive workflow.

The SWORD project (Shankar et al. 2002) of Stanford University is also exploring techniques for composing services. In SWORD, a service is represented as a logical rule (precondition => post condition) that expresses the inputs and the outputs associated with it. A rule-based expert system is used to automatically determine whether a process could be realized with the given services. It also returns a process plan that realizes the composition. There can be multiple ways of deriving the same composition, but SWORD currently returns an arbitrary plan since it does not have a cost model to evaluate alternative plans.

13

The existing work on Web service composition discussed above could be classified based on composition type and search type as shown in Table 2.2.

Table 2.2. Classification of Service Composition Works

	Keyword based Searching		Semantic Searching	
	Languages	Systems	Languages	Systems
Static Composition	WSFL, XLANG, BPEL4WS	SWORD, XL	DAML-S	eWorkFlow
Dynamic Composition	WSFL, BPEL4WS	eFlow		

Fensel and Bussler (2002) have proposed a Web Service Modeling Framework (WSMF), that provides a rich conceptual model for the development and description of the Web services technology. WSMF consists of four main elements: *ontology*[2], which is a key enabling technology for the semantic Web (Berners-Lee et al. 2001), *goal repositories* that define the problems that should be solved by Web services, *Web services descriptions* that define the various aspects of Web services and *mediators* which resolve interoperability problems among different services.

Testing and simulating Web processes for performance evaluation are two new areas with little previous work. Testing Web services for performance evaluation has been classified into load testing, stress testing and spike testing (Daniel et al. 2001). Simulation of composite Web services closely models simulation of workflow models (Miller et al. 1995; Miller et al. 2002).

---

[2] An ontology is a formal, explicit specification of a shared conceptualization (Gruber 1993)

The work in simulation that most closely relates to ours is described by Narayanan et al. (2002). In their work, DAML-S service descriptions of composite services are represented as Petri Nets (Petri 1962), providing decision procedures for Web services simulation, verification and composition.

4. Issues in Process Composition

This section analyzes some of the current issues involved in composing Web services. They include determination of composition type, specification of services, representation of composition, interoperability between services, efficiency and execution of the composition.

4.1 Composition Types

Web service composition can be done either statically or dynamically (Benatallah et al. 2002). The decision of whether to make a static or dynamic composition depends on the type of process being composed. If the process to be composed is of a fixed nature wherein the business partners/alliances and their service components rarely change, static composition will satisfy the needs. Static composition is done in three stages: planning, definition and implementation. In the planning phase, elementary or composite services are discovered and checked for compatibility and conformance with the process specification. In the definition phase, the specification of the composition (with the predetermined services) is produced and the implementation phase executes the composition based on the specification.

On the other hand, if the process has a loosely defined set of functions to perform, or if it has to dynamically adapt to unpredictable changes in the environment, then static composition may be too restrictive. This is because changes in statically composed systems have to be done manually, interrupting the operation of the process. This is difficult if the static process realizes

an important mission critical application. Dynamic composition is an important step in overcoming this problem. In dynamic composition, the Web services to be used for the process are decided at run-time by, for example, the process execution engine. Dynamic composition involves run-time searching of service registries to discover services.

4.2 Web Service Specification

Service providers advertise their service descriptions in public Web service registries. This enables service requesters to search for services, matching their requirements. This specification of a Web service, should include syntactic (what does it look like?), semantic (what does it mean?) and QoS (how well does it perform?) information. Quality of Service (QoS) (Cardoso, Sheth et al. 2002) attributes, which are timeliness, cost of service, and reliability provide a description of the quality that can be expected from the service.

Currently, WSDL is the widely used language for describing Web services. WSDL is used to specify a Web service's interface. It defines the syntactic information about a service such as the set of available operations, network end-points, etc. Researchers are also developing DAML-S, which is an ontology based interface description language that can describe the syntactic as well as the semantic content of a service. DAML-S also describes limited nonfunctional QoS related attributes of a service.

As explained earlier, WSFL, XLANG, BPEL4WS and DAML-S are among the languages that can specify composition of Web services. WSFL was chosen as the language for specifying composed processes in our system. We have extended WSFL's specification to include time, cost and reliability QoS attributes for the activity nodes in the process description. WSFL precedes BPEL4WS, which is a recently proposed language supporting stateful, and long

running asynchronous interactions between services involved in a process. BPEL4WS was proposed after the completion of this project, but much of the contributions of our work remain.

4.3 Representation of Composition

A Web service composition can be represented in a variety of ways. While digraphs have been used widely to model business processes (Casati et al. 2000; Sheth et al. 1996), other methods like Petri-Nets (Aalst et al. 1994), Activity/State charts (Harel 1987) are also being employed.

In our system, we represent the composition of Web services as a digraph. This representation, which is also frequently used in Workflow Management Systems (Sheth et al. 1996), represents a process in terms of activity nodes, control links and data links. The activity nodes in the digraph represent the tasks in the Web process. The control links specify the control flow within the process. They also capture various constructs such as XOR splits, AND splits, XOR joins, and AND joins that are normally associated with process designs. For example, an XOR split represents a node in the process, where based on the conditions the control can flow in only one of the several outgoing links. The data links in a digraph represent the data that flows between two activities and also the information about any mapping that needs to be applied between the output of the one activity and the input of the another activity. The details of how a Web service composition is represented as a digraph in our system is further explained in Section 5.3.

4.4 Interoperability Between Web services

In choosing Web services for a process, one has to consider how the chosen Web services will interoperate with each other, with respect to the data mappings that need to be applied between the output of one service and the input of another service. For static compositions, the process

composer can add the necessary logic to his process definition, which can handle data mapping between the pre-defined Web services. This program logic serves as an intermediary/adapter performing tasks such as extracting relevant information or changing the structure of a Web service's output, so that it is compatible with the subsequent service. This approach is restrictive in dynamic compositions, where the services are decided only at run time. In a dynamic composition, the service matching algorithms should efficiently search registries and consider services that are suitable for automatic composition.

Currently, as the registries being used for Web services base their search only on keywords and categories. This will result in returning a large set of hits for a search on a service. To address this problem, Paolucci et al. (2002) explore the addition of a semantics layer to a UDDI Web service registry, wherein services are described semantically using ontologies. This allows for more efficient matching of services. In cases where a fully automated dynamic composition is not possible because of the inability to automatically choose services or accommodate for interoperability between the services, a semi-automated composition technique could be adopted. A semi-automated technique is a computer-aided process composition technique, in which the composer will also have a role in managing the execution of the process. The composer may have to manually choose services or specify the mappings from the output of one Web service to the input of another, if the system is not able to automate the task.

4.5 Process Execution

A composed Web process can be executed either via a centralized approach or a distributed approach. The centralized approach is based on the client/server architecture, with a scheduler, which controls the execution of the components of the Web process. The controller (client) can reside either in the host where the composition is made or at a separate host to allow for a better

load balancing. The controller/scheduler invokes a Web service, gets the results, and based on the results and the Web process design specification, the controller then invokes the next appropriate Web service. This is the easiest approach for executing Web processes. eFlow (Casati et al. 200) is such a system with a centralized execution engine.

The distributed approach for Web process execution is more complex. In a distributed approach, the Web services involved are expected to collaborate and share the execution context to realize the distributed execution. In this approach each of Web service involved, hosts a coordinator component which collaborates with other coordinators to realize the execution (Benatallah et al. 2001, 2002). A slightly modified version of distributed execution involves coordinators which control a set of Web services. The process is executed in a distributed fashion by these coordinators, but internally each coordinator implements a centralized model for executing the tasks controlled by it (Benatallah et al. 2001).

4.6 Process Efficiency

It is important to know the efficiency and the Quality of Service (QoS) of a Web process before making it available to its customers. To estimate the QoS of a process, Cardoso, Miller et al. (2002) have developed a comprehensive model for the specification of QoS of workflow and Web processes. They have investigated dimensions such as time, cost, reliability and fidelity required to develop a real and usable QoS model. In our work, we have investigated how the invocation time could be practically computed and analyzed to estimate the individual components involved in it, allowing users to use the mathematical models explained in the above mentioned work. This can serve as a basis for realizing the time dimension of their QoS model.

5. System Architecture and Implementation

In this section, we introduce the architecture of our system for composing and executing Web services and explain it with an example scenario.

5.1 Architectural Overview

Figure 2.2 shows the system architecture of our composition and execution tool. The main modules in our system are the Service Composition and Execution Tool (SCET), the JSIM simulator and the Perl execution controller. SCET in turn has a process designer, simulation
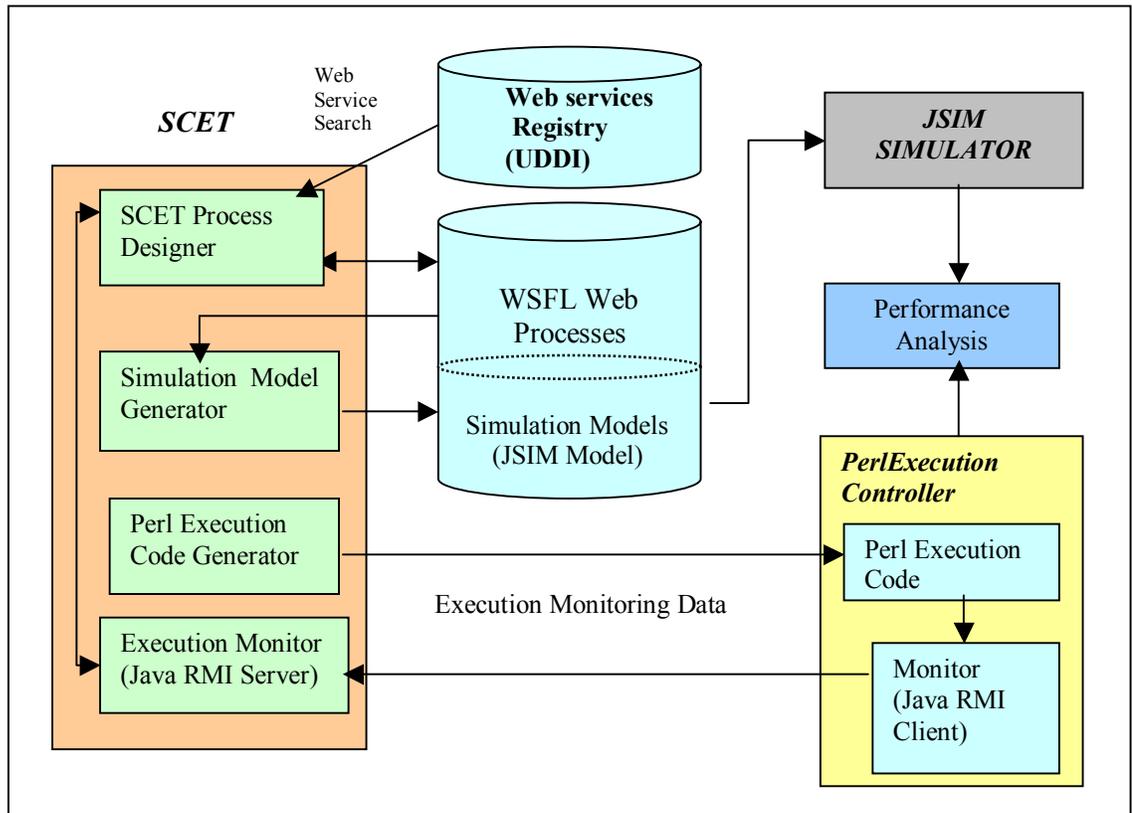

Figure 2.2. System Architecture for SCET

model generator, perl code generator and execution monitor sub-modules to help the composer to easily compose, simulate, execute and monitor a process respectively. The primary sub-modules of SCET are discussed in the following subsections.

5.1.1 Process Designer

Central to the architecture is the SCET's Process Designer, a graphical design tool allowing users to statically compose processes. Users can design a digraph representing the process and specify the necessary information associated with the nodes and the links of the graph. The designer can store the composed process as a WSFL based specification file or as an XML document in a repository (e.g., Db4XML (Sipani et al. 2002) native XML database).

5.1.2 Simulation Model Generator

The JSIM simulator used in our system requires a Java based specification of the model that is to be simulated. In SCET, we represent a composed process as a WSFL based specification. Thus, we need to convert this WSFL based process specification to a model, which the JSIM simulator can interpret The Simulation Model Generator of SCET does this work of automatically transforming the WSFL based process designs, generated by the process designer to JSIM simulation models.

5.1.3 Perl Execution Code Generator

The Perl Execution Code Generator sub module of SCET is capable of generating the Perl execution code for the composed process, allowing for easier execution after composition. The Perl Execution Code Generator traverses the WSFL specification of the process and generates a Perl program, with perl code blocks that correspond to the elements encountered in the

specification file. For example, an activity node is transformed to a Perl Web service invocation code block as shown in Figure 2.6.

5.2 Scenario

Figure 2.3 depicts the internal tasks involved in processing a customer's book purchase order ("BarnesBookPurchase" service).The activities in this process are BarnesGetPrice, CheckCredit, CheckInventory, GenerateBackOrder, ReleaseOrder and SendCreditLowInfo. The bookstore has a real Web service associated with each of the activities involved in this process. These Web services are composed to create a customer order processing Web process.
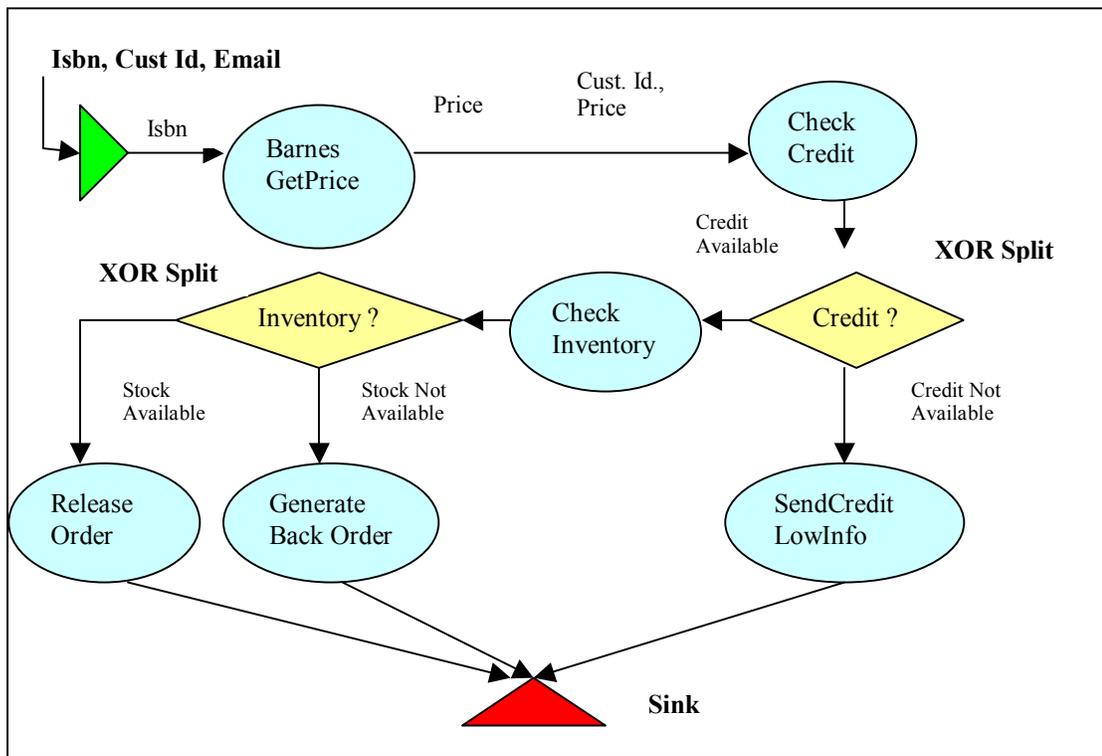

Figure 2.3. Barnes Book Purchase Web Services Composition

The price of the book chosen by the customer is retrieved using the BarnesGetPrice service. The user's account is then checked for sufficient funds using the CheckCredit service. CheckCredit is an example of an XOR split activity. After the CheckCredit service, the control flows in one of the two outgoing links depending on whether it returns success or failure. If the user has sufficient credit, the CheckInventory service is invoked; otherwise, the SendCreditLowInfo service is invoked. If the CheckInventory Web service returns true the ReleaseOrder service is invoked to send the books; otherwise the GenerateBackOrder service is invoked. The following sections explain in detail the composition of the above process using SCET.

5.3 SCET Process Composition

In our system, a Web process is represented in a digraph form using the Process designer's source nodes, sink nodes, activity nodes, data links and control links (Figure 2.4). The black transition links in the figure represent control/data links while the green transition links (not in the picture) represent the data links between the activities.

The process composer apart from laying out the process structure also provides information about the activities (as shown in the Activity Definition dialog box in Figure 2.4) and the links used in the process. An activity node stores information about the Web service implementing it. This includes the Web service's WSDL file location, the operation being invoked, and QoS information (such as mean service time, reliability factor, cost associated with the activity etc.). The QoS information, which could be obtained either from the service provider or by performing analysis tests (as explained in Section 6), is used by the simulator of the system to simulate the process behavior. For control links, the composer specifies the condition on which

the control flows in that link, while for data links the composer specifies how the output of one activity is routed to the input of another activity.
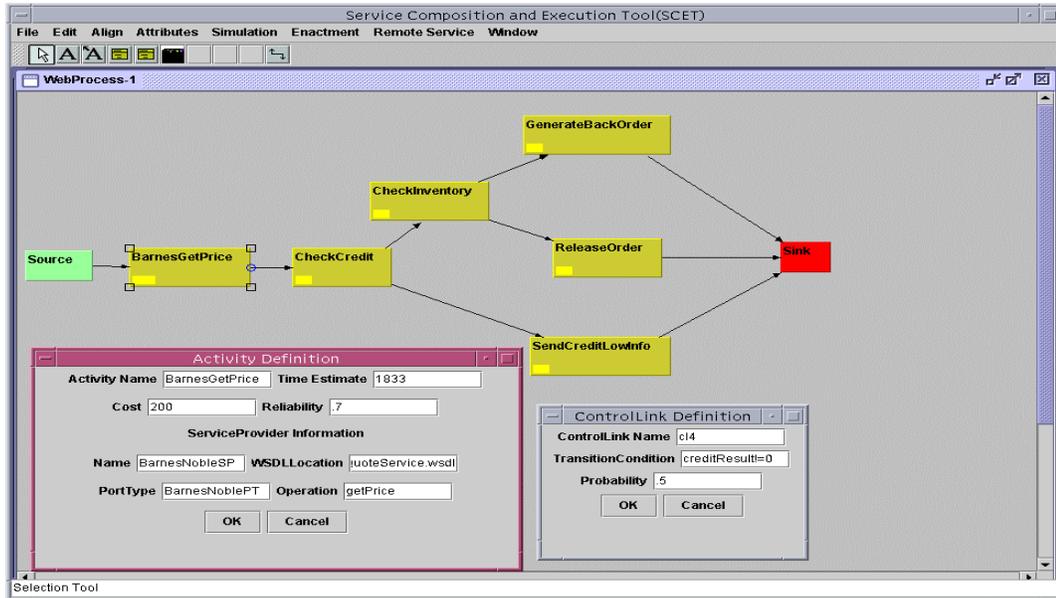

Figure 2.4. Web Process Composition using SCET

5.4 WSFL Based Specification Generation

SCET saves the process composition as a WSFL based specification. The transformation from the internal storage model to WSFL specification is straightforward. The source/sink nodes in the process design result in source/sink elements of the WSFL specification. The activity nodes of the design are converted to the activity elements of the WSFL specification. The input/output information associated with each activity is transformed to input/output message elements of the corresponding activities in WSFL. The links connecting the nodes in the digraph generate either control link or data link elements based on the type of link used by the composer while designing the process. Figure 2.5 shows a code fragment of the WSFL based specification generated for the

24

"BarnesBookPurchase" process composition. In the next section, we discuss the Web process

execution approach that has been followed in SCET.

```
<definitions>
<message name="CheckInventoryInput">
                <part name="bookIsbn" element="string"/>
</message>
.....................
.....................
<serviceProvider name="CheckInventorySP">
  <locator
        type="static" service=
                "http://www.cs.uga.edu/sent/xmethods/CheckInventory.wsdl"/>
</serviceProvider>
.....................
.....................
<activity name="CheckInventory" X1="336.0" Y1="104.0"
                X2="481.0" Y2="152.0" time="79" cost="200" reliability=".7" type="Facility">
        <input message="CheckInventoryInput"/>
        <output message="CheckInventoryOutput"/>
        <join condition="true" when="deferred"/>
        <performedBy serviceProvider="CheckInventorySP">
                <implement>
                        <export>
                                <plugLink>
                                <target
                                PortType="CheckInventoryPT"
                                Operation="hasBook"/>
                                </plugLink>
                        </export>
                </implement>
        </performedBy>
</activity>
.....................
.....................
<controlLink
        name="cl6" source="CheckInventory" target="GenerateBackOrder"
        condition="CheckInventoryResult==0" probability=".2"
        X1="431.0" Y1="104.0" X2="490.0" Y2="72.0" />
.....................
.....................
<dataLink name="dl2" X1="238.0" Y1="178.0" X2="336.0" Y2="149.0"
                source="SearchAmazon" target="CheckInventory">
        <mapInfo sourcePart="bookIsbn" targetPart="bookIsbn"/>
</dataLink>
</definitions>
```

Figure 2.5. WSFL Process Description

5.5 Web Process Execution

A Web process execution is similar to a workflow enactment, the difference being the

components of a workflow are activities while the components of a Web process are services.

Web services differ from workflow activities in their distribution, autonomy and heterogeneity.

Substantial research on workflow enactment has been done in the Large Scale Distributed

Information Systems Lab (LSDIS) at the University of Georgia (Sheth et al. 1996; Miller et al. 1996; Miller et al. 1998; and Kochut et al. 1999). Both centralized and distributed enactment engines were developed as part of the METEOR project (Kochut et al. 1999; Miller et al. 1998; and Sheth et al. 1996).

We have followed the centralized execution approach in our system. SCET is capable of automatically generating Perl execution code from WSFL based process descriptions. In our implementation, a Perl controller manages the entire Web process execution. The Perl execution controller interconnects the activities using a pipe-and-filter model, wherein the output of one Web service is appropriately routed to the input of another Web service based on the control flow and data mappings that were specified by the designer.

There are many alternatives for choosing a language for executing Web processes. Java, Perl, and Python were among the languages considered for our system. We chose Perl as our execution language, because its easy to use SOAP::Lite (Kulchenko 2002) modules help in quickly scripting the process from its WSFL description. Java is another viable option, which does include powerful execution features, but the execution code size of Java is bulkier than that of Perl (see http://www.cs.uga.edu/~jam/sent/thesis/compare.html for comparison).

Perl allows us to capture the execution logic of the process including AND/XOR splits and AND/XOR joins. We have tested the implementation of these constructs in Perl using its process management utilities: fork() and wait(). Forking a process helps realize parallel execution for AND splits. For joins in a process, a Perl process (representing a Web service invocation) can wait for other processes, thereby allowing one activity to synchronize with other activities. Perl's thread management features (Sugalski 1999) could also help realize the above-mentioned constructs.

As invocation of a Web service via Perl is simple, the transformation of WSFL process specification into Perl execution code, involves converting the WSFL constructs to Perl code blocks. In the "BarnesBookPurchase" Web process, we invoke the getPrice( ) method on the BarnesGetPrice Web service with bookIsbn as its string parameter. This Web service invocation is converted to the following simple Perl snippet (Figure 2.6), which is capable of realizing the BarnesGetPrice activity of the process and storing its result for further usage. In the next section, we discuss techniques for evaluating the performance of a composed process.

```
my $GetBNPrice = SOAP::Lite
        -> service ('http://www.xmethods.net/sd/2001/BNQuoteService.wsdl');
$GetBNPriceResult = $GetBNPrice->getPrice($isbn);
```

Figure 2.6. Perl Execution Code Snippet

6. Performance Evaluation

Performance evaluation of Web services can help implementers understand the behavior of the activities in a composed process. Since the performance of a single Web service has the potential to affect the performance of an entire Web process, it is wise to evaluate the performance of the services within a process before making it available for commercial usage. "The most commonly used approach to obtain performance results of a given Web service is performance testing, which means to run tests to determine the performance of the service under specific application and workload conditions" (Daniel et al. 2001). From the user's (process composer's) perspective, the total execution time of a process is a measure of its efficiency.

6.1 Time Analysis

The execution time taken by a single Web service invocation has three components: Service Time (*S*), Message Delay Time (*M*) and Waiting Time (*W*) (Cardoso and Sheth 2002; Chandrasekaran et al. 2002). Service Time is the time that the Web service takes to perform its task. Message Delay Time is the time taken by the SOAP messages, in being sent/received by the invocation call. It is determined by the size of the SOAP message being transmitted/received and the load on the network through which the message is being sent/received. Waiting Time is the Web service invocation delay caused by the load on the system where the Web service is deployed. Thus, the Total Invocation Time (*T*) for a Web service σ is given by the following formula.

$$T(\sigma) = M(\sigma) + W(\sigma) + S(\sigma)$$

Evaluating the above three components of *T* for a Web service invocation, will help in analyzing the efficiency of a Web process. We have performed tests to determine each of the above three components for all the Web service invocations used in the process. Message Delay Time was estimated by invoking a ping function for each Web service. XML messages were sent and received, but the Web service performed no work. Service Time was estimated by running tests against the Web service in an environment where the load and waiting delay for the service were controlled. Waiting Time was estimated by running the test in an environment where the Web service was loaded with requests.

Figure 2.7 shows a snapshot of the sample test result for the example Web process. The time values in the figure are the sum of the Service Time (*S*) and the Message Delay Time (*M*), as the hosts were controlled and the services were not loaded during the test. The *BarnesGetPrice* service used in our example is a Web service hosted by Xmethods (Hong et al. 2000), while the other Web services were hosted locally in our intranet The WSDL files for the three services used in the above test are available at, www.xmethods.net/sd/2001/BNQuoteService.wsdl,

www.cs.uga.edu/~jam/sent/CheckCredit.wsdl,

www.cs.uga.edu/~jam/sent/SendLowCreditInfo.wsdl respectively.

Similar tests are performed when the Web services are loaded to determine the three time measures (Service Time, Message Delay Time, Waiting Time) (Figure 2.8) for each service invocation. This provides information for analyzing the performance of individual Web services being used in the composed process. In Figure 2.8, which shows the distribution of the overall time for one of our tests, the *SendLowCreditInfo* Web service has a high Waiting Time. This may indicate that either the Web service or the system hosting that service is not able to handle the load. Replacing the *SendLowCreditInfo* Web service with an equivalent service, which can handle more load, may improve the quality of service of the entire Web process.

The above work on practically analyzing the time dimension of a Web service's QoS model serves as a starting point to realize the mathematical models of Cardoso, Miller et al. (2002). The individual time components estimated using our approach could be used in their model, to compute the overall time estimate of the entire process. The algorithm implemented in their work, could also be used to estimate the overall Service Time (S), Waiting Time (W) and Message Delay Time (M) for the entire Web process, providing more information on the QoS of the process.

```
10: Num of Iterations
        :BarnesGetPrice    :CheckCredit        :SendcreditLowInfo
        :1.913394          :1.342986           :0.759598
        :1.714697          :0.682717           :0.486978
        :1.562682          :1.42475            :0.692787
        :1.886854          :0.643781           :0.473613
        :1.828123          :0.816607           :0.525161
        :1.765729          :0.876123           :0.486001
        :1.502503          :0.93461            :0.534121
        :1.617242          :0.644926           :0.464647
        :2.485304          :0.82555            :0.691154
        :2.126191          :0.855064           :0.472463
        :1.84127           :0.918563           :0.57291


Total       :18.330595         :8.622691           :5.399835
Average     :1.8330595         :0.8622691          :0.5399835
Minimum     :1.502503          :0.643781           :0.464647
Maximum     :2.4853            :1.42475            :0.692787
STD         :0.276012          :0.2134661          :0.0824032
```

Figure 2.7. Test Results for Total Invocation Time (*T*)
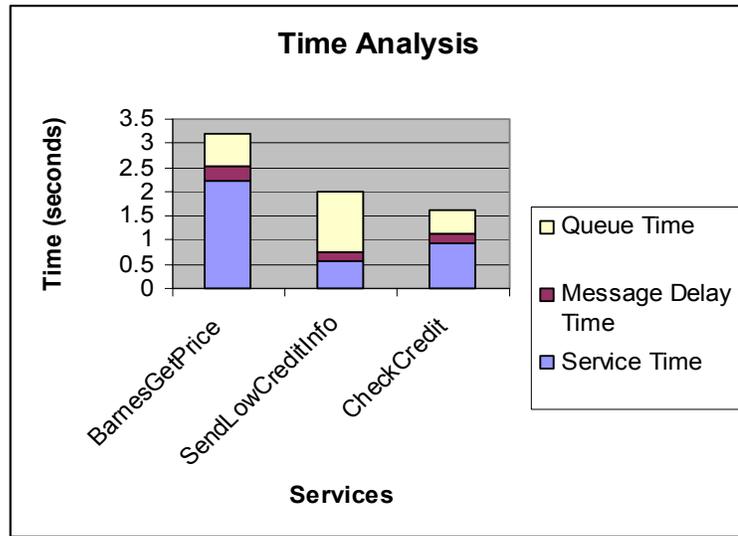

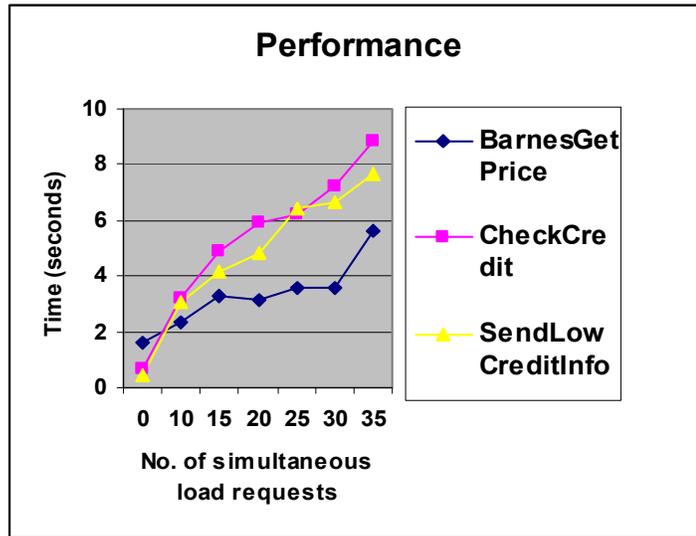Figure 2.8. Performance Time Analysis

Figure 2.9. Load Testing Results for the Barnes Book Purchase Web Process

Load Testing is another approach that can be used to measure the performance of a Web service. The Web services are gradually loaded with client invocations and the performance results are measured. After a certain load point the performance of the Web service will start degrading. This point is the load range to which the Web service is performing effectively. This testing is a useful means of comparing and determining which Web service to choose for a process. Figure 2.9 shows the time (sum of service time and message delay time) taken by the Web services with respect to the number of simultaneous load requests. Thus the above performance analysis tests, help the composer to

- identify overloaded Web service hosts, by estimating the Waiting Times for each service invocation;

- identify message communication overheads, by estimating the Message Delay Times for each service invocation;

- estimate the load to which each service is able to perform effectively.

31

6.2 Monitoring Process Invocation

SCET is capable of monitoring a process that is being executed. It can visually show the expected number of Web service invocations present in the host of the service provider. The expected number of Web service invocations $L_n(\sigma)$ (present in the host of a Web service σ, for the nth invocation of the process)

$$L_n(\sigma) \quad = \quad \frac{W(\sigma_n) + S(\sigma_n)}{\bar{S}_n(\,)}$$

where $S(\sigma_n)$, $W(\sigma_n)$, $\bar{S}_n(\sigma)$ are the Service Time, Waiting Time and the moving average Service Time. The moving average Service Time for a Web service σ, is computed by the following formula

$$\bar{S}_n(\sigma) \quad = \quad \frac{\sum_{i=1}^{n} (T(\sigma_i) - M(\sigma_i))}{n}$$

where n is the number of test iterations.

The queue bar (yellow, left-bottom part) in each activity of Figure 2.4, indicates this value of the expected number of invocations queued for each Web service. During execution of a Web process, the size of the queue varies according to the load on each Web service. The SCET designer is linked with the Perl execution controller to realize this.

We have used the Java RMI mechanism for realizing the execution monitoring functionality in our system. The Perl execution controller executes the process and obtains the monitoring data. This data is then passed to the designer through a Java RMI Client. The Java RMI client communicates the monitoring data with the Java RMI Server of the SCET designer to allow users to visually monitor the process execution. We have adopted this approach in our system as a bridge to communicate information between the Perl execution controller and our Java based SCET designer.

The performance analysis techniques described above provide feedback on the quality of the composed processes. These techniques also have some associated disadvantages. They can yield good estimates if the Web services involved are under our control. This is because all these tests need to be executed in a controlled manner, taking care of the load on the system when the testing is done. The distribution and autonomy of Web services makes meeting this requirement difficult. If there is a Web service over which we do not have control, then it is difficult to get accurate performance testing results. The server hosting the Web service may be heavily loaded with other programs and it may be incorrectly accounted as service time in the analysis.

Further, if the services involved in the process include "world-altering" services (such as flight-booking service, money transfer service) or if there is a cost involved in invoking the individual services, then analyzing the performance of the process by executing it may not be feasible. To overcome this difficulty, simulation based testing, which is described in the next section could be used as an alternative technique for evaluating the process.

7. Simulation

Simulation helps in determining how the composed Web services will perform when they are deployed. It plays an important role by exploring the "what-if" questions during the process composition phase, which may not be feasible or too costly to do performance tests with real Web processes. Simulation can provide feedback on the process that was composed allowing the composer to modify his/her process design by

- Replacing services that do not provide required service time averages.
- Modifying the process structure (such as altering the number of activities involved, and changing the control flow) based on the simulation runs.

7.1 JSIM Simulation

SCET is integrated with the latest version of JSIM, a Java-based simulation and animation environment (Nair et al. 1996; Miller et al. 1997) that contains several features to support simulation of Web processes. JSIM currently simulates both the centralized execution (Figure 2.10) of a Web process (implemented in our system), as well as the distributed version (Figure 2.11) of it. The basic components of a JSIM simulation model and how they are related to our Web process model is described below:

- Source Nodes: Generate entities using an inter-arrival time produced by a random variate. This is analogous to the Source activity nodes of the Web process specification. Source nodes serve as the origination point for simulation entities (objects), while Source activity nodes serve as the node from where the process executions start.

- Server Nodes/Facility Nodes: Provide service to simulation entities using a service time produced by a random variate. Servers have one or several service units without queuing functionality. Facility nodes provide a queue for waiting entities. The activities (Web services) in the process design, map to either Server or Facility nodes of the JSIM model depending on the queuing behavior of the service.

- Sink Nodes: Consume simulation entities and capture statistical information about the entities.

- Transports: Edges that connects two nodes. The control links in the process design model map to transports in the JSIM simulation model.

- SimObjects: Instances of simulation entities. An invocation of a Web process is represented as a SimObject in the JSIM model.
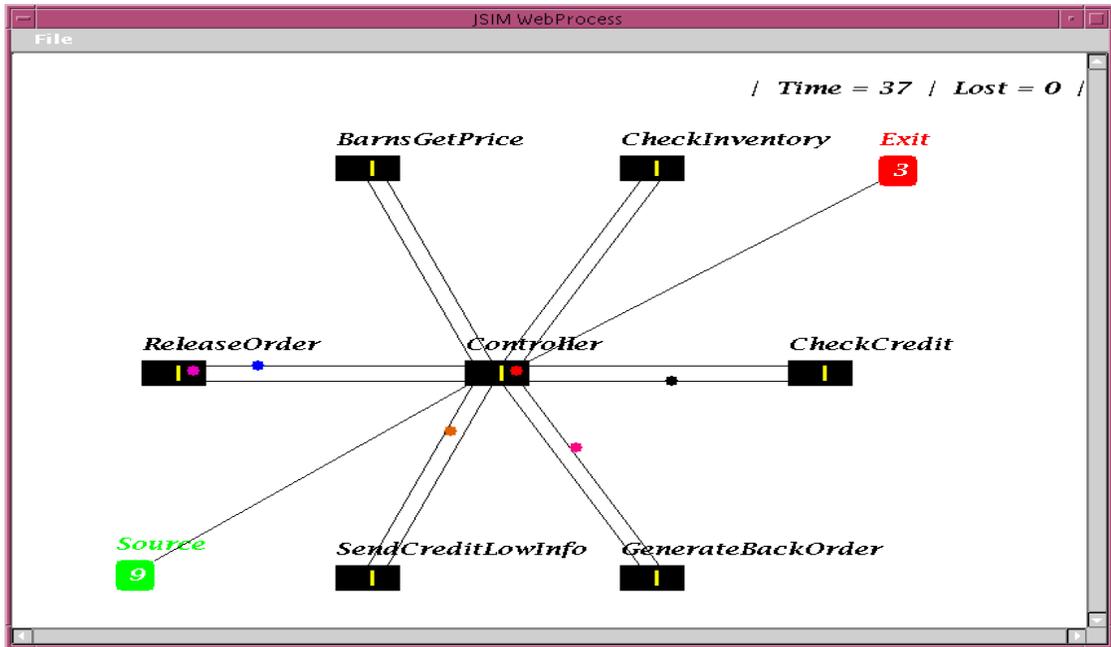
Figure 2.10. JSIM Centralized Execution Simulation

7.1.1 Simulating Processes

After composing the process using the designer, the composer can make use of the simulation

model generator to convert the WSFL process specification to JSIM simulation model. The JSIM

simulation model takes as input the service time distribution functions characterizing the Web

services. These distribution functions can be computed by performance analysis tests as explained

in the appendix or obtained from the service providers. For each of the control links involved in

the process, JSIM requires an associated probability value for simulating the process execution.

In our experiments, we have computed the probability values associated with each of the control

links, by executing the process on a test basis.

After its simulation run, JSIM generates statistical information (Figure 2.12) about the

completed simulation. This includes information about minimum, maximum, mean, and standard

deviation of the time estimates of each of the activities involved in the process model. During this

simulation phase, the composer can try and examine the various possible alternatives that could

be applied to the composition. This includes analyzing how the process and individual Web

services will perform when the process structure is changed, or when Web services with better

service time averages are being used in the process. The statistical information generated for

these test runs provides feedback to the composer about the behavior of the process for

hypothetical cases. This creates a feedback loop when composing Web services, thereby allowing

the composer to iterate through this feedback loop (as shown in Figure 2.1) until an efficient

process suitable for execution is obtained.
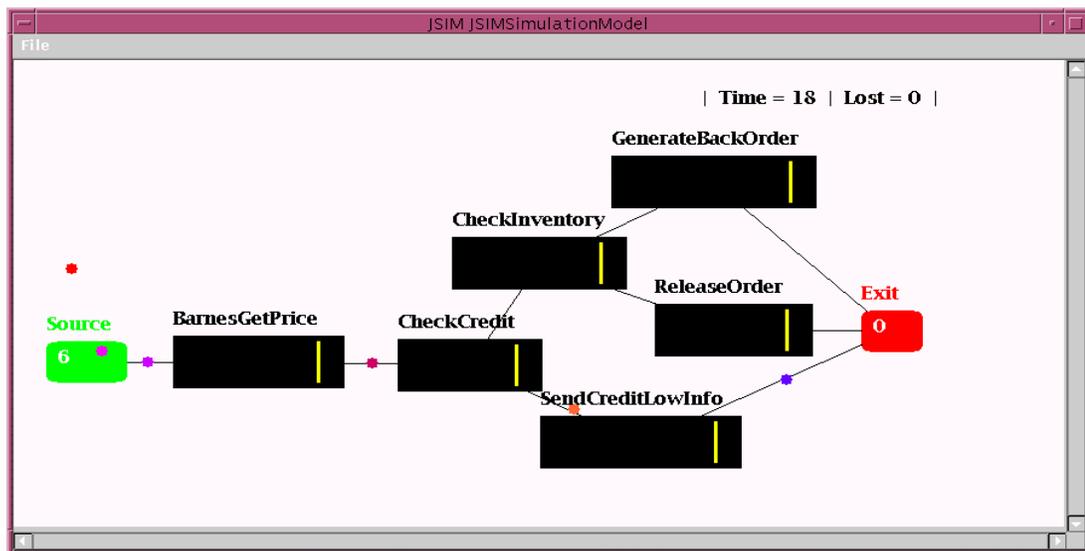

Figure 2.11. JSIM Distributed Execution Simulation

7.1.2 Comparing Simulation and Real Time Invocation

The JSIM Simulation displays dynamically the number of entities (invocations) being queued up

in an activity. The SCET designer is also capable of displaying the expected number of real

invocations present in the host of a Web service. Thus, we can do both simulation and execution

and can compare the two models quantitatively (Figure 2.12, Figure 2.7), as well as visually, to determine the validity of the simulation model. If the simulation model is found not to be valid, we can make changes to the simulation parameters to improve its accuracy.

In the comparison test made above, the actual process was executed several times with a three second time period between successive process executions. The corresponding simulation of the process, was done by modeling JSIM to generate several simulation entities with an inter-arrival time of three seconds. The actual service time averages were found to be 1.833, .862, and .539 seconds for BarnesGetPrice, CheckCredit and SendCreditLowInfo Web services respectively. The simulation results resulted in service time averages of 1.883, .861 and .538 seconds, which closely approximate the actual execution values. These initial tests suggest that the time measures from both the simulation runs and the actual execution, match closely for all the activities involved in the process, which is an indication, that the simulation closely models the actual behavior of the process. Additional studies are being carried out.

				Statistics			
NoSamples	MinValue	MaxValue	MeanValue	Deviation	Interval	Precision	StatName
10.0	3000.0	3000.0	3000.0	0.0	0.0	0.0	Source (dur)
10.0	0.0	0.0	0.0	0.0	0.0		Exit (dur)
10.0	1832.938	1833.028	1833.002	0.025	0.019	0.0	BarnesGetPrice (du
36033.0	0.0	1.0	0.509	0.5	0.005	0.010	BarnesGetPrice (oc
10.0	861.947	862.058	861.997	0.034	0.026	0.0	CheckCredit (dur)
38496.0	0.0	1.0	0.224	0.417	0.004	0.019	CheckCredit (occ)
0.0	0.0	0.0	0.0	0.0	0.0	0.0	CheckInventory (d
0.0	0.0	0.0	0.0	0.0	0.0	0.0	CheckInventory (oc
10.0	538.951	539.032	538.991	0.029	0.022	0.0	SendCreditLowInfo
41035.0	0.0	1.0	0.131	0.338	0.003	0.025	SendCreditLowInfo
0.0	0.0	0.0	0.0	0.0	0.0	0.0	GenerateBackOrde
0.0	0.0	0.0	0.0	0.0	0.0	0.0	GenerateBackOrde
0.0	0.0	0.0	0.0	0.0	0.0	0.0	ReleaseOrder (dur)
0.0	0.0	0.0	0.0	0.0	0.0	0.0	ReleaseOrder (occ)
30.0	400.0	400.0	400.0	0.0	0.0	0.0	path1 (dur)
0.0	0.0	0.0	0.0	0.0	0.0	0.0	path2 (dur)
0.0	0.0	0.0	0.0	0.0	0.0	0.0	path3 (dur)
0.0	0.0	0.0	0.0	0.0	0.0	0.0	path4 (dur)
0.0	0.0	0.0	0.0	0.0	0.0	0.0	path5 (dur)
0.0	0.0	0.0	0.0	0.0	0.0	0.0	path6 (dur)
140.0	400.0	400.0	400.0	0.0	0.0	0.0	path7 (dur)
50.0	400.0	400.0	400.0	0.0	0.0	0.0	path8 (dur)
40.0	400.0	400.000	400.000	0.0	0.0	0.0	path9 (dur)

Figure 2.12. JSIM Statistics Table

8. Conclusions and Future Work

Web services composition is a new research area that combines Web services technology and process composition. In this paper, we have focused on problems related to composition representation, specification of service compositions and execution of processes, using our Service Composition and Execution Tool (SCET). We have described the key components of SCET and how it can be used to compose efficient Web processes. SCET allows users to statically compose services to form processes and store them as WSFL based specifications. As QoS specification is an integral part of describing a process, we have enhanced WSFL to include QoS attributes for specifying requirements on quality such as time, cost and reliability. SCET executes a process by following a centralized process execution technique. Services should be composed taking into consideration the overall performance efficiency of the resulting process. In this direction, we present performance analysis approaches that help a composer to evaluate the performance of a process. We also describe how simulation can be used with a Web process composition system to carry out efficient compositions.

However, there are still several issues that need to be addressed. First, SCET is a static composition tool. SCET needs to support dynamic composition, so that users can just specify their requirements on the services and not the actual service. Currently, SCET supports WSFL based process specification. BPEL4WS is a newly proposed composition standard for Web services. SCET needs to be enhanced to support this new standard. SCET also needs changes to its framework, to accommodate for this kind of evolving nature of process specification languages. From the execution point of view, the execution code generator has restrictions on its functionality. It is capable of handling services which return primitive data types. This needs to be improved to support array values and objects. Also, we need to address the issue of data mapping between Web services, so that users can specify their data transformation function to

map information between Web services. In this work, we have provided approaches for realizing the time QoS dimension of a Web process. We would need to enhance it to address other QoS dimensions such as reliability and cost.

CHAPTER 3

CONCLUSIONS

Web services composition is a new research area that combines Web services technology and process composition. In this paper, we have focused on problems related to composition representation, specification of service compositions and execution of processes, using our Service Composition and Execution Tool (SCET). We have described the key components of SCET and how it can be used to compose efficient Web processes. SCET allows users to statically compose services to form processes and store them as WSFL based specifications. As QoS specification is an integral part of describing a process, we have enhanced WSFL to include QoS attributes for specifying requirements on quality such as time, cost and reliability. SCET executes a process by following a centralized process execution technique. Services should be composed taking into consideration the overall performance efficiency of the resulting process. In this direction, we present performance analysis approaches that help a composer to evaluate the performance of a process. We also describe how simulation can be used with a Web process composition system to carry out efficient compositions.

However, there are still several issues that need to be addressed. First, SCET is a static composition tool. SCET needs to support dynamic composition, so that users can just specify their requirements on the services and not the actual service. Currently, SCET supports WSFL based process specification. BPEL4WS is a newly proposed composition standard for Web services. SCET needs to be enhanced to support this new standard. SCET also needs changes to its framework, to accommodate for this kind of evolving nature of process specification

languages. From the execution point of view, the execution code generator has restrictions on its functionality. It is capable of handling services which return primitive data types. This needs to be improved to support array values and objects. Also, we need to address the issue of data mapping between Web services, so that users can specify their data transformation function to map information between Web services. In this work, we have provided approaches for realizing the time QoS dimension of a Web process. We would need to enhance it to address other QoS dimensions such as reliability and cost.

REFERENCES

Aalst, W., Hee, V., and Houben, G. (1994) 'Modelling Workflow Management Systems with High-level Petri Nets', in *Proceedings of the second Workshop onComputer Supported Cooperative Work, Petrinets and related formalisms*, pp. 31-50.

Ankolekar, A., Burstein, M., Hobbs, J., Lassila, O., Martin, D., McDermott, D., McIlraith, S., Narayanan, S., Paolucci, M., Payne, T., and Sycara, K. (2002) 'DAML-S: Web Service Description for the Semantic Web', in *International Semantic Web Conference*, Sardinia, Italy, pp. 348-363.

Austin, D., Barbin, A., Ferris, C., and Garg, S. (2002) 'Web Services Architecture Requirements', http://www.w3c.org/TR/wsa-reqs, accessed 22 September 2002.

Benatallah, B., Dumas, M., Fauvet, M., and Paik, H. (2001) 'Self-Coordinated and Self-Traced Composite Services with Dynamic Provider Selection', Technical report, School of Computer Science & Engineering, The University of New South Wales.

Benatallah, B., Dumas, M., Fauvet, M., and Rabhi, F. (2002) 'Towards Patterns of Web Services Composition', in *Patterns and Skeletons for Parallel and Distributed Computing,*Springer Verlag, UK (to appear).

Berners-Lee, T., Handler, J., and Lassila, O. (2001) 'The Semantic Web', http://www.sciam.com/article.cfm?articleID=00048144-10D2-1C70-84A9809EC588EF21, accessed 22 September 2002.

Cardoso, J., Miller, J., Sheth, A., and Arnold, J. (2002) 'Modeling Quality of Service for Workflows and Web Service Processes', in *The VLDB Journal* (submitted).

Cardoso, J. and Sheth, A. (2002) 'Semantic e-Workflow Composition', in *Journal of Intelligent Information Systems* (submitted).

Cardoso, J., Sheth, A., and Miller, J. (2002) 'Workflow Quality of Service', in *International Conference on Enterprise Integration and Modeling Technology and International Enterprise Modeling Conference*, Valencia, Spain, Kluwer Publishers (to appear).

Casati, F., Ilnicki, S., Jin, L., Krishnamoorthy, V., and Shan, M. (2000) 'Adaptive and Dynamic Service Composition in eFlow', in *Proceedings of the International Conference on Advanced Information Systems Engineering*, Stockholm, Sweden, pp. 13-31.

Chandrasekaran, S., Silver, G., Miller, J., Cardoso, J., and Sheth, A. (2002) 'Web Service Technologies and their Synergy with Simulation' in *Proceedings of the 2002 Winter Simulation Conference, San* Diego, CA. (to appear).

Christensen, E., Curbera, F., Meredith, G., and Weerawarana, S. (2001) 'Web Services Description Language (WSDL) 1.1.', http://www.w3.org/TR/wsdl, accessed 22 September 2002.

Colan, M. (2001) 'An Overview of Web Services', http://www-106.ibm.com/developerworks/webservices, accessed 22 September 2002.

Curbera, F., Goland, Y., Klein, J., Leymann, F., Roller, D., Thatte, S., and Weerawarana, S. (2002) 'Business Process Execution Language for Web Services', http://msdn.microsoft.com/webservices/default.asp?pull=/library/en-us/dnbiz2k2/html/bpel1-0.asp, accessed 22 September 2002.

Daniel, A. and Virgilio, A. (2001) *Capacity Planning for Web Services: metrics, models, and methods*, Prentice Hall, Englewood Cliffs, NJ.

Fensel, D. and Bussler, C. (2002) 'The Web Service Modeling Framework WSMF',
http://www.cs.vu.nl/~dieter/ftp/paper/wsmf.pdf, accessed 22 September 2002.

Florescu, D., Grunhagen, A., and Kossman, D. (2002) 'XL: An XML Programming
Language for Web Service Specification and Composition', in *Proceedings of the
Eleventh International World Wide Web Conference*, Honolulu, HI, pp. 65-76.

Gruber, T. (1993) 'Towards Principles of the Design of Ontologies used for Knowledge
Sharing', Technical Report KSL-93-04, Knowledge Systems Laboratory, Stanford
University.

Harel, D. (1987) 'State Charts: A Visual Formalism for Complex Systems', *Science of
Computer Programming*, Vol. 8, pp. 231-274.

Hong, T. and Hong, J. (2000) 'X Methods', http://www.xmethods.net, accessed 22
September 2002.

Jonathan, R., Joe, L., and David, S. (1998) 'XML Query Language (XQL)',
http://www.w3.org/TandS/QL/QL98/pp/xql.html, accessed 22 September 2002.

Kochut, K., Sheth, A., and Miller, J. (1999) 'Optimizing Workflow', in *Component
Strategies*, Vol. 1, No. 9, pp. 45-57 (SIGS Publications Inc).

Kulchenko, P. (2002) 'SOAP::Lite for Perl', http://www.soaplite.com, accessed 22
September 2002.

Leymann, F.(2001) 'Web service flow language (WSFL) 1.0',
http://www-4.ibm.com/software/solutions/webservices/pdf/WSFL.pdf, accessed 22
September 2002.

Miller, J., Cardoso, J., and Silver, G. (2002) 'Using Simulation to Facilitate Effective
Workflow Adaptation', in *Proceedings of 35th Annual Simulation Symposium*, San
Diego, CA, pp. 177-181.

Miller, J., Nair, R., Zhang, Z., and Zhao, H. (1997) 'JSIM: A Java-based Simulation and Animation Environment', in *Proceedings of the 30th Annual Simulation Symposium*, Atlanta, GA, pp. 31-42.

Miller, J., Palaniswami, D., Sheth, A., Kochut, K., and Singh, H. (1998) 'WebWork: METEOR's Web-based Workflow Management System', in *Journal of the Intelligent Information Management Systems*, Vol. 10-2, pp. 185-215.

Miller, J., Sheth, A., Kochut, K., Wang, X., and Murugan, A. (1995) 'Simulation Modeling within Workflow Technology', in *Proceedings of the 1995 Winter SimulationConference*, Arlington, Virginia, pp. 612-619.

Miller, J., Sheth, A., Kochut, K., and Wang, X. (1996) 'Corba-Based Run-Time Architectures for Workflow Management Systems' in *Journal of Database Management*, Special Issue on Multidatabases, Vol. 7-1, pp. 16-27.

Nair, R., Miller, J., and Zhang, Z. (1996) 'JSIM: A Java-based Query Driven Simulation Environment', in *Proceedings of the 1996 Winter Simulation Conference*, Coronado CA, pp. 786-793.

Narayanan, S. and Mcllraith, S. (2002) 'Simulation, Verification and Automated Composition of Web Services', in *Proceedings of the Eleventh International World Wide Web Conference*, Honolulu, HI, pp. 77-88.

Paolucci, M., Kawamura, T., Payne, T., and Sycara, K. (2002) 'Semantic Matching of Web Services Capabilities', in *Proceedings of the First International Semantic Web Conference,* Sardinia, Italia, pp. 333-347.

Petri, C. (1962) 'Kommunickation mit Automaten', PhD thesis, Institut fur instrumentelle Mathematik, Bonn.

Piccinelli, G. (1999) 'Service Provision and Composition in Virtual Business Communities',

    Technical Report HPL-1999-84, Hewlett-Packard,

    http://www.hpl.hp.com/techreports/1999/HPL-1999-84.html, accessed 22

    September 2002.

Sahni, S. and Franta, W. (1977) 'An Efficient Algorithm for the Kolmogorov-Smirnov and

    Lilliefors Tests', *ACM Transactions on Mathematical Software,* Vol. 3, No. 1,

    pp. 60 - 64.

Shankar, P. and Fox, A. (2002) 'SWORD: A Developer Toolkit for Web Service Composition',

    in *Proceedings of the Eleventh International World Wide Web Conference*, Honolulu, HI.

Sheth, A., Kochut, K., Miller, J., Worah, D., Das, S., Lin, C., Lynch, J., Palaniswami, D., and

    Shevchenko, I. (1996) 'Supporting State-wide Immunization Tracking using Multi-

    paradigm Workflow Technology', in *Proceedings of the 22nd International Conference*

    on Very Large Databases, Bombay, India, pp. 263-273.

Sipani, S., Verma, K., Chandrasekaran, S., Zeng, X., Zhu, J., Che, D., and Wong, K. (2002)

    'Designing an XML Database Engine: API and Performance', in *Proceedings of the 40th*

    Annual Southeast ACM Conference, Raleigh, NC, pp. 239-245.

Sugalski, D. (1999) 'Tutorial on Threads in Perl',

    http://www.xav.com/perl/lib/Pod/perlthrtut.html, accessed 22 September 2002.

Tidwell, D. (2000) 'Web Services – The Web's Next Revolution',

    http://www-106.ibm.com/developerworks/webservices, accessed 22 September 2002.

Thatte, S. (2001), 'XLANG: Web Services for Business Process Design',

    http://www.gotdotnetcom/team/xml_wsspecs/xlang-c/default.htm, accessed 22

    September 2002.

Waldt, D. and Drummond, R. (2001) 'EBXML: The Global Standard for Electronic

    Business', http://www.ebxml.org/presentations/global_standard.htm, accessed 22

    September 2002.

Yang, J. and Papazoglou, M. (2001) 'Web Components: A Substrate for Web Service Reuse and

    Composition', in *Proceedings of the 14th International Conference on Advanced*

    Information Systems Engineering, Toronto, Canada, pp. 21-36.

APPENDIX A: DISTRIBUTION FUNCTION ANALYSIS

If a Web process is invoked many times over a period of time, we can determine the distribution

function for the service times associated with each Web service. This distribution function is

mainly deduced so that it could be plugged into the simulation model, to correctly simulate the

process. Our tests on the example process showed that the *BarnesGetPrice's* Service Time (S)

closely models a Normal distribution with mean = 1.84, standard deviation = .277, and p = 0.32.

(shown in Figure A.1). We used the Kolmogorov-Smirnov Goodness-of-Fit test (Sahini et al.

1977) in our analysis to check whether the data models the normal distribution. This distribution

function (normal) computed from the performance analysis tests or obtained from the service
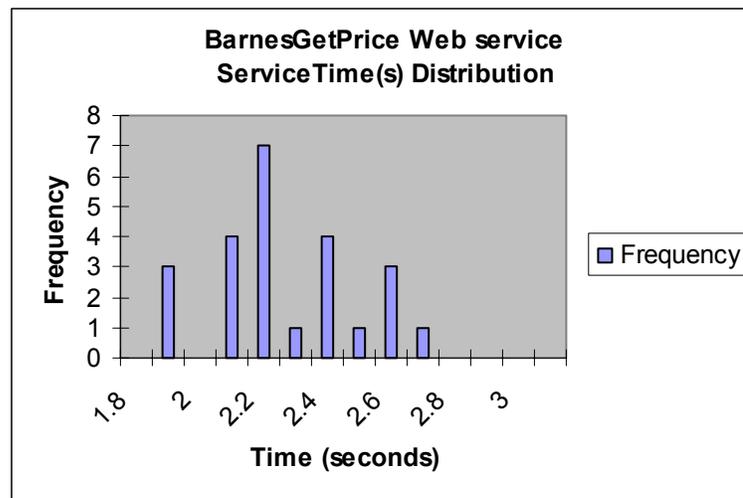
provider is used in the simulation model to simulate a more statistically appropriate model.

Figure A.1 Distribution Analysis