

# Reducing Data Transfer in Service-Oriented Architectures: The Circulate Approach

Adam Barker, Jon B. Weissman and Jano I. van Hemert



**Abstract**—As the number of services and the size of data involved in workflows increases, centralised orchestration techniques are reaching the limits of scalability. When relying on web services without third-party data transfer, a standard orchestration model needs to pass all data through a centralised engine, which results in unnecessary data transfer and the engine to become a bottleneck to the execution of a workflow.

As a solution, this paper presents and evaluates Circulate, an alternative service-oriented architecture which facilitates an orchestration model of central control in combination with a choreography model of optimised distributed data transport. Extensive performance analysis through the PlanetLab framework is conducted on a Web service-based implementation over a range of Internet-scale configurations which mirror scientific workflow environments. Performance analysis concludes that our architecture's optimised model of data transport speeds up the execution time of workflows, consistently outperforms standard orchestration and scales with data and node size. Furthermore, Circulate is a less-intrusive solution as individual services do not have to be reconfigured in order to take part in a workflow.

**Index Terms**—Service-oriented architecture, orchestration, choreography, workflow optimisation.

## 1 INTRODUCTION

Many problems at the forefront of science, engineering and medicine require the integration of large-scale data and computing. The majority of these data sets are physically distributed from one another, owned and maintained by different institutions, scattered throughout the globe [15]. Scientists and engineers require the ability to access, compose and process these distributed data sets in order to discover correlations, enable decision making and ultimately progress scientific discovery.

In order to integrate software and data, academia and industry have gravitated towards service-oriented architectures. Service-oriented architectures are an architectural paradigm for building software applications from a number of loosely coupled distributed services. This paradigm has seen wide spread adoption through the Web services approach, which has a suite of simple standards (e.g., XML, WSDL and SOAP) to facilitate interoperability.

These core standards do not provide the rich behavioural detail necessary to describe the role an individual service plays as part of a larger, more complex collaboration. Co-ordination of services is often achieved through the use of workflow technologies. As defined by the Workflow Management Coalition [16], a workflow is the automation of a business process, in whole or part, during which documents, information or tasks are passed from one participant (a resource either human or machine) to another for action, according to a set of procedural rules. Workflow is usually specified from the view of a single participant using centralised *orchestration* or from a global perspective using decentralised *choreography*.

Orchestration languages explicitly describe the interactions between services by identifying messages, branching logic and invocation sequences. Orchestration is described from the view of a single participant, which can be another service. Therefore a central process always acts as a controller to the involved services. The vast majority of workflow tools are based on orchestrating services through a centralised workflow engine: the Business Process Execution Language (BPEL) [29] is the current de-facto standard orchestration language. Workflow tools based on a centralised enactment engine can easily become a performance bottleneck for service-oriented workflows: all data are routed via the workflow engine, these data consume network bandwidth and overwhelm the central engine which becomes a bottleneck to the execution of a workflow. Instead, a solution is desired that permits data output from one service to be forwarded directly to where it is needed at the next service in a workflow.

Choreography on the other hand is more collaborative in nature. A service choreography is a description of the externally observable peer-to-peer interactions that exist between services, therefore choreography does not typically rely on a central co-ordinator. Refer to [3] for a summary of the differences. By adopting a choreography model, the output of a service invocation can be passed directly to where it is required, as input to the next service in the workflow; not through a centralised workflow engine as is the case with orchestration. However, although optimal in terms of data transfer, in practice, the design process and execution infrastructure

- Adam Barker is with the University of St Andrews, UK.  
E-mail: adam@cs.st-andrews.ac.uk
- Jon B. Weissman is with the University of Minnesota, Twin Cities, USA.  
E-mail: jon@cs.umn.edu
- Jano I. van Hemert is with Optos, UK.  
E-mail: jvanhemert@optos.com

for service choreography models are inherently much more complex than orchestration. Decentralised control brings a new set of problems, which are the result of message passing between asynchronous distributed and concurrent processes. Furthermore, current choreography techniques are invasive, in that each individual service needs to be re-engineered in order to take part in a choreography. There are relatively few decentralised choreography languages and even fewer implementations (discussed further in Section 8); the most prevalent being the Web Services Choreography Description Language (WS-CDL) [20].

This paper presents and evaluates the Circulate architecture, which sits in between pure orchestration (completely centralised) and pure choreography (completely decentralised). This centralised control flow, distributed data flow model maintains the robustness and simplicity of centralised orchestration but facilitates choreography by allowing services to transfer data amongst themselves, without the complications associated with modelling and deploying service choreographies. The Circulate architecture reduces data transfer between services (which don't contain functionality for third-party data transfer), which in turn speeds up the execution time of workflows and removes the bottlenecks associated with centralised orchestration.

## 1.1 Paper Contributions

This paper makes the following core contributions:

- *Hybrid architecture*: Circulate is a hybrid between orchestration and choreography techniques: this model maintains the robustness and simplicity of centralised orchestration but facilitates choreography by allowing Web services to transfer data amongst themselves. Importantly, the Circulate architecture is a general architecture and can therefore be implemented using different technologies and integrated into existing systems. However in this paper we will focus on an implementation based on Web services, which is used as the basis for our Internet-scale evaluation.

- *Internet-scale evaluation*: By avoiding the need to pass large quantities of intermediate data through a centralised server, we demonstrate through Internet-scale experimentation how the Circulate architecture reduces data transfer and therefore speeds up the execution time of a workflow. Our evaluation demonstrates how Circulate scales in terms of data size and node size across a range of common workflow topologies.

- *Less-intrusive solution*: In contrast with current service choreography techniques, Circulate is a less-intrusive solution. Our architecture is decoupled from the services they interact with and can be deployed without disrupting existing infrastructure; this means that services do not have to be altered before execution.

The remainder of this paper is structured as follows: Section 2 introduces Circulate by discussing the architecture, a Web services implementation and providing a

concrete example. Section 3 extracts a set of recurring workflow patterns which will be referred to throughout the remainder of this paper. Section 4 discusses the experimental set up used as the basis for the performance analysis across Local Area Network (LAN) and Internet-scale network configurations; a cross product of workflow pattern, node size and network configuration.

Section 5 presents the results from our LAN configuration, covering the remote LAN case, where all services are deployed within a LAN but the workflow engine is remote (common Cloud configuration) and the local LAN case where both the services and engine are deployed on the same LAN. Section 6 presents the results from our Internet-scale configurations, executed over the PlanetLab network, these experiments are broken down into national (all nodes in the same country), continental (all nodes in the same continent) and world wide configurations. In Section 7 the Circulate architecture is applied to an end-to-end application, Montage [17], a benchmark in the High Performance Computing community. Section 8 presents related work covering: decentralised choreography languages, data flow optimisation techniques and Grid frameworks which contain third-party data transfers. Last, Section 9 presents conclusions and future work.

## 2 THE CIRCULATE ARCHITECTURE

This Section describes the Circulate architecture's hybrid model, Web services implementation, API, a corresponding example of use.

### 2.1 Circulate Actors

The Circulate architecture is a hybrid between orchestration and choreography techniques. In order to provide Web services with the required functionality proxies are introduced. A proxy is a lightweight middleware that provides a gateway and standard API to Web service invocation. Proxies are less-intrusive than existing choreography techniques as individual services do not have to be reconfigured in order to take part in a workflow.

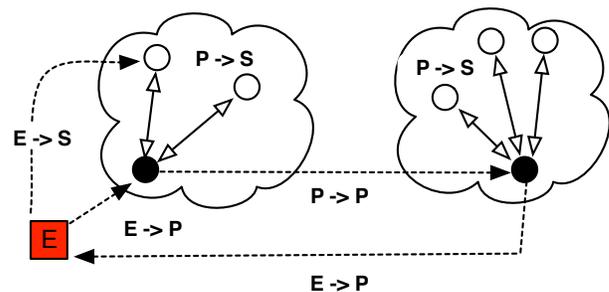


Fig. 1. Actors and interactions of the Circulate architecture. Web services are represented by hollow circles, proxies by solid circles, the workflow engine as a square, dashed lines are WAN hops, solid lines LAN hops.

Proxies are controlled through a centralised workflow engine, running an arbitrary workflow language; allowing standards-based approaches and tooling to be utilised. Proxies exchange references to data with the workflow engine and pass actual data directly to where they are required for the next service invocation. This allows the workflow engine to monitor the progress and make changes to the execution of a workflow.

To utilise a proxy, a Web service must first be registered. The Circulate actors and interactions are illustrated by Figure 1 and described below:

- *Engine to proxy (E→P) interaction*: In the standard orchestration model, the workflow engine interacts directly with all Web services, which for the remainder of the paper we will denote the engine to Web service interaction (E→S). Using the Circulate architecture, the workflow engine remains the centralised orchestrator for the workflow, however the task of service invocation is delegated to a proxy (E→P).

- *Proxy-Web service (P→S) interaction*: Proxies neither create Web service requests, nor do they utilise their responses. Proxies invoke services on behalf of a workflow engine, store the result (if required) and return a reference to the workflow engine.

- *Proxy to proxy (P→P) interaction*: Proxies invoke Web services on behalf of a workflow engine, instead of sending the results of a service invocation back to the workflow engine, they can be stored (if required) within the proxy. In order for a workflow to progress, i.e., the output of a service invocation is needed as input to another service invocation, proxies pass data amongst themselves, moving it closer to the source of the next Web service invocation.

Proxies are ideally installed as “near” as possible to enrolled Web services; by near we mean in terms of network distance, so that the communication overhead between a proxy and a Web service is minimised. Depending on the preference of an administrator, a proxy can be responsible for one Web service, 1:1 or multiple Web services, 1:N. Although it is ideal to place a proxy as closely as possible to an enrolled service (e.g., within the same network domain) it may not always be possible due to the network policy of a particular organisation. Performance benefit can still be accrued simply by harnessing the connectivity of proxies scattered across a network; this is demonstrated throughout our performance analysis across real networks, discussed further in Section 5, 6 and 7.

## 2.2 Web Services Implementation

WS-Circulate is implemented using a combination of Java and the Apache Axis Web services toolkit [2]. Proxies are simple to install and can be configured remotely, no specialised programming needs to take place in order to exploit their functionality. The only changes required are to the workflow specification itself which invokes methods on the proxy rather

than the services directly. WS-Circulate is multi-threaded and allows several applications to invoke methods concurrently. Results from Web service invocations are stored at a proxy by tagging them with a UUID (Universally Unique Identifier) and writing them to disk. There is an assumption that there is sufficient disk space and that storage is temporary, which the proxy can clean up/delete afterwards. Proxies are made available through a standard WSDL interface. This interface contains the following operations: `invoke`, `stage`, `returnData`, `flushTempData`, `addService`, `removeService`, `listOps`, `listOpParams`, `listOpReturnTypes`, and `listServices`. Full details of the API and implementation of the proxy can be found in a complementary paper [4]. In order to simplify the development and deployment, issues of security have not been taken into account, this is left to future work.

## 2.3 Circulate Example

In order to demonstrate the Circulate architecture and API, consider the following simple scenario (an example of a fan-in pattern, discussed further in Section 3): three sources are queried for data via Web service interfaces, these data are combined and used as input to a final sink service, which processes these data and returns a results set. Using UML Sequence diagram notation, standard orchestration is illustrated by Figure 2 and Circulate by Figure 3. Thicker arrows represent data movement, data sizes are arbitrary and used for illustrative purposes only.

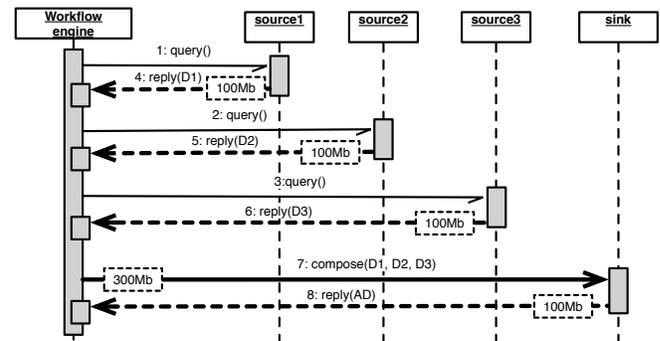


Fig. 2. UML Sequence diagram—orchestrated fan-in.

With reference to Figure 3, the first step in the workflow pattern involves making an invocation to the three source Web services `source1`–`source3`. Instead of contacting the service directly, a call is made to a proxy (`source-proxy`) which has been installed on the same server as the `source1` Web service. The proxy invokes the query operation on the `source1` service, the output is passed back to the proxy, tagged with a UUID (for reference later, e.g., retrieval, deletion etc.) and stored.

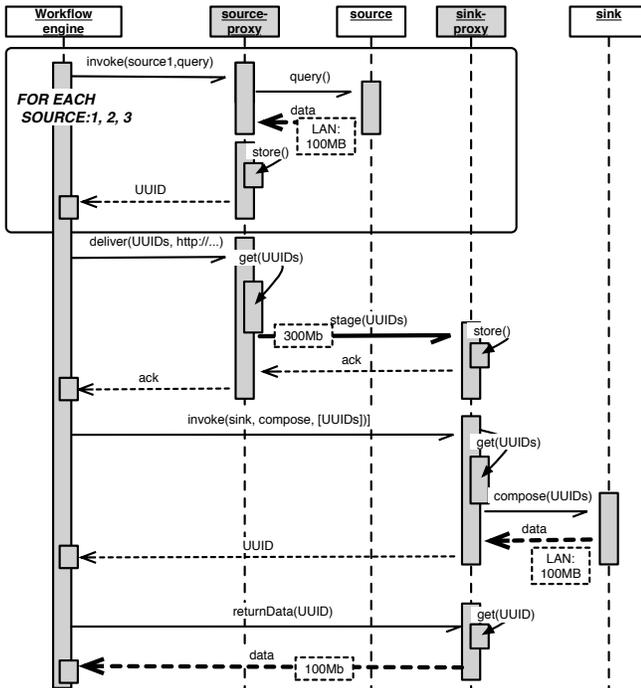


Fig. 3. UML Sequence diagram—Circulate fan-in.

The UUID (not actual data) is returned to the workflow engine. This process is repeated (either serially or in parallel) for `source2` and `source3` which could be served through the same proxy or an independent proxy.

The output from the Web service invocations are needed as input to the next service in the workflow, in this case the `sink` Web service. The workflow engine invokes the `deliver` operation on the `source-proxy` passing in the three UUID references along with the WSDL address of the `sink-proxy`. The `source-proxy` retrieves the stored data and transfers it across the network by invoking a `stage` operation on `sink-proxy`. Data are then stored at `sink-proxy`, if successful an acknowledgement message is sent back to `source-proxy` which is returned to the workflow engine.

The final stage in the workflow pattern requires using the output from the first three services as input to the `sink` Web service. The workflow engine passes the name of the service (`sink`) and operation (`compose`) to invoke and the UUID references, which are required as input. The proxy then moves the data across the local network and invokes the `compose` operation on the `sink` service. The output is again stored locally on the proxy and a UUID reference generated and passed back to the workflow engine. The workflow engine can then retrieve actual data from the proxy when necessary using the `returnData` operation.

### 3 WORKFLOW PATTERNS

As with software design patterns, workflow patterns refer to recurrent problems and proven solutions in the development of workflow applications. There is a large body of workflow patterns research detailing a comprehensive set of patterns from both a control flow and data flow perspective, the most prevalent being the work by van der Aalst and Hofstede et al. [1] and the Service Interaction Patterns set by Barros et al. [6], a collection of thirteen recurring patterns derived from insights into business-to-business transaction processing.

Workflows in the scientific community are commonly modelled as Directed Acyclic Graphs (DAGs), formed from a collection of vertices (units of computation) and directed edges. DAGs present a *dataflow view*, here data are the primary concern, workflows are constructed from data processing (vertices) and data transport (edges). DAGs may be used to model processes in which information flows in a consistent direction through a network of processors. The Genome Analysis and Database Update system (GADU) [26], the Southern California Earthquake Centre (SCEC) [13] CyberShake project, and the Laser Interferometer Gravitational-Wave Observatory (LIGO) [28] are all examples of High Performance Computing applications composed using DAGs.

OMII-Taverna [25] is an example of a popular tool used in the life sciences community in which workflows are represented as DAGs and executed using the service-oriented paradigm.

This paper focuses purely on optimising service-oriented workflows, where services are relatively simple and are not equipped to handle third-party transfers. For the remainder of this paper we take inspiration (i.e., patterns, input-output data relationships, scenarios) from DAG-based workflows but stress that we are focused purely on optimising service-oriented workflows.

#### 3.1 Patterns

From a DAG one can extract a number of isolated workflow patterns (sequence, fan-in and fan-out) which will be used as the basis for performance analysis throughout the remainder of this paper. It is important to note these patterns can be considered primitive or isolated patterns, many isolated patterns can in combination, form a macro-pattern, e.g., a fan-in followed by a fan-out. This situation will be addressed when we discuss the Montage application in Section 7.

DAG-based workflows are not the only possible representation for workflows, however, they are used as the basis for evaluation in this paper as they are commonly used to represent scientific workflows.

$$TC_{seq,nc} = \sum_{S_i=1}^n \left( \overleftrightarrow{C}_{E,S_i} \right) \quad (1)$$

$$TC_{seq,circ} = \overrightarrow{C}_{E,P(S_1)} + \sum_{S_i=1}^n \left( \overleftrightarrow{C}_{P(S_i),S_i} \right) + \sum_{S_i=1}^{n-1} \left( \overrightarrow{C}_{P(S_i),P(S_{i+1})} \right) + \overrightarrow{C}_{P(S_n),E} \quad (2)$$

$$TC_{fi,nc} = \sum_{S_i=1}^n \left( \overrightarrow{C}_{S_i,E} \right) + \overleftrightarrow{C}_{E,S_{FI}} \quad (3)$$

$$TC_{fi,circ} = \sum_{S_i=1}^n \left( \overleftrightarrow{C}_{P(S_i),S_i} \right) + \sum_{S_i=1}^n \left( \overrightarrow{C}_{P(S_i),P(S_{FI})} \right) + \overleftrightarrow{C}_{P(S_{FI}),S_{FI}} + \overrightarrow{C}_{P(S_{FI}),E} \quad (4)$$

$$TC_{fo,nc} = \overrightarrow{C}_{S_{FO},E} + \sum_{S_i=1}^n \left( \overleftrightarrow{C}_{E,S_i} \right) \quad (5)$$

$$TC_{fo,circ} = \overleftrightarrow{C}_{P(S_{FO}),S_{FO}} + \sum_{S_i=1}^n \left( \overrightarrow{C}_{P(S_{FO}),P(S_i)} \right) + \sum_{S_i=1}^n \left( \overleftrightarrow{C}_{P(S_i),S_i} \right) + \sum_{S_i=1}^n \left( \overrightarrow{C}_{P(S_i),E} \right) \quad (6)$$

Fig. 4. Equations modelling the standard orchestration and Circulate case for sequence (1 and 2), fan-in (3 and 4) and fan-out (5 and 6) patterns. Control flow is omitted.

Before we discuss each pattern we introduce the following mathematical notation which is used to provide a concrete representation of each pattern:

- $P_j$  : proxy.
- $S_i$  : service.
- $P(S_i)$  : proxy for service  $S_i$ .
- $\overleftrightarrow{C}_{i,j}$  : round-trip communication cost between entities  $i$  and  $j$ .
- $\overrightarrow{C}_{i,j}$  : one-way communication cost between entities  $i$  and  $j$ .

•  $n$  represents the number of services.  $S_{FI}$  is the fan-in service, where all sources are sent,  $S_{FO}$  is the fan-out service, where source data are extracted.

- $E$  : workflow engine.
- $TC_{pat,nc||circ}$  : total communication cost of executing a workflow pattern (seq, fo, fi - sequence, fan-out, or fan-in, respectively) in Circulate (circ) or non-Circulate (nc), i.e., standard orchestration architectures.

A mathematical representation of each pattern is provided in Figure 4 and illustrated further by Figure 5. We have chosen this method in order to identify generic trends of recurring patterns, which are both commonplace and reflect a realistic representation of how workflows are represented and executed. Complementing the pattern-based evaluation we also execute an end-to-end application, discussed in Section 7. With reference to Figure 4 and Figure 5 each pattern will now be explained in detail:

- *Sequential pattern*: This pattern involves the chaining of services together, where the output of one service invocation is used directly as input to another, i.e., serially. Data are sent from the workflow engine to the

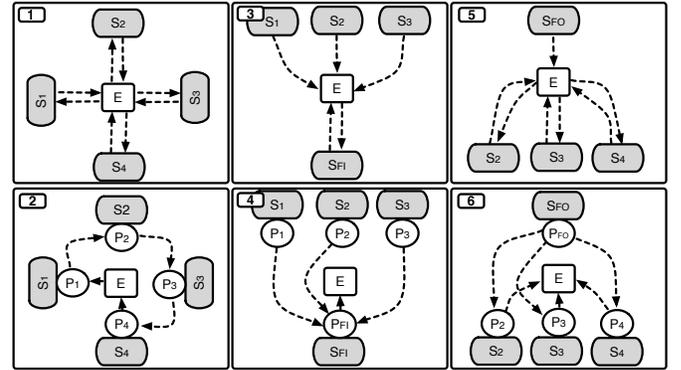


Fig. 5. Experiment setup used for the analytical model and the pattern-based performance analysis. Data flow in the sequential (first column), fan-in (second column) and fan-out (third column) patterns for the centralised architecture using standard services (1,3,5) and the Circulate architecture (2,4,6). Control flow messages are omitted. This example shows 4 services, each proxy is deployed on the same server as the service it is invoking.

first service in the chain and returned from the final service in the chain to the workflow engine.

For the standard orchestration model (equation 1, phase 1 of Figure 5)  $TC_{seq,nc}$  involves round-trip communications ( $\overleftrightarrow{C}$ ) between services  $S_1 \dots S_n$ . For the Circulate architecture (equation 2, phase 2 of Figure 5)  $TC_{seq,circ}$  is calculated as follows: initial data are sent one-way ( $\overrightarrow{C}$ ) from the workflow engine  $E$  to the first proxy  $P(S_1)$ , then round-trip ( $\overleftrightarrow{C}$ ) between all proxies and services  $P(S_1), S_1 \dots P(S_n), S_n$ , one-way ( $\overrightarrow{C}$ ) be-

tween  $n - 1$  proxies, and finally one-way ( $\vec{C}$ ) between the final proxy  $P(S_n)$  and the workflow engine  $E$ .

- *Fan-in pattern*: The fan-in pattern explores what happens when data are gathered from multiple distributed sources, concatenated and sent to a service acting as a sink. Multiple services are invoked with a control flow (no data are sent) message asynchronously, in parallel, data are returned from each service. Once data has been received from all source services it is concatenated and sent to the sink service as input, which in turn returns final output data to the workflow engine.

For the standard orchestration model (equation 3, phase 3 of Figure 5)  $TC_{fi,nc}$  involves one-way ( $\vec{C}$ ) communications between all sources of data  $S_1 \dots S_n$  and the workflow engine  $E$ , all source data are then sent round-trip ( $\overleftrightarrow{C}$ ) between the workflow engine and the final sink service  $S_{FI}$ .

For the circulate architecture (equation 4, phase 4 of Figure 5)  $TC_{fi,circ}$  involves round-trip ( $\overleftrightarrow{C}$ ) communications between all source proxies and all source services  $P(S_1) \dots P(S_n)$ , one-way communication ( $\vec{C}$ ) between all source proxies  $P(S_1) \dots P(S_n)$  and the sink proxy  $P(S_{FI})$ , once received by the sink proxy round-trip communication ( $\overleftrightarrow{C}$ ) takes place between the sink proxy  $P(S_{FI})$  and the sink service  $S_{FI}$ , finally data are sent one-way ( $\vec{C}$ ) back to the workflow engine  $E$ .

- *Fan-out pattern*: This pattern is the reverse of the fan-in pattern, here the output from a single source is sent to multiple sinks. An initial source service is invoked with a control flow message (again no actual data are sent), the service returns data as output. These data are sent, asynchronously in parallel to multiple services as input, each service returns final data to the workflow engine.

For the standard orchestration model (equation 5, phase 5 of Figure 5)  $TC_{fo,nc}$  involves one-way communication ( $\vec{C}$ ) between the source service  $S_{FO}$  and the workflow engine  $E$ , followed by round-trip communication ( $\overleftrightarrow{C}$ ) between the workflow engine  $E$  and all sink services  $S_1 \dots S_n$ .

For the Circulate architecture (equation 6, phase 6 of Figure 5)  $TC_{fo,circ}$  involves round-trip communication ( $\overleftrightarrow{C}$ ) between the source proxy  $P(S_{FO})$  and the source service  $S_{FO}$ , one-way communication ( $\vec{C}$ ) between the source proxy  $P(S_{FO})$  and all sink proxies  $P(S_1) \dots P(S_n)$ , round-trip communication ( $\overleftrightarrow{C}$ ) between all sink proxies  $P(S_1) \dots P(S_n)$  and all sink services  $S_1 \dots S_n$ , finally one-way communication ( $\vec{C}$ ) between all sink proxies  $P(S_1) \dots P(S_n)$  and the workflow engine  $E$ .

## 4 EXPERIMENTAL SETUP

A number of performance analysis experiments have been devised in order to observe and analyse the behaviour of various workflow patterns and variables to compare the Circulate architecture with that of a standard centralised orchestration model. The following sections describe the structure of the resulting performance

analysis experiments which are based the workflow patterns described in Section 3.

### 4.1 Node Size and Network Configurations

In our experimental set up, data flows with no data transformations (i.e., the output of one method invocation is the input to another). We use non-blocking asynchronous communication between the Web services, although data forwarding between proxies occurs after all Web services have finished execution.

For each of the 3 workflow patterns: sequence, fan-in and fan-out, the time taken for the pattern to complete using centralised orchestration and using the Circulate architecture is recorded (in milliseconds) as the size of the input data (in Megabytes) is increased. This basic set up is then run incrementally over a number of workflow node sizes: 4, 8 and 16 nodes in order to explore the affects of scalability. Each node size (i.e., 4, 8 or 16) is then run over different network configurations to explore how network configuration affects the performance of a workflow.

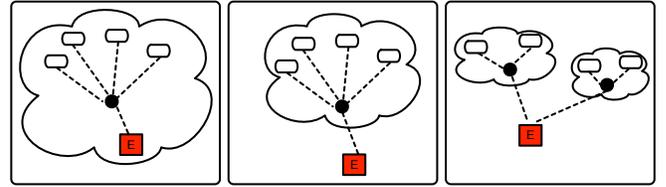


Fig. 6. Configurations from left to right: local LAN, remote LAN, Internet-scale. Services, proxies and the workflow engine are always deployed on separate machines.

We have selected the following network configurations (illustrated by Figure 6) which mimic common scenarios when composing sets of geographically distributed services:

- *LAN experiments*: Moving the Web services to a relatively uniform network topology in terms of speed, e.g., a LAN, allows for a simplified analysis of the two models. In a LAN, it is expected that the cost of the communication links  $E \rightarrow S$ ,  $E \rightarrow P$ ,  $P \rightarrow S$  and  $P \rightarrow P$  have little variance in terms of bandwidth. In this case the performance benefit with respect to the different workflow patterns may be exposed more readily. We explore this with a LAN configuration with a local workflow engine (same LAN) and a LAN configuration with a remote workflow engine (connected through a WAN).

- *Internet-scale experiments*: PlanetLab [9] is a global research network of distributed servers that supports the development of new network services. We have heavily utilised this framework in order to evaluate the performance of the Circulate architecture across Internet-scale networks. We have subdivided the pattern-based PlanetLab experiments into 3 configurations which mimic the typical geographical distribution patterns found in

workflow applications: ‘National’ (all nodes in the same country), ‘Continental’ (nodes in different countries on the same continent), and ‘World-wide’ (nodes spread throughout the globe). Finally, to demonstrate end-to-end performance, an end-to-end application is executed across the PlanetLab framework.

## 4.2 Experiment Cross Product

The configuration of our experiments mirror that of a typical workflow scenario, where collections of physically distributed services need to be composed into a higher level application, scaling from LAN to Internet-scale network configurations. In order to create a uniform test environment, proxies invoke Web services with one operation which inputs and outputs Java byte arrays transferred using SOAP. The cross product of each workflow pattern, node size and network configuration has been executed 100 times (combined standard deviation can be seen on graphs where a mean performance benefit is reported) over a set of distributed Linux machines running the WS-Circulate architecture discussed in Section 2.2. The experiments and graphs (throughout this paper) can be summarised as follows:

- 1)  $x$ -axis displays the size of the initial input file in Megabytes: for the sequence pattern this displays the size of data *sent* to the first service in the workflow, for fan-in and fan-out this represents the size of data *returned* by the first service.
- 2) For the LAN experiments the size of the initial input file ranges from 2MB to 96MB with a total of 12 data points collected. For the Internet-scale experiments the size of the initial input file ranges from 2MB to 64MB with a total of 10 data points collected. Each experiment is executed 100 times.
- 3) The  $y$ -axis displays the mean performance benefit. The mean performance benefit is calculated by dividing the mean time taken to execute the workflow using standard orchestration (i.e., non-proxy, fully centralised) and dividing it by the mean time taken to execute the workflow using the Circulate architecture. For example, a result of 2.0 means that the Circulate architecture executes a workflow twice as fast as standard orchestration.
- 4) Graphs that represent a performance benefit are shown together with a combined standard deviation of the two populations: orchestration and Circulate, represented by equation 7.

$$\sqrt{\left(\frac{stdev-non-circ}{mean-non-circ}\right)^2 + \left(\frac{stdev-circ}{mean-circ}\right)^2} \quad (7)$$

## 5 LAN CONFIGURATION

A pool of computers from the University of Edinburgh Distributed Informatics Computing Environment (DICE)<sup>1</sup> LAN were selected as workflow nodes for these configurations. These machines are all located in the same building and are connected to the network via a 100 Mbit network connection. All these machines share the same hardware/operating environment: Intel Core 2 Duo with 2GB RAM running Ubuntu version 7.04. In addition to these matched machines, similar servers (separate machines) were chosen to act as the workflow engine, allowing us to explore the remote orchestration case and the local orchestration case.

### 5.1 LAN - Remote Orchestration

By moving the workflow engine outside of the LAN configuration, we can explore workflow engine to Web service (E→S) communication. This reflects a very common real world scenario where a particular organisation provides all of the Web services used to compose a workflow, but these services are being orchestrated remotely. In particular this is applicable to Cloud Computing scenarios (for example Amazon Elastic Compute Clouds<sup>2</sup>) where services are typically hosted on groups of co-located machines and where the end user requiring these services is remote.

Figure 7(a) displays the mean performance benefit across each of the workflow patterns, on a LAN with 4 nodes, where all Web services are mapped to a single proxy and the workflow engine is remote from both proxies and services.

As one can observe from Figure 7(a), for all patterns the Circulate architecture outperforms the standard orchestration model, i.e., the mean performance benefit is larger than 1 even at low data sizes. The fan-in pattern shows the greatest benefit.

To justify, as the workflow engine lies outside the LAN where the Web services and proxies are deployed, the E→S and E→P links can be considered “expensive” WAN hops. As the proxy is deployed on the same LAN as the Web services it is invoking the P→S link is a less expensive LAN hop. Within a standard orchestration model, all data pass through a centralised workflow engine, hence all data are transferred over the expensive E→S link. Using the Circulate architecture, most of the processing for each of the patterns takes place over the less expensive P→S link; intermediate data are housed within the proxy. Data are only sent over the expensive E→P link at the start of the sequence pattern and the end of all patterns.

One can make the following observations from the LAN-scale experiments:

- *Patterns:* With reference to Figure 7(a), the worst performing pattern in this scenario was the sequence

1. <http://www.dice.inf.ed.ac.uk> [24/02/11]

2. <http://aws.amazon.com/ec2> [24/02/11]

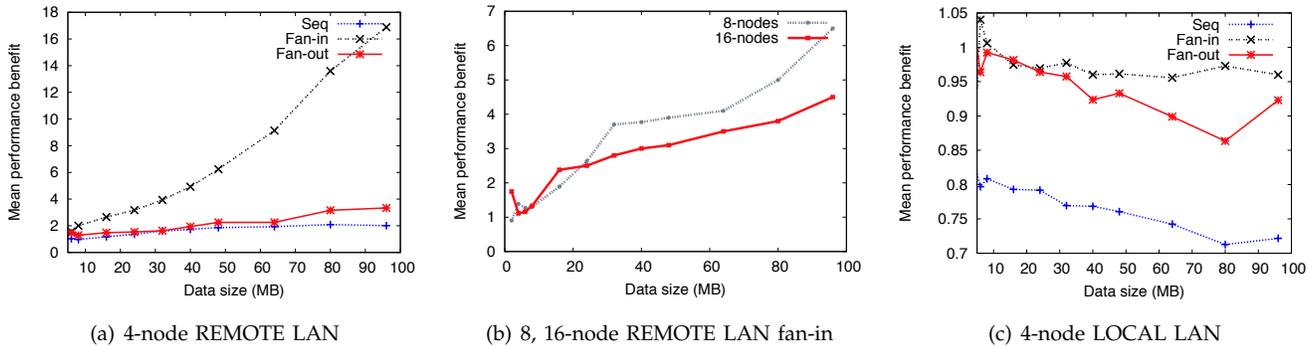


Fig. 7. Remote LAN configurations

pattern, however even this pattern demonstrates the significant advantage of using the Circulate architecture. At 2MB, the mean performance benefit is just 1.03, at 96MB this benefit increases to 2.0. Looking at the best performing pattern, fan-in, at 2MB the mean performance benefit is 2.0, at 96MB the mean performance benefit has risen to 16.8.

- *Node size*: In order to explore scalability, the same experiment was run over 8-node and 16-node configurations. One proxy is assigned to each group of four Web services, resulting in 2 proxies for the 8-node experiment and 4 proxies for the 16-node experiment. As per the previous setup, the workflow engine is remote from all services and proxies. Figure 7(b) illustrates the mean performance benefit of each of the node configurations when compared to the centralised orchestration model for the fan-in pattern. We have selected the fan-in pattern as it demonstrates the most consistent improvement in performance.

The Circulate architecture always outperforms the centralised orchestration model for each of the node configurations. To justify, as the number of proxies increases, so does the  $P \rightarrow P$  communication, which adds a minimal additional cost to the execution of a workflow. However, this minimal overhead only adds an additional cost in comparison with using a configuration where all services share the same proxy. Our scaling experiment demonstrates that Circulate outperforms traditional orchestration for all node sizes in the remote LAN case.

- *Data size*: As the size of data involved in each of the patterns increases, the cost of processing the expensive WAN hops also increases. As the Circulate architecture reduces these more expensive WAN hops, the benefit of utilising the architecture increases in proportion to the size of data involved in a workflow. To quantify (with reference to Figure 7(b)), at 16MB the 8-node fan-in pattern’s mean performance benefit was 1.9, at 96MB this increased to 6.4. At 16MB the 16-node fan-in pattern’s mean performance benefit was 2.2, at 96MB this increased to 4.4.

## 5.2 LAN - Local Orchestration

In order to explore the limits of our approach, the workflow engine is deployed on a computer that is also connected to the proxies and Web services via a network switch. This is a suitable experimental setup to test the assumption of communication link equality but does not necessarily reflect common deployed patterns of Web services, normally the workflow engine is remote.

Figure 7(c) displays the mean performance benefit as the data size increases across each of the 3 workflow patterns on a LAN running a local workflow engine, with 4-nodes, where 4 Web services each share a proxy.

In a LAN environment we make the assumption that the cost of communication is relatively uniform, therefore the cost of  $E \rightarrow S$ ,  $E \rightarrow P$ ,  $P \rightarrow S$  and  $P \rightarrow P$  can be considered approximately equal. The Circulate architecture introduces extra communication links (between  $P \rightarrow S$ ) across this uniform network topology, which in turn degrades the execution time of a workflow when compared to standard orchestration. With reference to Figure 7(c) the fan-in pattern demonstrates similar performance to the standard orchestration model, while the standard orchestration model outperforms the fan-out and sequence patterns.

## 6 INTERNET-SCALE CONFIGURATIONS

Moving the experiments to the PlanetLab configurations allows the Circulate architecture to be evaluated over Internet-scale networks. The PlanetLab environment configurations are based on the geographical location of the nodes, which in turn are used as an indicator of communication link cost. By grouping the nodes, certain realistic scenarios can be constructed for the experiments. For example, by using a group of nodes all located in France, one can execute a workflow simulating the interactions between collaborating French universities; such scenarios are common place in large-scale workflows. The PlanetLab configurations all use a remote workflow engine, this is a common feature of scientific workflows, as one can imagine a scientist in a remote location orchestrating resources from a number of collaborating institutions.

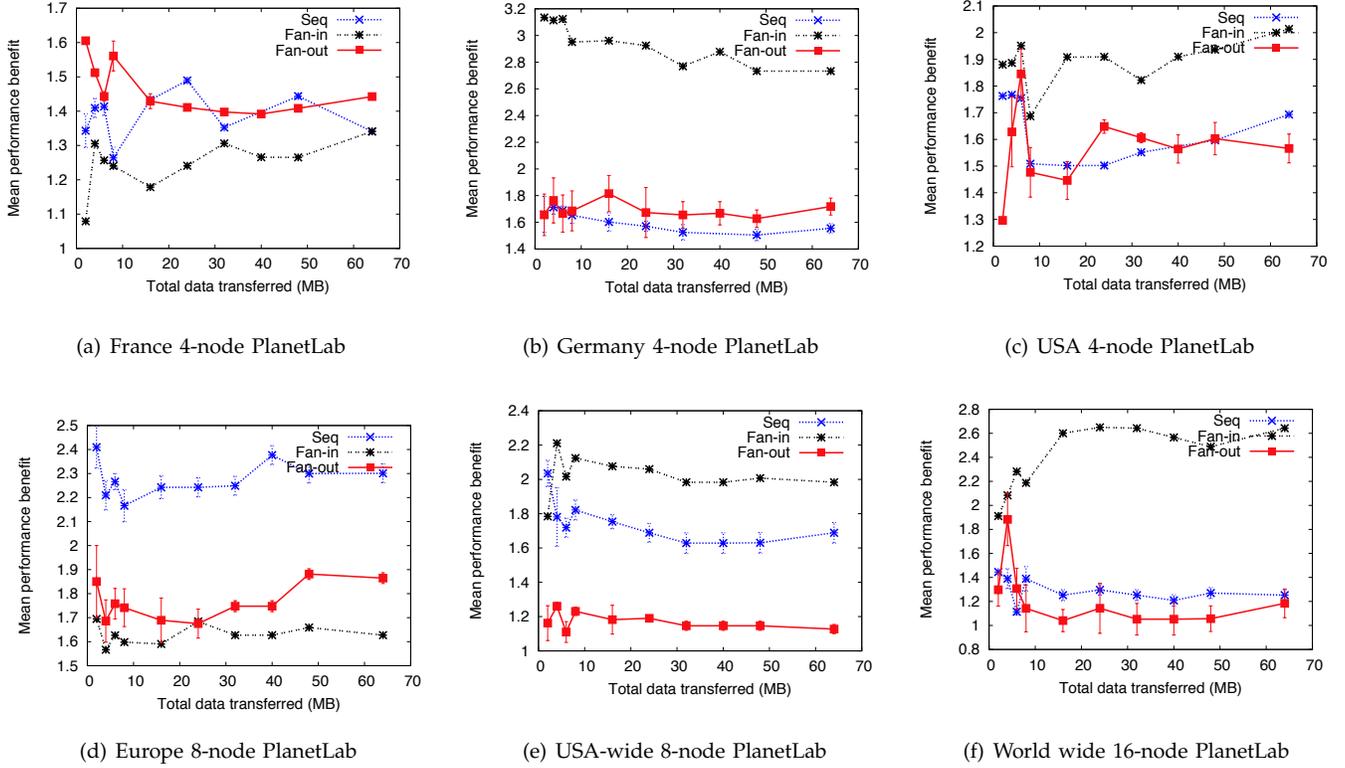


Fig. 8. National (4-node), continental (8-node) and global (16-node) PlanetLab configurations

Due to the variable nature of PlanetLab, node selection was a complex matter. We ran extensive tests to locate groups of nodes that we could reliably access throughout the duration of our experiments. After these preliminary tests nodes from France, Germany and the USA were selected, each nodes are shown in Table 1. Each workflow pattern was then executed over nodes from Table 1, mashed up using different geographical network configurations. The breakdown of each pattern, network configuration and mean performance benefit is displayed in Table 2.

TABLE 1  
PlanetLab node selection.

France	Germany	USA
inisa.fr	uni-konstanz.de	brown.edu
inria.fr	uni-goettingen.de	mit.edu
utt.fr	uni-paderborn.de	poly.edu
	fraunhofer.de	umd.edu
	tu-darmstadt.de	byu.edu
		postel.org
		iit-tech.net

- *National-4-node*: Within the national configuration, 4 nodes were selected to deploy the Web services, a further node acted as a shared proxy and the workflow engine was deployed on a final node that is a separate node from the proxy and services. The mean performance benefit for the France configuration is displayed in Fig-

TABLE 2  
National (4-node), continental (8-node) and world wide (16-node) PlanetLab mean performance benefits.

Configuration	Best	Worst	Mean
<b>4-node – France</b>			
Seq	1.49	1.27	1.39
Fan-in	1.30	1.08	1.27
Fan-out	1.61	1.4	1.43
Overall			1.36
<b>4-node – Germany</b>			
Seq	1.72	1.50	1.58
Fan-in	3.13	2.72	2.94
Fan-out	1.82	1.64	1.69
Overall			2.07
<b>4-node – USA</b>			
Seq	1.76	1.50	1.61
Fan-in	1.95	1.69	1.85
Fan-out	1.85	1.30	1.61
Overall			1.69
<b>8-node – Europe</b>			
Seq	2.42	2.18	2.27
Fan-in	1.69	1.56	1.61
Fan-out	1.87	1.67	1.72
Overall			1.87
<b>8-node – USA</b>			
Seq	2.04	1.64	1.75
Fan-in	2.22	1.79	2.04
Fan-out	1.27	1.11	1.18
Overall			1.66
<b>16-node – World wide</b>			
Seq	1.44	1.12	1.29
Fan-in	2.63	1.92	2.32
Fan-out	1.89	1.04	1.22
Overall			1.61

ure 8(a), the Germany configuration in Figure 8(b) and the USA configuration in Figure 8(c).

- *Continental-8-node*: In order to separate the nodes further, each workflow pattern was executed across 2 continental configurations. The first was a European set, which contained nodes from both France and Germany. The second utilised nodes across the USA. 8 Web services were deployed across 8 nodes, a further 2 nodes acted as proxies for groups of 4 Web services, a final node acted as a workflow engine that is a separate node from any of the proxies and services. The mean performance benefit for the European configuration is displayed in Figure 8(d) and USA-wide configuration in Figure 8(e).

- *World wide-16-node*: The final PlanetLab experiment executed each pattern across all available nodes in the previous configurations. 16 Web services were deployed across 16 nodes, 4 proxies were deployed across separate nodes, which manage 4 Web services each, a final node acted as the workflow engine that is a separate node from any of the proxies and services. The mean performance benefit for the world wide configuration is displayed in Figure 8(f).

One can make the following observations from the Internet-scale experiments:

The Circulate architecture always outperformed the centralised orchestration model. The greatest overall improvements in performance were seen on the 4-node Germany configuration (2.07 overall mean performance benefit), 8-node Europe configuration (1.87 overall mean performance benefit), and 4-node USA configuration (1.69 overall mean performance benefit).

- *Patterns*: If we calculate the mean across all Internet-scale network configurations (i.e., the 4, 8 and 16-node) from Table 2 the following trends can be observed: the sequence pattern demonstrated a performance benefit of 1.65, fan-in a performance benefit of 2.01 and fan-out a performance benefit of 1.48.

- *Data size*: As one can observe from Figures 8(a) to 8(f), the Internet-scale configurations were more variable than the LAN experiments. However the general trend was that as the data size increased the benefit either improved or remained relatively constant.

- *Node size*: Obtaining general trends on the PlanetLab results was not as straight forward as the LAN configurations. Unlike the relatively uniform LAN environment, the performance across PlanetLab is heavily dependent on the quality of links utilised and current load of the network (observed by the error bars in Figure 8). As we have discovered PlanetLab links vary radically in quality. For example, the overall improvement in the 4-node German configuration (2.07 mean performance benefit) is higher than in the 4-node France configuration (1.36 mean performance benefit). Analysing individual runs it was found that the E→S and E→P communication cost was higher in the German nodes. As the Circulate architecture reduces these expensive links, the mean performance benefit across pattern was lower.

For standard orchestration the quality of the E→S links

is the overall factor affecting performance, for Circulate it is the quality of the E→P and P→P links. Therefore one has to place proxies carefully in order to ensure that the P→P links are faster than the E→S links, these optimisation and proxy placement strategies are left to future work.

Most promising however was that in all cases, when all workflow patterns were taken into account, the Circulate architecture always speedup the execution time of a workflow. This is important as it demonstrates that even if one pattern is misbehaving, a workflow, which will be composed of an arbitrary number of patterns will still experience performance improvements.

## 7 END-TO-END APPLICATION: MONTAGE

Although the focus of our paper has primarily been on pattern-based performance analysis, it is important to demonstrate the Circulate architecture on an end-to-end application. Figure 9 illustrates the Montage workflow, a benchmark in the High Performance Computing community and representative of a class of large-scale, data-intensive workflows. Montage constructs custom “science-grade” astronomical image mosaics from a set of input image samples. It illustrates several features of data-intensive scientific workflows and it can result in large data flow requirements, intermediate data can be 3 times the size of input data.

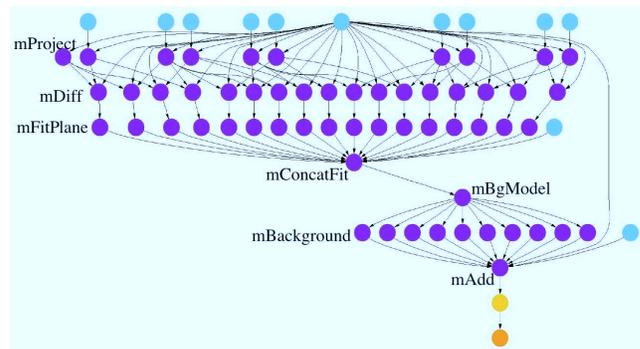


Fig. 9. Montage use-case scenario.

The Montage workflow is represented as a DAG, each component of the DAG (along with corresponding input-output data ratios) is explained as follows:

- 1) *mProject*: re-projects an image to the coordinate system defined in a header file. (output = input)
- 2) *mDiff/mFitPlane*: 3 inputs (1 header and 2 images) *fan-in* to the *mDiff* function, which finds the difference between the two images, the output is then passed through the *mFitPlane* function (usually executed on the same machine), which fits a plane to the difference image. Output = 15–20% of a typical image for each image triplet.

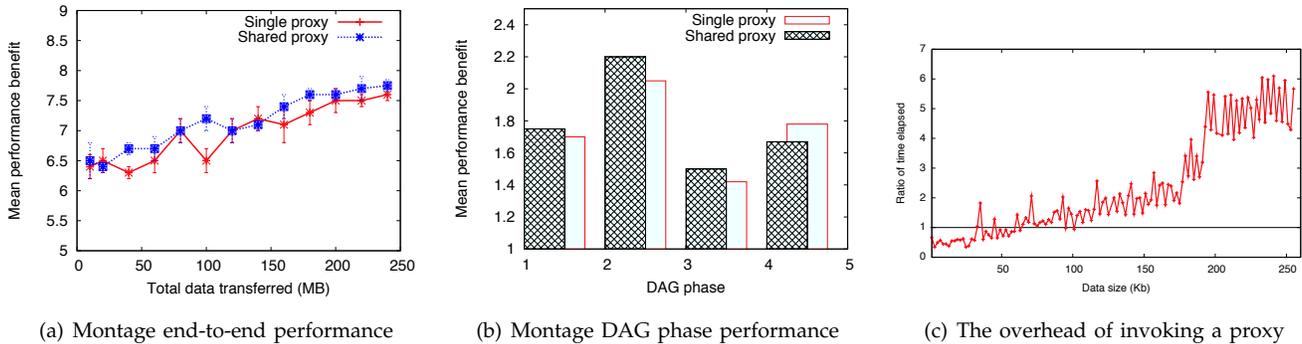


Fig. 10. Montage end-to-end performance (a), Montage individual phases performance (b) and proxy overhead (c).

- 3) mConcatFit: a simple concatenation of the plane fit parameters from multiple mDiff/mFitPlane jobs into a single file. *fan-in* pattern with 18 inputs (from different resources), which are passed through the mConcatFit function.
- 4) mBgModel: models the sky background using the plane fit parameters from mDiff/mFitPlane and computes planar corrections for the input images that will rectify the background across the entire mosaic. mBgModel is a *fan-out* pattern, where the output is distributed to 10 sinks.
- 5) mBackground/mAdd: the mBackground function rectifies the background in a single image, output = input. Data from each mBackground computation are sent to the mAdd function, which co-adds a set of reprojected images to produce a mosaic as specified in a template header file. This forms a *fan-in* pattern with 10 inputs to the mAdd function (output = 70–90% the size of inputs put together).

The Montage application (Web services and data) along with a set of proxies were deployed on PlanetLab nodes spanning the USA. Our deployment maintains the number of services (i.e., fan-ins and fan-outs), data sizes and importantly the input-output relationships of Montage.

Two experiments were performed: first, “single proxy” where each Web service is maintained by its own proxy, and “shared proxy” where groups of 4 Web services were maintained by a single proxy. Proxies and services were scattered across PlanetLab nodes spanning the USA, proxies were always deployed on a separate machine within the same domain as the Web service it is invoking and the workflow engine was always remote from both proxy and Web service. This deployment was executed 50 times across the PlanetLab framework. As with previous experiments the  $x$ -axis represents the size of the input file and 95% confidence intervals are provided for every mean performance benefit. A total of 13 data points were collected from 10MB to 240MB. As we are focusing purely on optimising data transfer, processing times in both models have not been taken into account as these remain the same regardless of the number of

services served per proxy.

Figure 10(a) illustrates the end-to-end mean performance benefit of the Montage application using Circulate and standard centralised orchestration. Figure 10(b) illustrates the mean performance benefit (average across 13 data points and 50 runs) per phase of the Montage application and demonstrates how each phase collectively provides an end-to-end benefit. Phase 1 represents mProject to mFitPlane, phase 2 mConcatFit, phase 3 mBgModel, and phase 4 mBackground to mAdd.

The end-to-end Montage application resulted in a mean performance benefit of 6.95 for the single proxy configuration over all 13 data volumes tested.

This benefit increases slightly as the input data size increased: from an average of 6.4 (30 seconds for orchestration, 192 seconds for Circulate) at 10MB to an average of 7.6 (246 seconds for orchestration, 1869.6 for Circulate) at 240MB. The shared proxy configuration resulted in a mean performance benefit of 7.12, again the benefit increased slightly as the data size increased: from 6.5 (24 seconds for orchestration, 156 seconds for Circulate) at 10MB to 7.75 (216 seconds for orchestration, 1674 seconds for Circulate) at 240MB. The shared proxy resulted in a marginal benefit over the single proxy configuration due to reduced data transfer to and from proxies. With reference to the pattern-based performance analysis, we confirm that the benefit of using the Circulate architecture increases when isolated patterns are placed together to form a larger application.

## 7.1 Proxy Overhead

Figure 10(c) displays the average time (as a ratio: non-proxy centralised time divided by Circulate elapsed time) it takes to make a single invocation to a vanilla Web service and obtain the result vs. an invocation to a proxy that invokes the service on the orchestration engines behalf and returns a reference to its data. The workflow engine is remote to both proxy and service. Results under the horizontal line indicate the vanilla approach is optimal, results over the line show a benefit of using the Circulate architecture.

Circulate effectively separates out the control flow and data flow messages from one another; the work-

flow engine invokes proxies with smaller control flow messages and proxies pass larger data flow messages directly to one another. Although workflow languages such as BPEL do not explicitly separate out these two types of messages it is useful to think of data transfer in this way. The proxy overhead experiment demonstrates that Circulate is only suited to workflows where large quantities of intermediate data flow between services in a workflow. From the results we conclude that due to the overhead of the proxy, when dealing with input data sizes of less than  $\sim 100\text{K}$  of data the Circulate architecture offers no performance benefit to Web services. Anything over  $\sim 100\text{K}$  of data the proxy begins to speedup the execution time of the invocation due to the storage of the results within the proxy.

## 8 RELATED WORK

This Section discusses all related work from the literature, spanning pure choreography languages, enhancements to widely used modelling techniques, i.e., BPMN, decentralised orchestration, data flow optimisation architectures and Grid toolkits.

### 8.1 Choreography Languages

For completeness it is important to discuss the current state-of-the-art in choreography techniques. There are an overwhelming number of pure orchestration languages. However, relatively few targeted specifically at choreography, the most prevalent being WS-CDL [5], BPEL4Chor [12] and Let's Dance [30].

There are even fewer complete implementations of choreography languages, this means that choreography techniques are rarely deployed in practice. For example, there are only two documented prototype implementations of the WS-CDL specification. WS-CDL+, an extended specification [18] has been implemented in prototype form, although only one version, version 0.1 has been released. A further partial implementation [14] of the WS-CDL specification is currently in the prototype phase. The other widely known implementation is pi4soa<sup>3</sup>, an Eclipse plugin that provides a graphical editor to compose WS-CDL choreographies and generate from them compliant BPEL. Maestro [11] is an implementation of the Let's Dance language and supports the static analysis of global models, the generation of local models from global ones, and the interactive simulation (not enactment) of both local and global modes. In the BPEL4Chor space, a Web-based editor<sup>4</sup> allows engineers to graphically build choreography models.

### 8.2 Techniques in Data Transfer Optimisation

There are a limited number of research papers that have identified the problem of a centralised approach to service orchestration.

The *Flow-based Infrastructure for Composing Autonomous Services* or FICAS [23] is a distributed data-flow architecture for composing software services. Composition of the services in the FICAS architecture is specified using the Compositional Language for Autonomous Services (CLAS), which is essentially a sequential specification of the relationships among collaborating services. This CLAS program is then translated by the build-time environment into a control sequence that can be executed by the FICAS runtime environment.

FICAS is intrusive to the application code as each application that is to be deployed needs to be wrapped with a FICAS interface. In contrast, our proxy approach is more flexible as the services themselves require no alteration and do not even need to know that they are interacting with a proxy. Furthermore our proxy approach introduces the concept of passing references to data around and deals with modern workflow standards.

*Service Invocation Triggers* [7], or simply Triggers are also a response to the bottleneck problem caused by centralised workflow engines. Triggers collect the required input data before they invoke a service, forwarding the results directly to where these data are required. For this decentralised execution to take place, a workflow must be deconstructed into sequential fragments that contain neither loops nor conditionals and the data dependencies must be encoded within the triggers themselves. This is a rigid and limiting solution and is a barrier to entry for the use of proxy technology. In contrast with our proxy approach, because data references are passed around, nothing in the workflow has to be deconstructed or altered, which means standard orchestration languages such as BPEL can be used to co-ordinate the proxies.

In [24] a similar (pure choreography) approach is also proposed. Authors introduce a methodology for transforming the orchestration logic in BPEL into a set of individual activities that co-ordinate themselves by passing tokens over shared, distributed tuple spaces. The model suitable for execution is called Executable Workow Networks (EWFN), a Petri nets dialect. This approach utilises a pure choreography model which has many additional modelling and enactment problems associated with it, due primarily to the complexity of message passing between distributed, concurrent processes.

*Data caching techniques* in Grid workflows are proposed in [8]. This architecture caches "virtual data" of previous queries, so any overlapping queries and processing do not have to be repeated. The Circulate architecture stores data at a proxy so that it can be transferred directly to the next stage of the workflow, avoiding costly network hops back to the workflow engine. It is then up to the user to clean up stored data afterwards, by using the methods discussed in Section 2.2. The automated data caching techniques proposed could be applied to the Circulate proxies to further enhance performance.

3. <http://sourceforge.net/projects/pi4soa> [24/02/11]

4. <http://www.bpel4chor.org/editor> [24/02/11]

### 8.3 Third-party Data Transfers

This paper focuses primarily on optimising service-oriented workflows, where services are: not equipped to handle third-party transfers, owned and maintained by different organisations, and cannot be altered in anyway prior to enactment. For completeness it is important to discuss engines that support third-party transfers between nodes in task-based workflows.

*Directed Acyclic Graph Manager (DAGMan)* [10] submits jobs represented as a DAG to a Condor pool of resources. DAGMan removes the workflow bottleneck as data are transferred directly between vertices in a DAG however focuses purely on Condor based Grid jobs and does not address the bottleneck problems associated with orchestrating service-oriented workflows.

*Triana* [27] is an open-source problem solving environment. It is designed to define, process, analyse, manage, execute and monitor workflows. Triana can distribute sections of a workflow to remote machines through a connected peer-to-peer network.

*OGSA-DAI* [19] is a middleware product that supports the exposure of data resources on to Grids. This middleware facilitates data streaming between local OGSA-DAI instances. Our architecture could be implemented on this platform to take advantages of its streaming model.

*Grid Services Flow Language (GSFL)* [21] addresses some of the issues discussed in this paper in the context of Grid services, in particular services adopt a peer-to-peer data flow model. However, individual services have to be altered prior to enactment, which is an invasive and custom solution, something that is avoided in the Circulate architecture.

## 9 CONCLUSIONS

As the number of services and the size of data involved in workflows increases, centralised orchestration techniques are reaching the limits of scalability, in standard orchestration: all data passes through a centralised engine resulting in unnecessary data transfer, wasted bandwidth and the engine to become a bottleneck to the execution of a workflow. Decentralised choreography techniques, although optimal in terms of data transfer are far more complex to build due to message passing between distributed, concurrent process and in practice and rarely deployed.

This paper has presented the Circulate architecture; a centralised orchestration model of control with a peer-to-peer choreography model of data transfer. This extra functionality is achieved through the deployment of a lightweight proxy that provides a gateway and standard API to Web service invocation. Importantly, proxies can be deployed without disrupting current services and with minimal changes in the workflows that make use of them. This flexibility allows a gradual change of infrastructures, where one could concentrate first on improving data transfers between services that handle large amounts data.

An open-source Web services implementation, WS-Circulate, served as the platform for our evaluation across LAN and Internet-scale configurations through the PlanetLab network. Through our experimentation we have demonstrated that by reducing data transfer, the Circulate architecture significantly speeds up the execution time of workflows across common workflow patterns and network topologies. Moreover, it scales with data size and common workflow topologies, and consistently outperforms centralised orchestration techniques. The Montage DAG demonstrated that the benefit of using the Circulate architecture increases when isolated patterns are placed together to form a larger application.

Future work includes the following challenges:

- *Architecture evolution*: Although this paper has discussed a Web services-based implementation, Circulate is a general architecture and mappings could be provide to multiple back end technologies, e.g., Condor [22]. This would allow multiple technology sets to be optimised and orchestrated via a standard workow language and workflow engine.

- *Compressing Web service content*: Techniques e.g., gzipped SOAP<sup>5</sup> have been proposed to reduce the quantity of data transferred between Web services. We plan to integrate such techniques into the Circulate architecture.

- *Proxy placement*: This architecture also opens up a rich set of additional optimisations with respect to dynamic proxy deployment i.e., load balancing depending on network traffic.

- *Data transformation*: A common workow task is to change the output data from one service into a slightly different format to use as input into another service, this process is known as Shimming. This will affect the performance of a workow further as it introduces an extra service into the workow chain. If a shim is implemented as a custom transformation Web service this step will simply be included in the workflow specification. Future work includes implementing a shim loader component, where custom shims can be uploaded to the proxy, allowing shims to be performed locally at a proxy, avoiding the introduction of further network hops.

## REFERENCES

- [1] W.M.P van der Aalst, A.H.M. ter Hofstede, B. Kiepuszewski, and A.P. Barros. Workflow patterns. *Distributed and Parallel Databases*, 14:5–51, 2003.
- [2] Apache Axis. <http://ws.apache.org/axis> [24/02/11].
- [3] A. Barker and J. van Hemert. Scientific Workflow: A Survey and Research Directions. In Roman Wyrzykowski and et al., editors, *Seventh International Conference on Parallel Processing and Applied Mathematics, Revised Selected Papers*, volume 4967 of LNCS, pages 746–753. Springer, 2008.
- [4] A. Barker, J. B. Weissman, and J. van Hemert. Orchestrating Data-Centric Workflows. In *The 8th IEEE International Symposium on Cluster Computing and the Grid (CCGrid)*, pages 210–217. IEEE Computer Society, 2008.

5. <http://www.ibm.com/developerworks/webservices/library/ws-sqzsoap.html> [24/02/11]

- [5] A. Barros, M. Dumas, and P. Oaks. A Critical Overview of the Web Services Choreography Description Language (WS-CDL). *BP Trends Newsletter* 3, 2005.
- [6] A. Barros, M. Dumas, and A. ter Hofstede. Service Interaction Patterns. In *Proceedings of the 3rd International Conference on Business Process Management*, pages 302–318. Springer, 2005.
- [7] W. Binder, I. Constantinescu, and B. Faltings. Decentralized Orchestration of Composite Web Services. In *Proceedings of the International Conference on Web Services, ICWS'06*, pages 869–876. IEEE Computer Society, 2006.
- [8] D. Chiu and G. Agrawal. Hierarchical caches for grid workflows. In *CCGRID '09: Proceedings of the 2009 9th IEEE/ACM International Symposium on Cluster Computing and the Grid*, pages 228–235, Washington, DC, USA, 2009. IEEE Computer Society.
- [9] B. Chun, D. Culler, T. Roscoe, A. Bavler, L. Peterson, M. Wawrzoniak, and M. Bowman. Planetlab: an overlay testbed for broad-coverage services. *SIGCOMM Comput. Commun. Rev.*, 33(3):3–12, 2003.
- [10] Condor Team. [www.cs.wisc.edu/condor/dagman](http://www.cs.wisc.edu/condor/dagman) [07/06/10].
- [11] G. Decker, M. Kirov, J. Maria Zaha, and M. Dumas. Maestro for Let's Dance: An Environment for Modeling Service Interactions. In *Demonstration Session of the 4th International Conference on Business Process Management (BPM)*, 2006.
- [12] G. Decker, O. Kopp, F. Leymann, and M. Weske. BPEL4Chor: Extending BPEL for Modeling Choreographies. In *Proceedings of the IEEE 2007 International Conference on Web Services (ICWS 2007)*, pages 296–303, 2007.
- [13] E. Deelman et al. Managing Large-Scale Workflow Execution from Resource Provisioning to Provenance tracking: The CyberShake Example. In *Proceedings of the Second IEEE International Conference on e-Science and Grid Computing*, 2006.
- [14] L. Fredlund. Implementing WS-CDL. In *Proceedings of the second Spanish workshop on Web Technologies (JSWEB 2006)*, 2006.
- [15] A. J. G. Hey and A. E. Trefethen. The Data Deluge: An e-Science Perspective. In *Grid Computing—Making the Global Infrastructure a Reality*, pages 809–824. Wiley and Sons, 2003.
- [16] D. Hollingsworth. *The Workflow Reference Model*. Workflow Management Coalition, 1995.
- [17] J. C. Jacob, D.S. Katz, and et. al. The Montage Architecture for Grid-Enabled Science Processing of Large, Distributed Datasets. In *Proceedings of the Earth Science Technology Conference*, June 2004.
- [18] Z. Kang, H. Wang, and P. C. Hung. WS-CDL+: An Extended WS-CDL Execution Engine for Web Service Collaboration. In *IEEE International Conference on Web Services (ICWS 2007)*, pages 928–935, 2007.
- [19] K. Karasavvas, M. Antonioletti, M. Atkinson, N. C. Hong, T. Sugden, A. Hume, M. Jackson, A. Krause, and C. Palansuriya. Introduction to OGSA-DAI Services. In *LNCS*, volume 3458, pages 1–12, 2005.
- [20] N. Kavantzias, D. Burdett, G. Ritzinger, and Y. Lafon. Web Services Choreography Description Language (WS-CDL) Version 1.0. W3C Candidate Recommendation, 2005.
- [21] S. Krishnan, P. Wagstrom, and G. von Laszewski. GSFL: A Workflow Framework for Grid Services. Technical report, Argonne National Argonne National Laboratory, 2002.
- [22] M. Litzkow, M. Livny, and M. Mutka. Condor - a hunter of idle workstations. In *Proceedings of the 8th International Conference of Distributed Computing Systems*, June 1988.
- [23] D. Liu, K. H. Law, and G. Wiederhold. Data-flow Distribution in FICAS Service Composition Infrastructure. In *Proceedings of the 15th International Conference on Parallel and Distributed Computing Systems*, 2002.
- [24] D. Martin, D. Wutke, and F. Leymann. A Novel Approach to Decentralized Workflow Enactment. *EDOC '08. 12th International IEEE Conference on Enterprise Distributed Object Computing*, pages 127–136, 2008.
- [25] T. Oinn, M. Addis, J. Ferris, D. Marvin, M. Senger, M. Greenwood, T. Carver, K. Glover, M. R. Pocock, A. Wipat, and P. Li. Taverna: a Tool for the Composition and Enactment of Bioinformatics Workflows. *Bioinformatics*, 20:3045–3054, 2004.
- [26] D. Sulakhe, A. Rodriguez, M. Wilde, I. T. Foster, and N. Maltsev. Interoperability of GADU in Using Heterogeneous Grid Resources for Bioinformatics Applications. *IEEE Transactions on Information Technology in Biomedicine*, 12(2):241–246, 2008.
- [27] I. Taylor, M. Shields, I. Wang, and R. Philp. Distributed P2P Computing within Triana: A Galaxy Visualization Test Case. In

*17th International Parallel and Distributed Processing Symposium (IPDPS 2003)*, pages 16–27. IEEE Computer Society, 2003.

- [28] I. J. Taylor, E. Deelman, D. B. Gannon, and M. Shields, editors. *Workflows for e-Science: Scientific Workflows for Grids: Scientific Workflows for Grids*. Springer-Verlag, 2006.
- [29] The OASIS Committee. Web Services Business Process Execution Language (WS-BPEL) Version 2.0, 2007.
- [30] J. M. Zaha, A. Barros, M. Dumas, and A. ter Hofstede. Let's Dance: A Language for Service Behavior Modelling. In R Meersman and Tari Z, editors, *OTM Conferences (1)*, volume 4275 of *LNCS*, pages 145–162. Springer, 2006.



**Adam Barker** is a Lecturer (Assistant Professor) in Computer Science at the University of St Andrews, UK. Prior to obtaining a faculty position, he worked as a postdoctoral researcher at the University of Melbourne, the University of Oxford and the National e-Science Centre (NeSC), University of Edinburgh. He received his PhD from the School of Informatics, University of Edinburgh in 2007. His broad research interests concentrate on the theoretical foundations and effective engineering of large-scale distributed systems. His research agenda is pursued by building and evaluating novel architectures and algorithms based on sound foundations in systems research. For a more detailed biography and list of publications please visit: [www.adambarker.org](http://www.adambarker.org)



**Jon B. Weissman** is a leading researcher in the area of Cloud and Grid computing. His involvement dates back to the influential Legion project at the University of Virginia during his Ph.D. He is currently an Associate Professor of Computer Science at the University of Minnesota where he leads the Distributed Computing Systems Group. His current research interests are in Grid computing, distributed systems, high performance computing, resource management, reliability, and e-science applications. He works primarily at the boundary between applications and systems. He received his B.S. degree from Carnegie-Mellon University in 1984, and his M.S. and Ph.D. degrees from the University of Virginia in 1989 and 1995, respectively, all in computer science. He is a senior member of the IEEE. As well as being a Visitor at the Institute, Jon has been appointed an Honorary Fellow of the College of Science and Engineering at the University of Edinburgh.



**Jano I. van Hemert** has a PhD in Mathematics and Physical Sciences from the Leiden University, The Netherlands (2002). He is a project manager and academic liaison at Optoson innovative retinal imaging company with a vision to be recognised as the leading provider of retinal diagnostics. He has held research positions at the University of Edinburgh (UK), Leiden University (NL), the Vienna University of Technology (AT) and the National Research Institute for Mathematics and Computer Science (NL). In 2004, he was awarded the Talented Young Researcher Fellowship by the Netherlands Organization for Scientific Research. In 2009, he was recognised as a promising young research leader with a Scottish Crucible. In September 2010, he became an Honorary Fellow of the College of Science and Engineering at the University of Edinburgh in recognition of his work for the School of Informatics.