

The SAHARA Model for Service Composition Across Multiple Providers

Bhaskaran Raman, Sharad Agarwal, Yan Chen, Matthew Caesar, Weidong Cui,
Per Johansson*, Kevin Lai, Tal Lavian[†], Sridhar Machiraju, Z. Morley Mao,
George Porter, Timothy Roscoe[‡], Mukund Seshadri, Jimmy Shih, Keith Sklower,
Lakshminarayanan Subramanian, Takashi Suzuki, Shelley Zhuang,
Anthony D. Joseph, Randy H. Katz, Ion Stoica

University of California at Berkeley

Abstract. Services are capabilities that enable applications and are of crucial importance to pervasive computing in next-generation networks. *Service Composition* is the construction of complex services from primitive ones; thus enabling rapid and flexible creation of new services. The presence of multiple independent service providers poses new and significant challenges. Managing trust across providers and verifying the performance of the components in composition become essential issues. Adapting the composed service to network and user dynamics by choosing service providers and instances is yet another challenge. In SAHARA¹, we are developing a comprehensive architecture for the creation, placement, and management of services for composition across independent providers. In this paper, we present a layered reference model for composition based on a classification of different kinds of composition. We then discuss the different overarching mechanisms necessary for the successful deployment of such an architecture through a variety of case-studies involving composition.

1 Introduction

Pervasive computing demands the all-encompassing exploitation of services inside the network. By services, we mean both the components of distributed applications and the glue that interconnects them as they function across the network. Services range from providing basic network reachability to creating overlay networks with enhanced qualities like predictable latencies and sustained bandwidths. Services also include instances of application building blocks, requiring processing and storage, judiciously placed in the network to control connection latencies and to achieve scale through load sharing. Such services may be simple format translators, interworking functions, or major subsystems for content distribution or Internet search, which are often regarded as applications in their own right. Composition via interconnection of services allows more sophisticated services and applications to be constructed hierarchically from more primitive ones. Since economics makes it unlikely that any single service provider will be able to provide all of the connectivity, applications building blocks, processing, and

* Ericsson Berkeley Wireless Center; [†] Nortel Networks; [‡] Intel Berkeley Research Lab

¹ Project supported by Sprint, Ericsson, NTTDoCoMo, HRL, and Calif. Micro Grant #01-042.

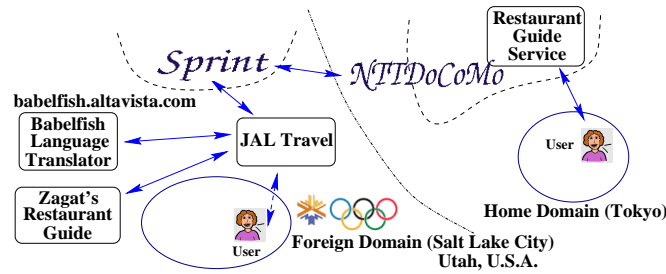


Fig. 1. Service Composition: A Multi-Provider Scenario

storage resources to effectively deploy a globe-spanning application, the composition of services across independent providers is essential. This paper proposes a comprehensive reference model for composed services in support of pervasive computing.

To illustrate our concept of service composition across service providers, consider the following scenario (Figure 1). Ms. Tanaka travels from Tokyo to Salt Lake City to attend the Winter Olympics. Her cellular provider, NTTDoCoMo, maintains roaming agreements with several foreign network operators, such as Sprint, so she can make and receive calls in the U.S.. But her new information appliance, built by Ericsson, is much more capable than just a phone – it is a gateway to extensive information, entertainment, and messaging services. In Japan, she subscribes to NTTDoCoMo’s restaurant recommendation service, and would like to use the same application in Utah. In the U.S., Sprint has an arrangement with Zagat’s Guide to present such information to its subscribers, but the text is in English and formatted for presentation on a different kind of display than Ms. Tanaka’s. A third party, JAL Travel, assembles a special new service for Japanese tourists at the Olympics from component services: Zagat’s Restaurant Guide, Japanese translation using Babelfish, and reformatting for Japanese-style information appliance displays. Ms. Tanaka subscribes to this service for two weeks, and the usage charges appear on her NTTDoCoMo bill back in Japan.

This scenario exemplifies several key points about services composition. In next generation networks, users will demand enhanced services like restaurant recommendations, but they expect them to look and feel the same whether they are at home or in a foreign network. A new service, *composed* from localized information sources, appropriate language translators, and content reformatters, makes the underlying differences transparent to the user. Such a service can be created quickly by simply connecting the components across the network. But many entities participate in the realization of this service, and they must be managed, authenticated and compensated in some way.

In the context of object-based systems, programming by composition across the network is hardly new. Yet there are critical new challenges. The first is compositions across independent service providers: Zagat owns the database and the web site, NTTDoCoMo maintains the relationship with the client and collects her monthly bill, and Sprint provides the local wireless access as well as the gateway access to local services. Creating the necessary mechanisms to support cooperative composition of

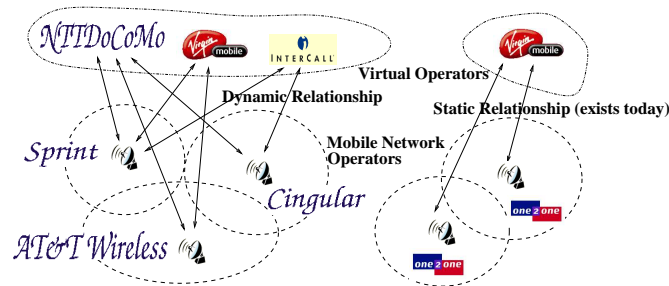


Fig. 2. Multi-provider Scenario in Wireless Connectivity Service

services across independent service providers, each with its own authentication, authorization, and accounting mechanisms, is an essential challenge.

JAL Travel created a new service from pieces, some of which were provided by other third parties, such as Babelfish, and all of which need to run on machines interconnected across the network, spanning Internet Data Center providers and Internet Service Providers. Herein lies the second challenge: the ability for third parties to discover components and to broker new services from constituent pieces, some of which may not even be aware of the composition in which they are participating. As the qualities of a composed service are no better than its weakest component, an essential need is for brokers to be able to verify the performance and behavior of the assembled components, whether or not these underlying participants are aware of their role in compositions. If a component does not meet its performance or behavioral specification, it must be “composed out”, and a new instance from a different provider “composed in”.

A third challenge for service providers is the need for an extensive set of new service composition management tools. From a provisioning viewpoint, sufficient instances of the components need to be placed at locations within the network to ensure scalable performance and high availability even in the face of site failures or network outages. Such placement also needs to ensure appropriate network and processing latencies to achieve adequate responsiveness for the supported applications. Such tools include a policy management mechanism for service providers to inform service composers about how their instances in their network should be used for providing fault tolerant and load balanced behavior. A pervasive monitoring and measurement infrastructure is needed to detect changing access patterns and shifting workloads, to drive redirection to unloaded service instances or to change the number and placement of deployed instances. Network topology-awareness is important, for availability as well as performance. Placement and connectivity issues are complicated since some service instances are anchored to fixed locations, while lighter weight services can be placed close to the user community.

Further challenges arise when we consider user dynamics. A large number of foreign roamers like Ms. Tanaka converge in Salt Lake City, yielding flash crowds and over-utilized spectrum. We need new ways of efficiently allocating resources in the context of new service provider business models. The 3rd Generation Partnership Project

(3GPP) defines the concept of a Mobile Virtual Network Operator (MVNO) [1], an entity with subscribers but no network. The MVNO provides wireless connectivity service composed from the physical network resources of underlying wireless operators. In our example, NTTDoCoMo acts as a virtual operator using spectrum for its subscribers from the already established network operators such as Sprint, Cingular, and AT&T Wireless. We view such a multi-provider relationship as just another case of *service composition*. Today, the relationships between the virtual operator and the Mobile Network Operator (MNO) are static and negotiated long in advance (e.g., between *Virgin Mobile* and *One2One* in UK [1], see Fig. 2). But this is inefficient when user dynamics are considered; and we expect much more dynamic formation and dissolution of relationships in the future. Thus, as a fourth challenge, we address issues in efficient resource allocation across providers, considering the dynamics of user communities. In our MVNO scenario, this translates to dynamically selecting (“roaming”) among co-located MNOs. Figure 2 illustrates such dynamic allocations in the MVNO context. We envision that dynamic relationships will last for short time-scales of minutes to hours, thus allowing for load balancing, and efficient resource usage.

Our overall goal is to define *a comprehensive reference model that is able to describe the assembly from components of end-to-end services with desirable, predictable, enforceable properties, yet spanning potentially uncooperating service providers*. We are developing these concepts in the context of the SAHARA² project, which is also the name of our prototype architecture. The next section summarizes the discussion above with the technical issues in composition. Sec. 3 presents a classification of the different kinds of composition and Sec. 4 presents a layered reference model for composition. We describe the different mechanisms we employ to address the technical issues in composition in Sec. 5. We discuss related work in Sec. 6 and conclude in Sec. 7.

2 Technical Issues

The scenario in our prior section can be understood along three dimensions in the choices in service composition: (a) what set of services to use for composition, (b) which service providers’ resources to use, and (c) which instances of each service to use for a particular client session. In addition, we also have the issue of who makes these decisions. Considering these three dimensions, the technical issues that must be addressed in a reference model for service composition are:

- *Trust management and behavior verification*: When multiple providers interact, it is important to establish mutual trust. This is not only for the purpose of user authentication and billing, but also to verify the behavior of the components in composition. Does a component meet its promises in terms of functionality, protocols, performance, availability, or other properties?
- *Adapting to network dynamics*: With network dynamics, workloads can shift, congestion can arise, and reachability to services can be lost. This implies the need for performance monitoring, modeling and prediction, and a performance-sensitive choice of providers and instances. We term this the *service selection* problem.

² SAHARA: Service Architecture for Heterogeneous Access, Resources, and Applications; <http://sahara.cs.berkeley.edu>

- *Adapting to user dynamics*: User dynamics can cause different providers to face varying demand for physical resources from their users. Allocation of resources (spectrum, network bandwidth, CPU, etc.) across providers, based on current demand is important to achieve fair/utility-driven resource allocation.
- *Resource Provisioning and Management*: For a given community of users and a set of performance, availability, and administrative constraints, how many instances of a service are needed? How can this be optimized given a knowledge of the provisioning done by other service providers, and network topology information? We call this the *service placement* problem. Also, how can service providers enforce their local policies of the use of their service instances in a composition performed by a third party? We term this the *policy management* issue.
- *Interoperability across multiple service providers*: When composing services across providers, we have to deal with heterogeneity in protocols and data formats, as well as authentication and authorization mechanisms.

3 Service Composition Models

We now classify service composition into two models based on the type of interaction between composed component service providers and analyze their pros and cons: (1) **Cooperative Model**: Service providers interact in a distributed fashion, with distributed responsibility, to provide an end-to-end composed service; (2) **Brokered Model**: A single provider, the broker, uses the functionalities provided by underlying service providers and encapsulates these to compose the end-to-end service.

In either case, the end-user subscribes to only one provider. However, the difference lies in the way the responsibility for the composed service is apportioned. The two possibilities represent different business models for composition. In the cooperative model, the properties of the composed service such as functionality, performance, and availability, are guaranteed by the design of the distributed interaction, and through service-level agreements between the interacting entities. Each service provider is only responsible for providing guarantees for the portion of the composed service within its domain. In the brokered model, the broker assumes responsibility for the properties of the composed service. We can imagine a broker entering into contracts with service providers, and using these to construct end-to-end composed services. The broker verifies the functionality of the individual pieces in the service path. This is because individual component providers may not trust each other – they may limit the information about the state of their service they expose to the other providers in the composition, or may actively seek to cheat on the quality of service they provide.

An example of cooperative composition is cellular roaming as a service, composed from the resources of multiple mobile network operators, as in Sec. 1. The distributed interaction between operators (NTTDoCoMo and Sprint) enables roaming. Another example is end-to-end connectivity service in the Internet. An inter-domain routing protocol allows cooperation between domains in a distributed fashion. In these examples, there are long-term, static, negotiated contracts between the participant providers.

An example of brokered composition is the restaurant guide service assembled by JAL Travel in our earlier scenario. It assumes responsibility for the functionality and

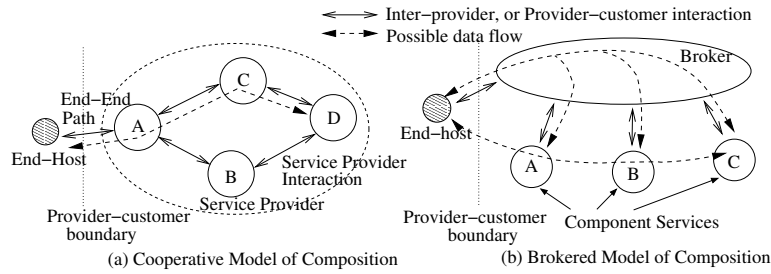


Fig. 3. Service Composition Models

performance of the composed service. Another example is the Yahoo portal service that composes third party services such as the Google search engine, stock ticker and news.

Fig. 3 illustrates these alternate composition models. The models of composition say nothing about the data flow, only the nature of interaction, or the business model, between providers. As an example, in the brokered model, we could have data flow through the broker, who assembles it (shown by the dotted lines on the top). Instead, the broker could set up the data exchange and not be in the data path (shown by the dotted line at the bottom of the figure).

The same composed service could be implemented in either model. Consider the provision of a connectivity service with QoS guarantees between two points on the Internet. In a cooperative model, ISPs enter into service level agreements that specify mutual QoS guarantees. These may be stitched together in a distributed fashion to offer end-to-end guarantees (www.merit.edu/working.groups/i2-qbone-bb). In a brokered model, a provider like InterNAP (www.internap.com) purchases pipes with specified guarantees from individual ISPs and uses them to provide QoS to its customers.

Each of these models is suited to a particular environment. In the cooperative model, providers work together and can share performance information to ensure end-to-end properties of the composed service. However, they must rely on each other, which leads to issues of trust. Since the responsibility for the composed service is distributed, each provider must continuously verify that the others with whom it has agreements meet their service specifications. These specifications are in terms of functionality, protocol, performance, and availability. Such comprehensive verification is absent in the cooperative composition of Internet connectivity service across network domains today.

In the brokered model, the broker composes an end-to-end service by selecting individual services residing in different domains. This simplifies service deployment since the members of the composition need not agree among themselves, only with the broker. This also enables the composition of services across competing service providers. The broker assumes responsibility for constructing the entire end-to-end service path. Brokering is a powerful tool to construct services from providers who are not necessarily aware that they are participating in a larger end-to-end service (e.g., in Figure 1, Zagat's restaurant guide service does not know it is being composed with Babelfish's language translation service). However, because the broker has limited visibility into the underlying provider resources, sub-optimal utilization of provider resources may

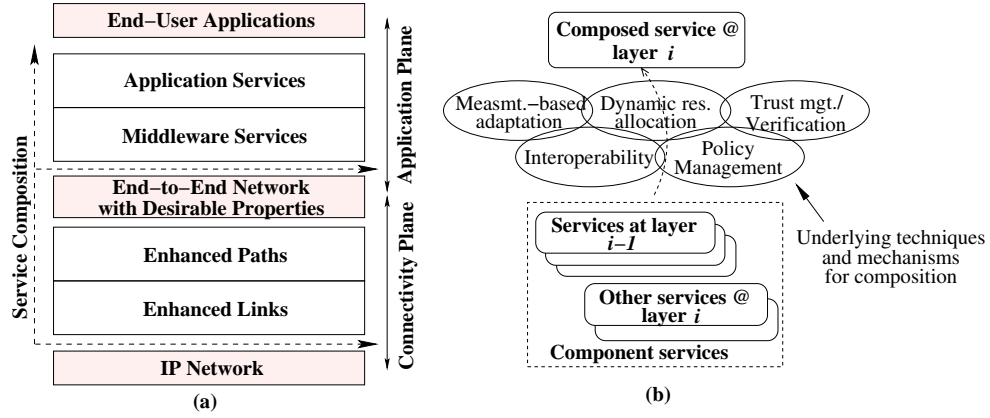


Fig. 4. A layered reference model for Service Composition

result. The nature of the composition and the relationships of the underlying service providers will determine which model is most appropriate.

4 Service Composition: A Layered Reference Model

The previous section classified service composition by how responsibility is shared across providers. Here we present a layered reference model for composition across different layers of concern.

Composed services build on top of connectivity provided by the IP layer, the bottom slice in Fig. 4(a). We assume that IP provides reachability, and build several “desirable” properties in the *connectivity plane*. These include features such as performance guarantees (e.g., latency, bandwidth, loss-rate bounds), availability guarantees (e.g., available 99.99% of the time), as well as functionality or protocol guarantees (e.g., guaranteeing that a particular advertised Internet route is valid). We achieve them through composition at the connectivity plane, across multiple providers. This results in an “end-to-end network with desirable properties”, as shown in the middle slice of the figure.

The connectivity plane is further divided into two layers: *Enhanced links*, and *Enhanced paths*. The enhanced links abstraction is built between two interacting entities: between two service providers, or between a provider and the end-user. This abstraction achieves desirable performance oriented properties, and verification of routing protocols between peering entities. Verification checks if the routes advertised by a peer are valid. An example of this composition is of the MVNO choosing between multiple MNOs. Here the enhancement is the improved performance through reduced call blocking rate, due to dynamic load sharing across multiple MNOs.³

³ Although it is not obvious here how this example builds on IP technology, it will be clear after a discussion of the mechanisms for resource allocation, in Sec. 5.

Enhanced paths build on enhanced links, and provide desirable properties in an end-to-end path between points on the Internet. The path spans multiple service domains, and can be chosen adaptively via resource allocation across providers to meet performance or availability constraints. Our reference model is independent of whether the performance guarantees are strict, or simply “enhancements” to the best-effort Internet. Alternative enhancements might be appropriate for different end applications.

The top half of the figure represents the application plane. These layers support end-user applications and are in turn built on top of the end-to-end network. The *middleware services* are enablers, such as the Babelfish language translation service, content-distribution networks or video/audio transcoders. The *application services* layer consists of services useful to end-users, such as the Zagat’s restaurant guide, search engines or a voice-mail service. Composition at the application plane results in enhanced functionality. In our example, the enhanced functionality is that of the restaurant guide appearing in the user’s native language (Japanese) and presentation style (NTTDoCoMo user interface) in a foreign network.

Service composition can be applied within and across these layers, as illustrated in Fig. 4(a). Fig. 4(b) shows this more explicitly. A composed service at a higher layer is composed of multiple services in the same layer, or of services at the layers below. In Fig. 1, composition takes place across the application (restaurant guide), middleware (language translation), and connectivity plane (roaming service enabled by Sprint) layers. In Fig. 2, composition is at the enhanced link layer, using the component connectivity services offered by co-located MNOs.

We note that this layerization is only a reference model, and compositions need not strictly adhere to it. Some application services can be composed directly of enhanced links, without using the enhanced paths abstraction or middleware services. This will occur typically in performance sensitive applications, where the composer needs full visibility via a flat rather than a hierarchical composition. We next discuss the techniques and mechanisms used in enabling composition as shown in Fig. 4(b).

5 Mechanisms for Service Composition

The issues we listed in Sec. 2 appear in different flavors in the alternate models of composition. SAHARA is our architectural prototype to explore the mechanisms required to address these issues. To understand the various mechanisms, we are working on several case studies of composition. These cover the dimensions of classifications in last two sections. We now describe the various mechanisms that we are designing to address the challenges presented earlier and how they are used in the individual case studies. We have performed in-depth evaluations of several of these mechanisms. Due to space limitations, we summarize the evaluations of only a subset of these mechanisms.

5.1 Measurement-Based Adaptation

For service composition, it is desirable to dynamically choose service providers, and service instances based on current network and server loads. Measurements can be carried out by a third party measurement service – a common element of the service infrastructure, or by the composer itself. This applies to both the cooperative and brokered

models of composition. We now describe two services that we are developing involving measurement-based adaptation.

We are developing a general end-to-end Internet host distance estimation service. Such a service is especially useful in a brokered composition since the broker does not have insight into the network characteristics of individual providers. Given the potential large numbers of service providers and instances, to scale the measurement service, we cluster the end hosts to be monitored based on the similarity of their perceived distance to the measurement points. The cluster center is then used as a single measurement target for future monitoring. Simulations with real Internet measurement data show that our scheme has good prediction accuracy and stability with a small communication and computation cost [2].

An application service that we have developed is the Universal Inbox [3], a metaphor for any-to-any communication across heterogeneous devices and networks. Data transformation services such as audio/video transcoders and text \leftrightarrow speech engines are extensively used to adapt content between communicating devices. For example, in our earlier scenario, Ms. Tanaka's email service could be composed with a text-to-speech conversion service so that she can listen to emails over her cellular-phone. We use the brokered model of composition here. To adaptively choose service providers and instances, we have designed a middleware measurement layer that exchanges network and server load using a link-state algorithm [4]. This exchange takes place across service execution platforms, enabling a dynamic choice of service instances, possibly in the middle of the user sessions, to hide network and server failures from end-users.

A critical challenge that we address in the context of such an application is *availability*. When composed services span multiple providers, data could traverse the wide-area Internet. We detect and recover from Internet path failures *quickly* by using the middleware measurement layer to choose alternate service instances for the client session. Our measurements in [4] show that Internet path failures that happen due to congestion or other factors can be detected reliably within about 2 *seconds*. Further, subsequent recovery by using alternate service instances can be completed in a *few hundred milliseconds*. Thus, network path failures lasting several tens of seconds to minutes [5] can be completely masked from the end client.

In addition to these two services, we have also designed and implemented a measurement methodology [6] to improve DNS-based server selection, which is a common technique used by Content Distribution Networks (CDNs) today. Our technique enables the collection of client to local DNS server mappings to allow more accurate server selection based on a client's local DNS server. Understanding the distribution of HTTP requests corresponding to local DNS servers also enables better load prediction given a DNS request and thus improved server selection mechanisms.

5.2 Utility-Based Resource Allocation Mechanisms

In a multi-provider environment, different providers may experience different demands for resources due to user dynamics. Demand or utility-based resource allocation can be applied within a service provider to manage its instances. In the brokered model, it can be used to allocate resources across providers. In SAHARA, we are exploring two resource allocation mechanisms.




Fig. 5. Congestion-sensitivity in QoS-based redirection

Auctions are one way of constructing a marketplace where a resource, such as bandwidth or physical spectrum, can be dynamically allocated. Auctions allocate resources to consumers based on their bids, which represent the value of the good to them. Furthermore, the resource can be subdivided into units, and multiple bidders can be allocated the resource until the resource pool is exhausted. Auctions can occur in rounds, where the allocation determined by each round can be for some future time period. Auction-based allocations in a multi-provider environment can provide the mechanism for demand-based resource allocation. For instance, in the MVNO scenario in Section 1, spectrum resources could be auctioned off every few minutes to competing virtual operators based on their current user-load in the area of coverage.

Congestion pricing is an allocation mechanism that assigns scarce resources to consumers using the abstraction of price as a means to moderate demand. During high demand, such a market ensures that the price increases. Only those consumers with the greatest need and having sufficient currency will obtain the needed resource. During low demand, the price drops, and access to the resource will be cheap and plentiful. This approach should yield an assignment of resources (supply) to the need (demand) that adapts to instantaneous demands [7].

As an application that uses the congestion pricing mechanism, we have looked at the selection of Voice-over-IP (VoIP) gateways across multiple providers [8]. These gateways are deployed by independent entities, and exchange dynamic pricing information as well as the peering relationships between the provider entities based on the IETF TRIP (Telephony Routing over IP) protocol. The price is decided based on the load, or congestion, at each gateway; and the user gets to choose between several gateways based on the price and the required quality of service. This achieves pricing-based load-sharing among the gateways. In [9], we looked at the trade-off between using QoS-based redirection (nearest VoIP gateway) and congestion-based redirection. This is shown in Figure 5 where we see that incorporating congestion-sensitivity in QoS-based redirection can improve call blocking rate by as much as a factor of three in comparison to random redirection. We also observe that increasing congestion-sensitivity does not significantly degrade the QoS of the VoIP calls.

5.3 Trust Management and Verification of Service/Usage

An important issue for composition is the establishment and monitoring of trust relationships between inherently untrusting entities. This is important in cooperative composition where providers have pairwise service agreements between them. This is also important in brokered composition where the agreements are with the broker. Typically a AAA (*Authentication, Authorization and Accounting*) server governs service instances and users within one administrative domain. However, in our scenario, we need to compose services across domains governed by multiple, different AAA servers.

We are investigating an authorization control scheme with credential transformations to enable cross-domain service invocation [10]. Federated administrative domains form credential transformation rules based on established peering agreements. These are used by a AAA server to make authorization decisions for a service request from an affiliated domain.

Another important issue in service composition is to verify whether the provided service adheres to the desirable properties advertised by its provider. Such properties can be specified in a bilateral Service Level Agreement (SLA) between provider and requester. We use parameter verification and usage monitoring as mechanisms to ensure that the properties specified in the SLA are being honored. For instance, in a case-study of connectivity composition across domains, we have border routers monitoring control traffic from different providers to detect malicious route advertisements.

5.4 Policy Management

An advantage of cooperative composition is that each provider has visibility into its network of services, while a broker does not. The disadvantage of this distributed form of composition is the lack of central control over the composition that the broker enjoys. This disadvantage can be minimized if some form of distributed policy management is in place. Specifically, in cooperative composition, the service composition policies of one service provider can be made visible to and applied at distant service providers further along the composed path. Such policies may include which service instances are for primary use and which are solely for use in various failure modes, and policies that govern load balancing between instances.

A case study involving this principle is the distributed application of policies in inter-domain routing on the Internet. The BGP protocol was not designed to allow local policies to be imposed at distant points in the network. The Internet infrastructure has been plagued of late with pathological routes that attempt to impose such inter-domain routing policies by tricking distant service providers' route selections [11].

To address this issue, we are developing an architecture [12] that allows distributed policy management between service providers. By negotiating policy changes at various points in the topology using a map, policy agents can improve load balancing and fault tolerance. We build an AS relationship map in [13] by using BGP measurements from multiple vantage points. The map indicates the nature of inter-provider relationship that exists between neighbor ASes: peer-to-peer or customer-provider. Figure 6 shows a summary of the results from this study.

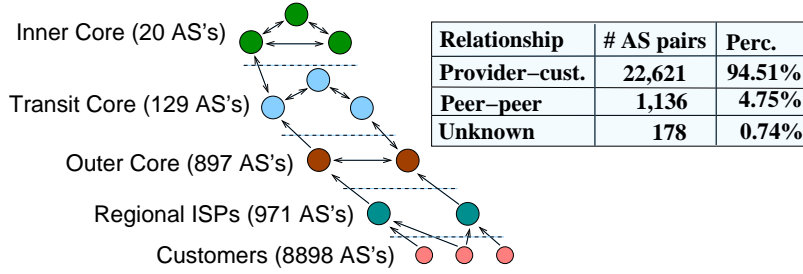


Fig. 6. Internet AS Hierarchy and Inferred Relationships for 23,935 AS Pairs

5.5 Interoperability through Transformation

Interoperability is important for service composition across different service providers. Data and protocols formats may need to be transformed for interoperability across heterogeneity. This is an important issue for composition in the connectivity plane and in the application plane where services have more complex interfaces. Defining the service interface and propagating it are two key challenges here. We have investigated one such interoperability service at the connectivity plane. In our *broadcast federation* work [14], a global multicast service is composed from the multicast implementations in different provider domains. We use protocol transformation gateways between administrative domains that have non-interoperable implementations of multicast.

5.6 Service Deployment

In any composition model, a service provider has to decide how many instances to place and where to place them. In SAHARA, we are exploring two mechanisms to enable such decisions. (1) First, we infer inter-provider relationships between Internet ASes by passively monitoring inter-domain routes from multiple locations [13]. Such a relationship map allows us to understand traffic flows in the underlying network topology by giving more information about reachability between points in the network, than a simple connectivity graph. This richer reachability information helps in service instance placement – the provider can ensure that there are sufficient instances with good reachability from various points on the Internet. (2) Next, we have designed a dynamic replica placement protocol that enables an application level multicast tree for Web content distribution. This can meet client QoS constraints and server capacity constraints, while retaining *efficient* and *balanced* resource consumption [15]. To provide a scalable solution, we aggregate the access patterns of clients that are topologically close in the network. We use incremental clustering and distribution based on client access patterns to adaptively add new documents and purge old ones from the content clusters [16].

6 Related Work

Related work falls into two categories: (a) Architectures for seamless integration of devices and services, and (b) Internet-based web-services initiatives.

Architectures for Seamless Device/Service Integration: The UMTS model [17] admits of a sophisticated accounting, billing, and settlements architecture to support third-party brokering between subscriber needs for service and multiple service providers. However, there is no explicit consideration of where service provision and service mediation should exist in this architecture, other than in the core network that ties together various access networks. The Virtual Home Environment (VHE) concept of IMT2000 [18] permits users to roam away from home, seeing the same service interface (service mobility). ICEBERG [19] looks at extensible personal and service mobility. Nevertheless, these efforts do not consider composition of services across multiple providers, efficiency through network awareness, or resource allocation issues.

TINA [20] is a CORBA-based [21] service architecture. The TINA reference architecture contributes the conceptual separation of the business model, the informational model, and the computational model. Its three layer model of applications, distributed processing, and network environment has influenced our layerization of composition in SAHARA. Key differences are that SAHARA adds elements of composition across heterogeneous providers, with a greater awareness and management of the underlying network topology. SAHARA also considers resource management via placement, allocation, redirection to services and resources.

Internet-Based Web-Services Initiatives: There are several industrial initiatives to enable web-services which “integrate PCs, other devices, databases, and networks into one virtual computing fabric that users could work with via browsers” [22]. These include HP’s e-speak (www.e-speak.net) and web-services platform, Microsoft’s .NET, and Sun ONE. These are based on a language for description of web services (WSDL), a common wire format for these descriptions (SOAP), and a registry to support service location (UDDI). Microsoft’s .NET also defines a language independent software platform for easy and secure interoperation of applications [22]. However, these do not define a wide-area service architecture; they are complementary to SAHARA’s goals of service placement, resource allocation, and network awareness aspects.

7 Conclusions

In this paper, we have presented a vision of distributed systems composed from services placed in the wide-area Internet, and spanning service providers at different levels. SAHARA is our evolving architectural prototype for the creation, placement, and management of services in next generation networks. Our goal is to enable end-to-end service composition with desirable, predictable and enforceable properties spanning multiple potentially distrusting service providers. We investigate two forms of service composition under different business models with varying degrees of cooperation and trust among providers. We classify component services and composed services into a layered hierarchy. The overarching themes in the various techniques and mechanisms that we use for composition include (a) measurement-based adaptation through dynamic choice among service providers and service instances, (b) utility-based resource allocation for demand-driven load sharing across provider resources, and (c) a trust-but-verify approach to management of trust and behavior verification when multiple providers interact to provide a composed service. We continue to develop these mech-

