On trust guided collaboration among cloud service providers

Liu Xin
School of Computer Engineering
Nanyang Technological University
Singapore
Email: liu_xin@pmail.ntu.edu.sg

Anwitaman Datta
School of Computer Engineering
Nanyang Technological University
Singapore
Email: anwitaman@ntu.edu.sg

***Abstract*—Cloud computing has emerged as a popular paradigm that offers computing resources (e.g. CPU, storage, bandwidth, software) as scalable and on-demand services over the Internet. As more players enter this emerging market, a heterogeneous cloud computing market is expected to evolve, where individual players will have different volumes of resources, and will provide specialized services, and with different levels of quality of services. It is expected that service providers will thus, besides competing, also collaborate to complement their resources in order to improve resource utilization and combine individual services to offer more complex value chains and end-to-end solutions required by the customers. It is challenging to select suitable partners in a decentralized setting due to various factors such as lack of global coordination or information, as well as diversity and scale. Trust is known to play an important role in promoting cooperation in many decentralized settings including the society at large, as well as on the Internet, e.g., in e-commerce, etc. In this paper, we explore how trust can promote collaboration among service providers. The novelty of our approach is a framework to combine disparate trust information - from direct interactions and from (indirect) references among service providers, as well as from customer feedbacks, depending on availability of these different kinds of information. Doing so provides decision making guidance to service providers to initialize collaborations by selecting trustworthy partners. Simulation results demonstrate the promise of our approach by showing that compared to random selection, our proposal can help effectively select trustworthy collaborators to achieve better quality of services.**

Keywords: cloud computing, collaboration, value chain, trust management, Dirichlet distribution

I. INTRODUCTION

Cloud computing provides resources such as computing or storage as well as software to end-users who can utilize these resources as and when they have the need, without having to invest in the infrastructure. Besides the technical solutions to provide such scalable and elastic solutions to end-users, there are several practical issues that can affect the business model and viability of cloud service providers. At present, the cloud computing market is dominated by big players such as Amazon Web Service [17], Google App Engine [4], GoGrid [5] or Windows Azure [2], who all create their own closed network. Such an environment is not conducive for smaller players, nor for end users because of various reasons, as described next.

From the perspective of service providers, if operating solo, service providers need to provision resources based on anticipated peak-hour needs. The trade-off is between inability to accept business opportunities or sacrifice on availability or quality of service provided. Provisioning for such extra resources needs not only larger investment, but furthermore such resources will stay idle for long stretches of time. This is particularly significant for small-scale service providers providing niche services, so that it can not multiplex the resource usage.

From the perspective of customers, they are restricted to services offered by a single provider at a time, and thus can not enjoy 'ready to use' multiple or collaborative cloud services. Moreover, customers may suffer from vendor lock-in issues.

Recent works such as [7] and [6] identify the motivations and possibility for collaboration among cloud service providers. This will lead to better amortization of their individual resources, and also will enable the service providers to compose their individual services in a value chain to offer the end-users more complex, end-to-end solutions.

Dealing with cloudbursts is a simple and obvious scenario where such collaboration is useful. A slightly more complex scenario is as follows. A company may need to put its business process application in a cloud which provides extremely high availability and powerful computational capability (and corresponding sophisticated analytic softwares) for routine business data analysis. It may also need an inexpensive solution for data archival, where availability and throughput are less critical since the archived data is accessed infrequently, but durability of the data and the cost of storage are very important. Smaller cloud service providers may specialize in each of these individual services. By collaborating together to provide a composite service, they can compete against bigger players who have in-house closed end-to-end solutions. The customer in turn gets cheaper alternatives, without the burden of negotiating and integrating piece-meal solutions provided by the individual service providers.

The formation of a successful collaborative group of service providers may be abstracted in four logical steps [7]: (1) A service provider identifies a great business opportunity or

The work presented in this paper has been supported by A*Star TSRP SERC Grant number 102 158 0038 for the pCloud project. http://sands.sce.ntu.edu.sg/pCloud

COLLABORATECOM 2010, October 9-12, Chicago, USA / Copyright © 2011 ICST 978-963-9995-24-6
DOI 10.4108/icst.trustcol.2010.6

other scenarios which need collaboration with other service providers to offer a set of new services to the customers. We call such service provider who initializes the collaboration *master service provider* (MSP) and other service providers that are invited in the collaboration *guest service provider* (GSP). (2) MSP selects GSPs from a pool of candidates based on any criterion. These service providers will sign an informal contract describing responsibilities taken by each provider to form a temporary collaborative group. (3) The collaborative group of service providers act as an entity to bid for a business opportunity. (4) Once winning an opportunity, the service providers will sign a formal contract and start collaboration.

This paper focuses on the second step, i.e., selecting trustworthy partners to form a good collaborative group. This is a non-trivial task because of various reasons. Complexities and conflicts of business interests may make it difficult to find service providers who are willing and able to collaborate. A service provider who needs to find a partner to provide storage service may be unable to find specific providers due to their own resource limitations, or because of government regulations and legislations [3], Service Level Agreement [12]; or such service providers may simply reject the collaboration according to their competitive strategies. From technical perspective, the service providers may not be able to offer high quality services (e.g. services are interrupted, data integrity is compromised, etc.) due to the underlying infrastructure and data management mechanisms, or even act maliciously for their own self-interests (e.g. stealing sensitive data). To address these technical issues, we propose trust based collaborator selection framework to help MSP select trustworthy GSPs by considering diverse information sources.

We identify three scenarios where MSP possesses different information sources about the potential GSPs: (1) MSP has direct experience with the candidates, i.e., MSP has collaborated with candidates before; (2) MSP has indirect experience, i.e., feedbacks about the candidates from other service providers; (3) MSP does not have any (direct or indirect) collaboration experience with these candidates. For the first scenario, we abstract every collaboration between MSP and a candidate as a transaction between the two providers. After completing such a transaction, MSP can rate it by specifying whether it was *good*, *medium* or *bad*, or other discrete rating representations. Using these ratings, MSP could apply Dirichlet distribution based statistical method [10] to estimate how trustworthy each candidate is in the future collaboration. When indirect experience is available, it is combined with direct experience to provide comprehensive assessment of the candidates by taking into account other service providers' opinions. Note that to combine direct experience and indirect experience, the ratings of all transactions should be represented in the same manner. In the worst scenario, i.e., when neither direct experience nor indirect experience is available, MSP makes prediction of the candidates' behavior by taking into account their customers' reviews. Customers' reviews may not accurately reflect candidates' real behavior in a collaboration, however, they indeed show the general quality of the candidates thus are useful particularly in absence of (more relevant) information.

The contributions of this paper are summarized as follows: (1) We emphasize and motivate the collaboration among service providers in cloud environment, and provide a generic trust based platform to promote such collaboration. (2) In the process of GSPs selection, we propose a framework to combine various trust information - from direct interactions to references among service providers, as well as from customer feedbacks, depending on availability of these information. (3) We conduct experiments with artificial workloads to make a preliminary assessment of the efficacy of our approach. Results show that our proposal indeed can help MSP find trustworthy GSPs, leading to better quality of collaboration.

The rest of this paper is organized as follows: Section II introduces the collaboration platform formed among service providers. In Section III, we present trust based GSP selection framework relying on service providers' past behavior. Our experiments and results are discussed in Section IV. Section V summarizes related works. Finally we conclude our work and suggest future research directions in Section VI.

II. Service provider collaboration platform

In order to cater to a dynamic cloud market and amortize individual service providers' infrastructure investments, as well as complement their specialities, collaborations among service providers are essential. Such collaboration can be carried out in a policy-driven manner. A well structured collaboration platform is needed to efficiently find and reasonably coordinate the activities of such service providers.

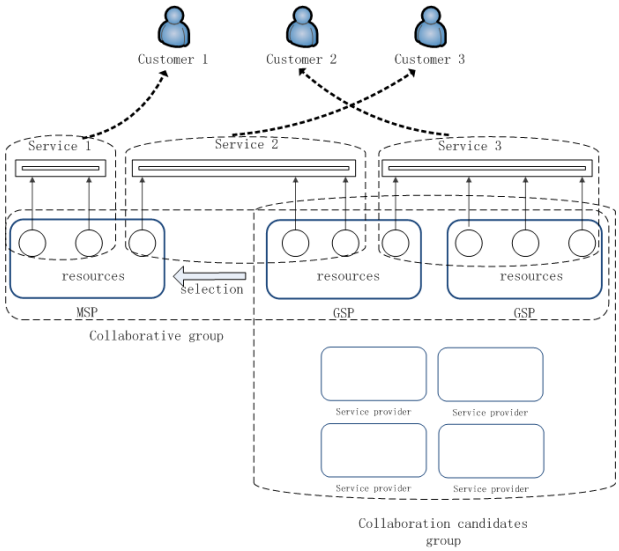

Fig. 1. A service provider collaboration platform.

Figure. 1 shows a tentative cloud service provider collaboration platform. A collaborative group is initiated by MSP on identifying a good business opportunity requiring collaboration with other relevant service providers to complement the MSP's competence. MSP selects trustworthy potential GSPs based on certain metrics, which will be described in the next section.

After the collaborative group is formed, the service providers act as an entity to try to get the business opportunity. Once successful, the collaborative service providers share their local resources (represented by circles in the figure) to produce new composite services as required by the customers. For instance (in Figure 1) MSP is able to provide service 1 independently using its own resources; but service 2, which is required by customer 3, needs to integrate resources from both MSP and GSP.

III. Trust based collaborator selection framework

Selecting trustworthy GSPs to facilitate collaboration effective and efficient is desirable. In this section, we present our trust based GSP selection framework. In Section III-A we introduce the notations used in this paper. In Section III-B, we discuss how trust can be modeled using Dirichlet distribution. Then, in Sections III-C and III-D we describe how our approach works when different kinds of trust information (i.e. direct collaboration experience, indirect collaboration experience and customers' feedbacks) are available .

A. Notations

We denote the set of all service providers in the cloud computing market as $S = \{SP_1, SP_2, SP_3, ...\}$. A service provider who is going to initiate a specific instance of collaboration is called the *master service provider* (MSP). MSP will select $SP_i \in S$ from a pool of candidates, which are denoted by S_c to form a collaborative group $S_g = \{GSP_1, GSP_2, ...\}$, where GSP_i represents *guest service provider* (GSP) who participates in the collaboration. Note that $S_g \subseteq S_c \subseteq S$. If MSP has collaborated with a service provider SP_i before, MSP assessed the experience with SP_i in the collaboration by assigning a discrete quantitative rating. For instance, the rating could be in the range of $[1, 2, 3, 4, 5]$, where 1 to 5 represents *not collaborative at all*, *not collaborative*, *medium*, *collaborative* and *very collaborative*. We denote each such collaboration between service provider SP_i and SP_j by C_{sp_i,sp_j}. If multiple collaboration instances between the two providers exist, we have $\mathcal{C}_{sp_i,sp_j} = \{C^1_{sp_i,sp_j}, C^2_{sp_i,sp_j}, ...\}$. A service provider SP_i may additionally obtain feedbacks about another service provider SP_j's performance in the collaborations from other service providers who have collaborated with SP_j before. We only consider one hop trust transitivity (i.e. we do not form a long "web of trust" [9]). We denote the feedback about SP_j obtained by SP_i by F_{sp_i,sp_j}. If there exist multiple such feedbacks, we have $\mathcal{F}_{sp_i,sp_j} = \{F^1_{sp_i,sp_j}, F^2_{sp_i,sp_j}, ...\}$. Please note that for the issue of compatibility, rating representation of feedbacks should be the same to that of collaboration. For instance, feedback rating also falls in the range of $[1, 2, 3, 4, 5]$, where 1 represents *not collaborative at all* and 5 represents *very collaborative*. Table I summarizes these notations.

B. Trust model

Trust is an important abstraction used in diverse scenarios including various distributed systems. Following the works [10], [11], [19], we model trust using Dirichlet distribution, which captures a set of n observations (i.e. collaborations or feedbacks in our approach) that have k possible outcomes. We define $\vec{p} = \{p_i | 1 \leq i \leq k\}$ to denote the k-component random probability variable, and $\vec{\alpha} = \{\alpha_i | 1 \leq i \leq k\}$ to denote the number of observations corresponding to one of the possible outcomes. The Dirichlet probability density function is then given as follows:

$$f(\vec{p}|\vec{\alpha}) = \frac{\Gamma(\sum_{i=1}^{k} \alpha_i)}{\prod_{i=1}^{k} \Gamma(\alpha_i)} \prod_{i=1}^{k} p_i^{\alpha_i - 1}, \quad (1)$$

Where $p_1, ..., p_k \geq 0$, $\sum_{i=1}^{k} p_i = 1$, $\alpha_1, ..., \alpha_k > 0$ and $\sum_{i=1}^{k} \alpha_i = n$. The expected value of any of the k random variable is:

$$E(p_i|\vec{\alpha}) = \frac{\alpha_i}{\sum_{i=1}^{k} \alpha_i} \quad (2)$$

Figure 2 shows Dirichlet probability density function (PDF) when k = 3 with different $\vec{\alpha}$. The curves express the relative likelihood of the probability that outcome of the observation is one specific result out of the k possibilities, i.e., the rating of the collaboration or feedback is *collaborative* or *medium* or *not collaborative*.

We choose Dirichlet distribution because it is a conjugate prior to the multinomial distribution, which describes the probability that each independent trial result is one of the finite number k of results. Since the ratings of past collaborations and feedbacks in our approach are also any one of fixed set of values, Dirichlet distribution is thus a natural choice in our solution. Actually, the trust in this work is represented in terms of ratings (i.e. trustworthiness of a service provider is *collaborative* or *medium* or *not collaborative* or any combination of these ratings). Trust is calculated by statistically updating Dirichlet probability density function. The posteriori (i.e. the

TABLE I
Notations summary

Notations	Description
SP_i	A common cloud service provider i.
MSP	The master service provider who initializes collaboration.
GSP	The guest service provider who is invited to participate in the collaboration.
S	Set of all service providers in cloud computing market.
S_c	Set of candidates for collaboration.
S_g	Set of GSPs for a collaboration.
$\mathcal{C}_{sp_i,sp_j}$	Set of historical collaborations between SP_i and SP_j.
$\mathcal{F}_{sp_i,sp_j}$	Set of feedbacks about SP_j obtained by SP_i.

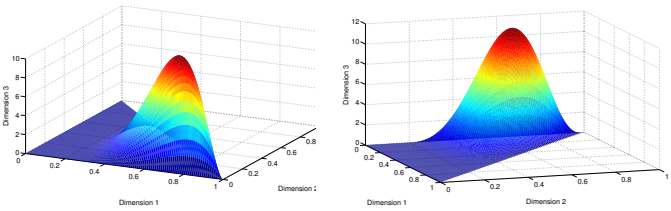

(a) $\vec{\alpha} = \{6, 2, 2\}$. (b) $\vec{\alpha} = \{2, 6, 3\}$.

Fig. 2. Dirichlet distribution probability density function.

updated) trust is derived by combining the priori (i.e. previous) trust with the new observations (i.e. new collaborations or feedbacks).

C. GSP selection with collaboration experience

We now present our trust based GSP selection framework. Our approach derives trustworthiness of GSP$_i$ according to its past behavior. Generally, if GSP$_i$ mostly behaves collaboratively, it is considered trustworthy with a high probability for the future collaboration, otherwise, it is considered untrustworthy with a high probability. We consider two ways in which a MSP may collect GSP$_i$'s historical information. One is based on direct experience, i.e., MSP's own past collaboration experience with GSP$_i$, and the other is indirect experience, i.e., feedbacks about GSP$_i$'s reputation obtained from other service providers who have interacted with GSP$_i$. We call the trust estimated using direct experience as *direct trust* and trust estimated using indirect experience as *indirect trust*.

1) Trust estimation using direct experience: Generally, direct experience is the most accurate information source for personalized trust estimation [13]. We assume that MSP has n_i past collaborations with GSP$_i$ and MSP rated these collaborations by assigning discrete quantitative ratings falling into the range $L = \{L_1, L_2, ..., L_k\}$. So the aggregate ratings for GSP$_i$ are represented as a vector $\vec{\mathcal{R}}_i = \{R_j(i)|j = 1...k\}$, where $R_j(i)$ represents number of collaborations whose ratings are L_j. Then by applying equation 1 and 2, We can derive the expected probability for each rating level in L for GSP$_i$:

$$P_j(i) = \frac{R_j(i)}{\sum_{j=1}^{k} R_j(i)} \tag{3}$$

Then the vector of expected probability for each rating level is defined by:

$$\vec{\mathcal{P}(i)} = \{P_j(i)|j = 1...k\} \tag{4}$$

Where $\sum_{j=1}^{k} P_j(i) = 1$.

Having such vector, we can estimate trustworthiness of GSP$_i$ in the future collaborations. There are two ways to predict the expected behavior (following the rating levels) using vector $\vec{\mathcal{P}(i)}$.

a) Most Likely Rating (MLR) approach: In this approach, we compare the expected probability $P_j(i)$ in $\vec{\mathcal{P}(i)}$. Since the vector can be interpreted like a multinomial probability measure to indicate how (to which level) a GSP will act in the future collaboration, we can conclude that if the expected probability is higher, the GSP is more likely to behave following the corresponding rating level. So we derive the trust of GSP$_i$ as:

$$Trust_i^{direct} = (L_j|max(\vec{\mathcal{P}(i)}) == P_j(i)) \tag{5}$$

Where the function $max(vector)$ returns maximum value of the vector.

Such approach has a drawback: if all expected probabilities in $\vec{\mathcal{P}(i)}$ are identical, trust of GSP$_i$ could be any rating level thus is meaningless. To address this issue, we propose weighted rating approach.

b) Weighted Rating (WR) approach: Instead of finding the most likely rating level, we derive trust of GSP$_i$ by taking into account all possible rating levels:

$$Trust_i^{direct} = \sum_{j=1}^{k} P_j(i)L_j \tag{6}$$

Note that in case the rating level is not quantitative, one can map it to quantitative variable. For instance, the rating level vector $\{bad, average, good\}$ could be transformed to $\{1, 2, 3\}$.

2) Trust estimation using indirect experience: When there are few or no direct experiences, MSP will also consider feedbacks about GSP$_i$ obtained from other third parties. To collect such feedbacks, MSP contacts other service providers who have interacted with GSP$_i$. We assume that the feedback represents GSP$_i$'s behavior using the same rating level discussed in the last sub section, i.e., $L = \{L_1, L_2, ..., L_k\}$. Using feedbacks $\mathcal{F}_{MSP,GSP_i}$ collected by MSP, we define the aggregate feedback ratings about GSP$_i$ as vector $\vec{R_i^F}(i) = \{R_j^F(i)|j = 1...k\}$, where $R_j^F(i)$ represents number of feedbacks whose ratings are L_j. In the same way as estimating trustworthiness using direct experience, we obtain vector of expected probability for each rating level: $\vec{\mathcal{P}^{\mathcal{F}}}(i) = \{P_j^F(i)|j = 1...k\}$. Then the indirect experience based trustworthiness $Trust_i^{indirect}$ of GSP$_i$ is derived using MLR approach or WR approach introduced before.

When third parties' feedbacks are reliable, MSP is also able to accurately predict service provider's future behavior. However, in the real world, feedbacks are not always reliable, and may even be misleading. For instance, for their own interests, service providers may provide incorrect information. To address such inaccurate feedbacks, we propose a lightweight defense mechanism.

After a collaboration is completed, MSP assesses GSP$_i$'s performance by assigning a rating level $L_m \in L$. Then MSP compares L_m with feedbacks provided by other parties. For a third party with feedback rating level L_n, we calculate the position difference between L_m and L_n:

$$D(L_m, L_n) = |Position(L_m) - Position(L_n)| \tag{7}$$

Where function $Position(L_j)$ returns position index of $L_j \in L$. MSP can estimate reliability of the corresponding feedback according to $D(L_m, L_n)$ by considering size of L. For instance, when $|L| = 5$ (i.e. there are 5 levels), MSP may consider the feedback is reliable if $D(L_m, L_n) \leqslant 1$, while when $|L| = 10$, MSP may consider the feedback is reliable if $D(L_m, L_n) \leqslant 2$.

3) Combining direct trust and indirect trust: Now we discuss how to combine direct trust and indirect trust to derive the final trust. Generally, for GSP$_i$, MSP should give more emphasis on direct trust built by personally observed behavior of GSP$_i$, while indirect trust derived from third party reports should have less importance. This is because, compared to others' opinions, local knowledge is more reliable as well as personalized. We thus need to estimate appropriate weights for reasonably aggregating the two kinds of trust information.

Although there are $k > 2$ rating levels, in order to estimate minimum number of direct collaborations to achieve certain level of prediction confidence, for simplicity, we categorize the levels into two general levels: *collaborative* and *non-collaborative*. This is reasonable, since no matter how many rating levels there are, MSP finally uses it to classify/predict new potential encounters into collaborative or non-collaborative. Please note that such binary classification does not conflict with the multiple level rating assigned by MSP because (1) multiple level rating indeed reveals more information about how the collaborator's performance is; (2) mapping to a binary rating, on other hand, is sufficient to estimate the weight for direct trust because MSP's action in the potential collaboration with a service provider is binary: collaboration or not collaboration.

Based on Chernoff bound theory [14], the minimum number of past collaborations necessary to achieve a certain level of confidence and error is calculated as follows:

$$N_{min} \geqslant -\frac{1}{2\varepsilon^2} ln(\frac{1-\gamma}{2}) \tag{8}$$

Here ε is the maximal level of error that can be accepted by MSP, and γ is the confidence measure. If the total number (denoted by N_{MSP,GSP_i}) of past collaborations between MSP and GSP$_i$ is larger than N_{min}, MSP is confident about direct trust, otherwise, it also considers indirect trust. The weights for direct and indirect trusts are determined according to N_{min}:

$$W_{MSP,GSP_i} = \begin{cases} \frac{N_{MSP,GSP_i}}{N_{min}} & \text{if } N_{MSP,GSP_i} < N_{min} \\ 1 & \text{otherwise} \end{cases}$$

So GSP$_i$'s final trust based on its historical collaboration experience is calculated by combining direct trust and indirect trust:

$$\begin{aligned} Trust_i = W_{MSP,GSP_i} \cdot Trust_i^{direct} + \\ (1 - W_{MSP,GSP_i}) \cdot Trust_i^{indirect} \end{aligned} \tag{9}$$

From equation 9 we can see that when there are more direct interactions, direct trust is more reliable thus more weight is given. Note that when N_{MSP,GSP_i} reaches N_{min}, MSP is confident about direct trust and no indirect trust is considered.

D. GSP selection without collaboration experience

In some scenarios, for GSP$_i$, MSP may lack any information regarding its past collaboration experience. For instance, it is MSP's first time to initialize a collaboration thus lacking direct experience, or GSP$_i$ never participated into any collaboration so indirect experience is unavailable. Hence, a mechanism to bootstrap trust estimates without any direct or indirect collaboration experience is needed.

We propose to derive trustworthiness of a certain GSP$_i$ for collaboration by taking into account its customers' reviews. Strictly speaking, customers' reviews should not be the criterion for accurately judging service provider's trustworthiness in the collaboration because service providers may behave differently in serving its customers and collaborating with other service providers due to its business strategy, priorities, etc. So customers' reviews may not accurately reflect the service providers' real performance in collaborations. However, we argue that customers' reviews are still useful in trust estimation because even though service providers may behave differently in delivering services to their customers and collaborating with other service providers, the behaviors in these two scenarios are expected to be positively correlated with each other, i.e., generally, the more trustworthy the service provider in delivering services to the customers, the more likely it will behave collaboratively in the interactions with other service providers (after all, the aim of collaboration is also to offer high quality services). So even though customers' reviews may not be the most accurate information source, they are useful particularly when other (more relevant) information is not available.

There are several ways to collect other service providers' customers' reviews. If possible, MSP could directly contact their customers for their reviews[1]. This is probably the best way to do so. Besides, MSP may resort to online resources such as relevant communities [15], [8], [18]. After the reviews are collected, similar to direct and indirect experience rating level representation, they are transformed to k level representation, i.e., $L = \{L_1, L_2, ..., L_k\}$ using any necessary techniques including text mining. We denote these ratings about GSP$_i$ collected by MSP by $\vec{R_i^C}(i) = \{R_j^C(i)|j = 1...k\}$, where $R_j^C(i)$ represents number of reviews whose ratings are L_j. In the same way as estimating trustworthiness using direct experience, we obtain vector of expected probability for each rating level: $\vec{\mathcal{P}^C}(i) = \{P_j^C(i)|j = 1...k\}$. Then the customers' reviews based trustworthiness $Trust_i^{customer}$ of GSP$_i$ is derived using either of the approaches introduced earlier in Section III-C1.

1) Improvements: Some improvements are necessary to make customers' reviews more accurate/relevent.

(1) Old reviews may not always be relevant for the trust estimation since the service providers may (willingly or unwillingly) change their service quality over time. So it is necessary to "forget" the reviews that are given long time ago. Two methods may be applied: (1) consider only the recent (e.g., last three or six months) reviews, and ignore all older ones. (2) use a decaying factor λ to assign weights of each review according to their ages. An example is $\lambda = b^a$, where $b \in [0, 1]$ and a is the age. Obviously, the larger the a, the less important the corresponding review is.

(2) Similar to inaccurate feedback filtering mechanism introduced in Section III-C2, if one review largely deviates from the real rating, the corresponding customer will not be requested next time (outlier elimination).

(3) In some online communities, the quality of reviews may be evaluated by other reviewers. This makes us obtain the accuracy of a review without efforts. For instance, in *Top*

[1]Incentive schemes such as rewards or any promotions/offers may be needed to elicit high quality review.

Web Hosting [8], for each review, the page shows that m out of n users found this review to be helpful. Such rate $\frac{m}{n}$ can be used as the weight of this review to indicate its accuracy.

IV. Evaluation

A. Simulation settings and methodology

Given the relative infancy of the cloud computing market, relevant real world datasets are not available, thus real trace driven simulation environments are not feasible at the moment. Following the simulation settings of previous work [7], we generate synthetic data to simulate a cloud market environment where we evaluate performance of our proposed mechanisms. We generate 100 service providers and 50 kinds of different resources. Each service provider possesses a set of resources and the set sizes of all service providers are uniformly distributed (i.e. from 1 to 10). One or more resources can be combined to produce a service. We also generate 10,000 customers who consume services offered by the service providers. Occasionally, customers may require new services, which need other resources as support. In this case, service provider may contact other service providers to form a collaborative group to share resources to fulfill new service requirements. Note that some service providers may reject the collaboration invitation due to various reasons such as limited profits, etc. If no service providers are willing to collaborate, the collaborative group will not be formed. We configure that some of the service providers may act maliciously during collaborations to degrade the quality of services. The fraction of such malicious service providers is denoted by P_m. We also denote the fraction of service providers who provide false feedbacks by P_f.

Initially, no collaboration exists, and all MSPs have to select GSPs by requesting their customers (see Section III-D). We denote the fraction of customers who provide inaccurate reviews by P_r. Since the focus of this work is in forming collaborative groups, we assume that once a collaborative group is properly formed, the service providers automatically obtain customers to sell the newly composed service.

The simulations were run for 100 synthetic time units and we configure that at each time point, each service provider may initialize collaborations according to its customers' new service requirements[2]. So, as time goes by, more and more collaborations occur among service providers. Each service provider records collaboration experience with other service providers by assigning ratings. We assume that each GSP acts in the collaboration in five performance levels, i.e., *terrible*, *bad*, *average*, *good* and *excellent*. So the rating also has such five levels. For convenience, we convert the rating levels to $\{1, 2, 3, 4, 5\}$, where 1 represents *terrible* and 5 represents *excellent*. Note that WR approach is used when calculating the trust value (see Section III-C1). When combining direct trust and indirect trust, we set N_{min} (see equation 8) as 15. In the GSP selection process, MSP first considers candidates for whom it has (direct or indirect) collaboration experiences, and only otherwise considers others.

[2]For simplicity, we assume that the collaborative group is formed within one time unit.

B. Results

1) Quality of new services: We evaluate quality of the newly provided services (through collaborations) in 50 time points (the results after time 50 are stable, thus are not shown). At each time point, each service provider may initialize a collaboration according to its customers' new requirement. After the collaboration is completed, MSP evaluates performance of each GSP by assigning ratings. We assume that such ratings reflect the real quality of performance of the corresponding GSP. So we measure quality of the collaboration (i.e. quality of the new service) using average rating of all GSPs' ratings in one collaboration.

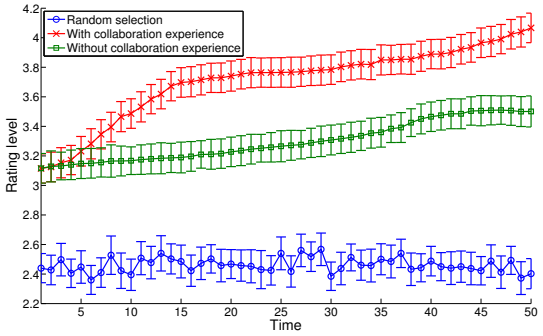

Fig. 3. Quality of new services through collaborations ($P_m = 0.3$, $P_f = 0.3$, $P_r = 0.5$).

We compare our proposed trust based approaches: (1) the approach without collaboration experience (see Section III-D) and (2) the approach with (direct or indirect) collaboration experience (see Section III-C), with the approach that only uses random selection. From Figure 3 we observe that during the whole simulation, random selection keeps the average ratings consistently at around 2.5 without any improvement. Comparatively, the approach without collaboration experience produces the average rating as high as around 3.5. Moreover, over time, more and more inaccurate review reporters are isolated so the average rating increases. This implies improved collaborations. At the very beginning, there is no direct or indirect collaboration experience so the approach with collaboration experience behaves similarly with that without experience. Then, when more collaborations are formed, more (direct or indirect) historical collaboration information about candidates is obtained, hence, our approach can more accurately predict a candidate's behavior for the future collaboration based on its past behavior. It is obvious that the approach with collaboration experience produces higher average rating than that without collaboration experience. This is because direct and indirect collaboration experience could more accurately predict candidate's future behavior in a collaboration than customers' reviews, which are the indicator of quality of service.

2) Effects of false feedbacks: Figure 4 demonstrates effectiveness of our approach in terms of eliminating false feedbacks. We vary fraction (P_f) of service providers who

provide false feedback from as low as 10% to as high as 50%. A larger value of P_f leads to higher fraction of false feedbacks, and a longer time before all such false feedback reporters are detected and isolated. However, in all cases, eventually MSPs are able to detect these service providers by comparing their feedbacks and the real ratings and then stop requesting them (i.e. the fraction of false feedbacks becomes 0). Such convergence occurs within time 30 to time 60 for the wide range of 10% to 50% of false feedback reporters.

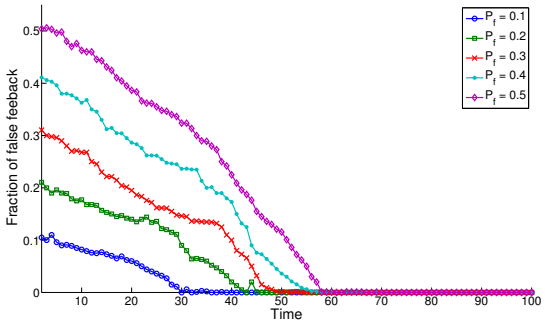

Fig. 4. Effects of false feedbacks.

3) Customer loss rate: We assume that customer loss rate of each service provider is mainly determined by the quality of its services. We note that customer loss rate may also be influenced by individual customers' service quality requirement. For instance, for storage service, business data that needs frequent processing must be strictly available all the time while archived data may tolerate some delay. In the simulations, the customers required service quality in the range of [1,5], following Normal distribution ($\mu = 3$, $\sigma = 1$).

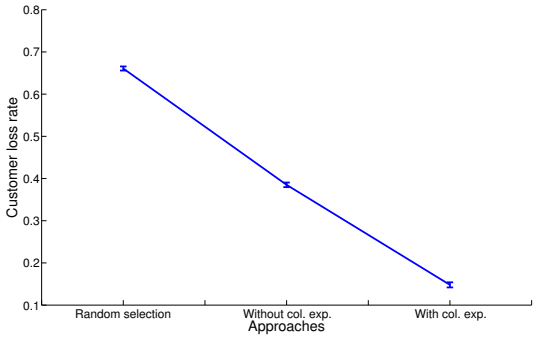

Fig. 5. Customer loss rate ($P_m = 0.3$, $P_f = 0.3$, $P_r = 0.5$).

We calculate customer loss rate using $\frac{N_{loss}}{N}$, where N_{loss} is the number of lost customers and N is the number of all customers who consume services offered through collaborations. That is to say, we only consider the customer loss due to collaboration[3]. From Figure 5 we can see that when random selection is used, the customer loss rate is as high as about 67%, while when the approach without collaboration experience is used, the loss rate drops to 38.5%. Finally, the approach with collaboration experience has the lowest customer attrition. This is because, by statistically investigating service providers' past collaboration behavior, MSP is capable of eliminating untrustworthy collaborators, thus improving quality of collaboration based services. We also observe that even though the approach with collaboration experience is used, there still exists around 14% customer loss rate due to two main reasons: (1) at the early stage of simulation, there is no sufficient historical collaboration information, in which case MSP may not find the most trustworthy collaborators; (2) to make simulation more realistic, we configure that each service provider rejects the collaboration invitation with the probability of 50% due to various reasons, which forces MSP to give up most trustworthy collaborators but select the lower ranked collaborators. Please note that the results here do not reflect customer loss rate in the real world, but just compare the plausible outcomes of different approaches.

4) Effects of malicious service providers: Malicious service providers are particularly harmful when they join a collaboration because they not only affect customers' user experience, but also damage all other collaborators' profits and reputation. Our historical information based approach is able to detect such malicious service providers by studying their past behavior, however, this process needs time to aggregate enough information. In order to make the simulations realistic, we set that malicious service providers act maliciously with a certain probability (80% in the simulation). Figure 6 shows how fraction of affected collaborations changes with time. We tried different populations of malicious service providers P_m (10%, 20% and 30%) to demonstrate effectiveness of our approach in different environments. Clearly, higher P_m affects a larger fraction of collaborations, and it takes longer time to mitigate the effects by identifying and isolating such malicious partners. Since we configure that malicious service providers do not always act maliciously, it is hard for all good providers to identify all malicious ones, so a very small fraction of collaborations continue to be affected for a long period.

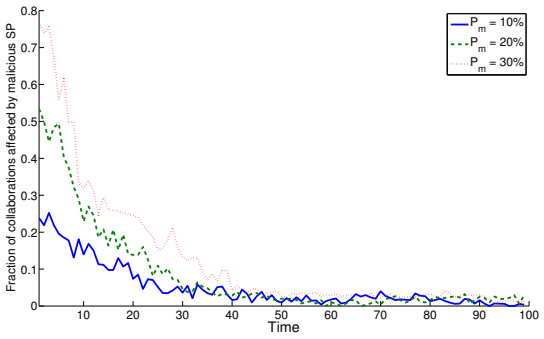

Fig. 6. Effects of malicious service providers.

V. Related Work

Collaboration among cloud providers has been proposed in [16]. The authors propose a modular and extensible cloud infrastructure in which resources and services are transparently provisioned across multiple clouds. Several technical issues

[3]We assume that when facing a new service requirement, if service provider does not collaborate, it will definitely lose the customers.

for cloud federation are identified. These include automation and fast deployment, dynamic elasticity, API compatibility, automated continuous optimization, virtualization technology independence, etc. The work *Open Cirrus* [1] takes a step forward in cloud federation by introducing an economic model. Some factors such as provider occupation and maintenance overheads are considered to evaluate a cloud provider partner.

Hassan et al. [7] propose a dynamic collaboration platform among cloud providers which can help address the interoperability and scalability issues for current cloud computing. The authors give the bidding function, which is determined by cost incurred by collaboration and expected profits to motivate the formation of a collaborative group. The authors also discuss the problem of partner selection, which is modeled as a multi-objective optimization problem, and genetic algorithms are applied to solve the same.

Goiri et al. [6] study cloud federation specifically for cloud storage by characterizing cloud federation to enhance service providers' profits. Decision equations with consideration of revenue, cost, utilization, etc. are discussed in three scenarios: (1) outsourcing resources to other federated providers; (2) renting extra resources to other federated providers and (3) shutting down unused machines to save power.

All the works mentioned above study the cloud federation/collaboration problem particularly from a technical perspective, however, they ignore the 'social' behavior of service providers, as investigated in this paper.

VI. Conclusion

Due to complexity of cloud market as well as individual service providers' resource constraints and specializations, it may be necessary to form service provider collaborations to maximize their profits and offer complex end-to-end integrated value chains required by customers. This paper presents a trust based service provider selection framework by combining disparate trust information derived from direct interactions and from (indirect) references among service providers, as well as from customer feedbacks, depending on availability of these different kinds of information: (1) When direct collaboration experience is available, selector uses Dirichlet distribution to model trust. Two approaches are introduced to predict candidate's future behavior. (2) When indirect collaboration experience is collected, similar methods are used to estimate candidate's indirect trust. Moreover, direct trust and indirect trust are combined to more reasonably derive trust taking into account certain level of error and confidence measure. (3) When neither direct nor indirect collaboration experience is there, we resort to customers' reviews of candidate's services to bootstrap the system and estimate whether a candidate is suitable for collaboration or not.

We conduct simulation experiments to evaluate performance of proposed approach by testing quality of collaboration based services, effect of false feedbacks, customer loss rate due to unsatisfactory collaborative service and effect of malicious collaborators. Simulation results show that compared to random selection, the approach without collaboration experience (i.e. rely on customers' reviews) greatly improves the service quality and lowers the customer loss rate. When (direct or indirect) historical collaboration information is considered, performance is further improved. Effects of false feedbacks and malicious service providers are also marginalized by our approach.

In the future, we plan to delve deeper into service provider collaborations to investigate how to efficiently and automatically manage resource sharing/scheduling among collaborators.

References

[1] A.I. Avetisyan, R. Campbell, I. Gupta, M.T. Heath, S.Y. Ko, G.R. Ganger, M.A. Kozuch, D. O'Hallaron, M. Kunze, T.T. Kwan, K. Lai, M. Lyons, D.S. Milojicic, Hing Yan Lee, Yeng Chai Soh, Ng Kwang Ming, J-Y. Luke, and Han Namgoong. Open cirrus: A global cloud computing testbed. *Computer*, 43(4):35–43, april 2010.

[2] Windows Azura. http://www.microsoft.com/windowsazure/windowsazure/.

[3] INFORMATION SECURITY BRIEFING 01/2010 CLOUD COMPUTING. http://www.cpni.gov.uk/docs/cloud-computing-briefing.pdf. Technical report, Centre for the Protection of National Infrastructure, 2010.

[4] Google App Engine. http://code.google.com/appengine/.

[5] GoGrid. http://www.gogrid.com/.

[6] I. Goiri, J. Guitart, and J. Torres. Characterizing cloud federation for enhancing providers' profit. In *3rd IEEE International Conference on Cloud Computing (CLOUD'10)*, pages 123–130, Miami, United States, Jul 2010. IEEE Computer Society.

[7] M.M. Hassan, B. Song, C. Yoon, H. W. Lee, and E-N. Huh. A novel market oriented dynamic collaborative cloud service infrastructure. *Services Part II, IEEE Congress on*, 0:9–16, 2009.

[8] Top Web Hosting. http://www.web-hosting-top.com.

[9] A. Jøsang, E. Gray, and M. Kinateder. Analysing topologies of transitive trust. In *Proceedings of the Workshop of Formal Aspects of Security and Trust (FAST)*, 2003.

[10] A. Jøsang and J. Haller. Dirichlet reputation systems. In *Availability, Reliability and Security, 2007. ARES 2007. The Second International Conference on*, pages 112 –119, 10-13 2007.

[11] A. Jøsang and W. Quattrociocchi. Advanced features in bayesian reputation systems. In *TrustBus '09: Proceedings of the 6th International Conference on Trust, Privacy and Security in Digital Business*, pages 105–114, Berlin, Heidelberg, 2009. Springer-Verlag.

[12] B.R. Kandukuri, V.R. Paturi, and A. Rakshit. Cloud security issues. In *Services Computing, 2009. SCC '09. IEEE International Conference on*, pages 517 –520, 21-25 2009.

[13] X. Liu, A. Datta, K. Rzadca, and E-P. Lim. Stereotrust: a group based personalized trust model. In *CIKM '09: Proceeding of the 18th ACM conference on Information and knowledge management*, pages 7–16, New York, NY, USA, 2009. ACM.

[14] L. Mui and M. Mohtashemi. A computational model of trust and reputation. In *In Proceedings of the 35th HICSS*, 2002.

[15] Cloud Hosting Reviews and Information. http://www.cloudhostingreview.com/.

[16] B. Rochwerger, D. Breitgand, E. Levy, A. Galis, K. Nagin, I. M. Llorente, R. Montero, Y. Wolfsthal, E. Elmroth, J. Caceres, M. Ben-Yehuda, W. Emmerich, and F. Galan. The reservoir model and architecture for open federated cloud computing. *IBM Journal of Research and Development*, 53(4):4:1 –4:11, july 2009.

[17] Amazon Web Services. http://aws.amazon.com/.

[18] Web Hosting Talk. http://www.webhostingtalk.com/.

[19] L. Yang and A. Cemerlic. Integrating dirichlet reputation into usage control. In *CSIIRW '09: Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research*, pages 1–4, New York, NY, USA, 2009. ACM.