Multi-Agent Service Description/Discovery Games* †

Jun Wang Les Gasser
Graduate School of Library and Information Science
University of Illinois at Urbana-Champaign
501 E. Daniel St., Champaign, IL 61820
{junwang4, gasser}@uiuc.edu

Abstract

Agent description and discovery (ADD) is a critical infrastructure for open multi-agent services. As multi-agent systems grow larger and more diverse, it becomes harder to locate useful service partner agents without efficient, effective ADD infrastructure. ADD functions rest on ontological foundations. Agent services are described and sought using terms that represent their identity, function, offerings, resource requirements, access rights, etc. In a truly distributed system, common global semantics for these terms can't be assumed, and there may be competition for using the same terms to describe different agent services; term choice matters. Descriptions clearly influence searchability, while search processes clearly favor some description regimes over others. In general, then, we need principles for describing agents and for structuring search processes based on clear principles of how description and discovery interact. This research aims at developing those principles. We study three specific questions: What equilibria exist in the space of service descriptions for populations of interacting ADD services; what rules for choosing descriptive terms yield the greatest social payoff (at equilibrium); and how close to optimum social payoff can an interacting ADD system get in principle? This research has many implications for understanding the collective properties of ADD processes and for the design of ADD systems.

1 Introduction

Agent description and discovery (ADD) is a critical infrastructure for open multi-agent services. As multi-agent systems grow larger and more diverse, it becomes harder to locate useful service partner agents without efficient, effective ADD infrastructure. In this paper, we consider rational design principles for distributed agent service description and discovery infrastructures. The task of ADD is to connect requesters of agent-based services with the agent services they need.

The ADD problem is a representation issue at its heart. Agents themselves are comprised of content and functions, but the only medium agents have to describe the services they offer or seek is an information medium—a language. Somehow, agents' service content/function must be transformed into a linguistic representation. Descriptive terms must be chosen and assigned as descriptions of agent services; requesters, too need to choose and assign terms to represent content/functions they seek. In a distributed system, there is no a-priori reason for agents to be consistent in choosing descriptive terms. In fact, there are many situations in which agents have good reasons for assigning random or irrelevant terms (e.g., to attract as many prospects as possible), and competition may develop to use the same terms for different services or *vice-versa* (synonymy/homonymy). Unfortunately, without consistent, meaningful associations between terms and content/functions, ADD is inefficient at best.

The research is part of our long-term efforts on dynamic distributed semantics in multi-agent systems. The basic issue of this paper is to understand the conditions under which rationally self-interested (payoff-maximizing) agents can collectively and efficiently describe their own services and services they seek. Specifically, we investigate equilibria and optimal social payoffs for service description games in which agents' payoffs derive from the effectiveness of the terms they choose for the services they offer or seek. Here we develop mathematical principles of structure and performance in such scenarios, which will serve as foundations for specific term-choice or term-modification algorithms.

We would like to know precisely how much performance is lost when inefficient strategies (like term competition, irrelevance, synonmy, etc.) are chosen by service providers, and why. We measure this loss as the ratio of the optimal effectiveness of socially coordinated descriptions to the ef-

*Work supported by NSF Grant IIS-0340996.
† STUDENT PAPER. Paper ID: 255.

fectiveness of purely self-interested descriptions, at Nash equilibrium. We call this ratio the *cost of selfish description*, because the more selfish (i.e. less coordinated) the description strategy, the more energy agents will spend weeding out useless service prospects. (This idea was inspired by [3] and [9].) We are also interested in how the cost of selfish description changes under alternative service description protocols.

By formalizing and analyzing the description competition among service providers as a non-cooperative game, we get the following result. Under some assumptions, the cost of selfish description can be as high as the number of terms that a service provider is allowed to use to describe its service. Surprisingly, in our symmetric single-mode distribution model, the lowest cost is typically achieved when the ADD system only allows agents to use one description term (the *mode* or *primary term*) for each service.

2 Preliminaries

In this section we present definitions and notations needed in the paper, including description language model, system protocols, game model, the cost of selfish description.

2.1 Description language model

For a given service, providers and requesters may choose descriptive terms stochastically. In general, language modeling provides an approach to characterize this kind of stochasticity by using a probability distribution to capture the regularities of language generation and use [1]. Similarly, in our context, we use a *description language model* or briefly *description model* to represent the statistical regularities of the selection of descriptive terms.

Let the universe of all available terms be $T = \{t_1, ..., t_m\}$. Then, given a service, its description model is a conditional probability distribution on T. Formally, we can denote by $P(t|d)$ the probability of using term t to describe or request the service d. We assume that for a specific service d, service-requesters and service-providers use this model $P(t|d)$ to select their descriptive terms or query terms. [1]

2.2 Protocols of service discovery

We now specify two aspects of the service description protocol for service providers: (a) What restrictions on the description of a service does the ADD system impose? (b) How does the ADD system rank multiple services when all of them are matched with a query?

Service description requirements. An ADD system places two requirements on any service description. One is that any description must be a term set, that is, a subset of T. The second requirement is that the system limits the cardinality of an acceptable term set. Later we will show how this restriction on the number of terms affects the cost of selfish description.

Ranking mechanism. When a service requester wants to locate a partner agent service, it submits a query term to the ADD system (the query term is generated according to the description model mentioned above, and recall that it may be a single word or a phrase). The system will return all services (precisely speaking, the addresses of the services) that match the query term. When there are several services matched with the query, the ADD system needs a decision procedure to specify the *rank* or the order of the matched services. Assume the requester *visits* the returned service to judge its actual quality or acceptability, in order of this rank. In this way, higher-ranked services that meet quality criteria will be more likely to be actually used, than the lower-ranked ones.

Since we have no information beyond descriptions to help rank matched services, we again take the stochastic approach. The rank of a service is represented as a random variable, denoted by R. When there are k matched services, we define the distribution of R as follows: for each $r = 1, 2, \cdots, k$,

$$\mathrm{P}(R = r|d, q) = \begin{cases} \frac{1}{k} & \text{if service } d \text{ matches with query } q, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

2.3 Visit possibility for a matched service

The visit possibility for a service depends on its rank. The higher the rank, the more the chance of a visit. In order to represent this relationship, we need to introduce a parameter, called *discount rate* $\delta \in [0, 1)$, as the rate with which the rank reduces the possibility of a visit. Suppose chance of visiting the highest-ranked service (whose $rank = 1$) is 1. Then the chance for visiting the second highest-ranked service is δ. Obviously, the nth-ranked service has the visit chance δ^{n-1}.

Thus, given a term t, suppose service d is one of k services that match with t, then the *visit possibility* that d gets visited is:

$$Q(d|t) = \sum_{r=1}^{k} \mathrm{P}(R = r|d, t)\delta^{r-1} = \frac{1 - \delta^k}{k(1 - \delta)} \,. \quad (2)$$

[1] This stochastic approach was also used in a probabilitic corpus model proposed by Papadimitriou and his colleagues, that aimed at interpreting the success of latent semantic indexing model in information retrieval[7].

Note that the *visit possibility* is not in general a probability (it is probability when $\delta = 0$) so we represent it using notation Q instead of P.

We also remark that a more accurate model of the visit possibility might consider the rank relation between a matched service and the desired service (the one that a requester intends to find). If the desired service has been found, there will be no chance of using the remaining services which are ranked after this desired service. For convenience, we will not treat this case in this paper.

2.4 Game model

We consider a game with n service providers: $\{1, 2, ...n\}$. Denote the service of provider i by d_i with description model $P(t|d_i)$. We define the strategy spaces and the payoff functions for them as follows.

Strategy space. Let the terms that provider i chooses to describe his service be $s_i \subset T$. Clearly, s_i can be represented as a binary vector: $s_i = (s_{i1}, \cdots, s_{im})$. $s_{ij} = 1$ indicates the term t_j is chosen as a descriptor for his service d_i. The strategy space of a provider is: $\{0, 1\}^m$. All providers have the same strategy space. Note that we only consider *pure* strategies.

Payoff function. Now we define the payoff functions of the game. Let the strategy of service provider i be $s_i = (s_{i1}, \cdots, s_{im})$. Let n_j be the number of service providers that choose the term t_j as a descriptor. Since $s_{ij} = 0$ when provider i doesn't use term t_j and $s_{ij} = 1$ when he does use it, we have

$$n_j = \sum_i s_{ij}. \tag{3}$$

From Eq. (2), we have the visit possibility of service d_i:

$$Q(d_i|t_j) = \frac{1 - \delta^{n_j}}{n_j(1 - \delta)}.$$

Since the distribution $P(t|d_i)$ is the description model of service d_i, requesters will discover d_i using term t_j with the probability $\mathrm{P}(t_j|d_i)$. In addition, let $P(d_i)$ denote the probability that service d_i will be sought (measured over all queries of all requesters). For a popular topic or service, its $P(d_i)$ might be higher.

Now, given the strategy combination of all providers being $(s_1, \cdots, s_n)$, we define the payoff of provider i as

$$\mu_i(s_1, ...s_n) = \mathrm{P}(d_i) \sum_j \mathrm{P}(t_j|d_i)\, \mathrm{Q}(d_i|t_j). \tag{4}$$

Clearly, this payoff definition reflects the chance that service d_i can be successfully discovered by requesters who intend to find d_i.

Nash equilibrium as game solution. Nash equilibrium is the most widely used solution to multi-player games[6]. A strategy combination $s = (s_1, \cdots, s_i, \cdots, s_n)$ is a Nash equilibrium if for any player i, the payoff of this player will not increase if he unilaterally changes his strategy. Formally speaking, for all i, and for all s'_i, we have

$$\mu_i(s_1, \cdots, s_i, \cdots, s_n) \geq \mu_i(s_1, \cdots, s'_i, \cdots, s_n).$$

2.5 Request effectiveness and social payoff

One of the design goals of an agent description and discovery system is to maximize the effectiveness of service requests. We give the definition of this effectiveness here, and show that in service description games, request effectiveness turns out to be equal to the *social payoff* of service providers, i.e., the sum of the individual payoffs of all providers in the game.

Request effectiveness. Request effectiveness is defined as the expected probability that a requester can find a desired service, computed as follows

$$\nu = \sum_i \mathrm{P}(d_i) \sum_j \mathrm{P}(t_j|d_i)\, \mathrm{Q}(d_i|t_j). \tag{5}$$

If we define two matrices $\boldsymbol{P}$ and $\boldsymbol{Q}$ as: $\boldsymbol{P}_{ij} = P(t_j|d_i)$ and $\boldsymbol{Q}_{ji} = Q(d_i|t_j)$; and we further assume that $P(d_i)$ is uniform, that is, $P(d_i) = \frac{1}{n}$. Then, ν can be rewritten as:

$$\nu = \frac{1}{n} TR(\boldsymbol{PQ})$$

where TR is the *trace* of a matrix, i.e., the sum of the diagonal entries. We will return to this representation later when we discuss the concept of optimal social payoff.

Social payoff. The social payoff is defined as the sum of payoffs of all providers in a game, shown as follows

$$\mu = \sum_i \mu_i. \tag{6}$$

Clearly, in our model the social payoff of the game is the same as the request effectiveness. That is, $\mu = \nu$.

Remark. Our service description game is a non-cooperative game where providers compete against each other to use descriptive terms. This is a type of *congestion game*. Typically, the solution (Nash equilibrium) of a non-cooperative game is not always the overall optimum, with the Prisoner's Dilemma being the best-known example. One purpose of this paper is therefore to analyze the cost of selfish description for the game and find how the cost changes under different system protocols so that we can adopt the best protocol.

2.6 Cost of selfish description

Let μ_{opt} denote the optimal social payoff, and μ_{worst} denote the social payoff at the worst possible Nash equilibrium. Then the *cost of selfish description* is defined as

$$\rho = \frac{\mu_{opt}}{\mu_{worst}}.$$

Note that our definition is superficially different from the one proposed in [3].

3 Nash Equilibria of Symmetric Single-Mode Service Description Games

From this point, for simplicity, we will focus only on symmetric single-mode service description games, leaving more general cases for future investigation.

3.1 Symmetric single-mode games

In this game, the number of services is the same as the number of terms. That is, $m = n$. Also, the prior probability that requesters seek service d_i is $P(d_i) = \frac{1}{n}$; this is uniformly distributed. For a given service, the probability that each term is used to describe that service is not uniformly distributed. Each service has one primary (or "mode") term for which the probability of use is $1-\epsilon$, and the other $n-1$ terms have the probability $\frac{\epsilon}{n-1}$ of being used as a descriptor.

We assume that each service has a *distinct* primary term. Without loss of generality, we state this formally as follows. Let t_i be the primary term of service d_i. Then $P(t_i|d_i) = 1-\epsilon$ and $P(t_j|d_i) = \frac{\epsilon}{n-1}$, for $j \neq i$. (In effect we are letting all the mode terms appear on the diagonal when this is written as a matrix, but they could of course be permuted.) See Table 1 for illustration.

	t_1	t_2	$\cdots$	t_n
d_1	$1-\epsilon$	$\frac{\epsilon}{n-1}$	$\cdots$	$\frac{\epsilon}{n-1}$
d_2	$\frac{\epsilon}{n-1}$	$1-\epsilon$	$\cdots$	$\frac{\epsilon}{n-1}$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
d_n	$\frac{\epsilon}{n-1}$	$\frac{\epsilon}{n-1}$	$\cdots$	$1-\epsilon$

Table 1. A symmetric single-mode description model, where t_i is the primary term (or mode) of d_i.

Note that for t_i to be the primary term of d_i, we require at least $1-\epsilon > \frac{\epsilon}{n-1}$, or equivalently $1-\epsilon > \frac{1}{n}$. For reasons which will be given later, it still makes sense to make a *primary term probability assumption*, namely that the probability of the primary term being be used to describe the service is at least slightly larger than the probability of the other terms; formally:

$$1-\epsilon > \frac{2}{n+1}.$$

3.2 Properties of Nash equilibria

Here we will first show some properties for Nash equilibrium in the description games. Then we will show (Theorem 3.5) that the social payoff in description games at Nash equilibrium is unique when the *primary term probability assumption* given above holds. Recall that our ADD protocol requires service providers to use a limited number of description terms, with the limit denoted by k.

Lemma 3.1. *A strategy at Nash equilibrium must have exactly k terms chosen.*

Proof. According to the system protocol, a provider can choose no more than k description terms. Suppose some provider chooses a term set of size less than k for its service. Then, if this provider adds to his strategy any term that is not currently in his term set his payoff will increase. This contradicts the principle of Nash equilibrium. So for a Nash equilibrium, a provider must use exactly k terms. □

Lemma 3.2. *Let n_j be the number of services that use the term t_j as one of their descriptors. Then in any Nash equilibrium, for $j = 1, \cdots, n$, we have $n_j = k$.*

Proof. Lemma 3.1 tells us that for n services, in a Nash equilibrium, the total number of used terms is: nk. Suppose there exists a term, say, t_1 with $n_1 \neq k$. Without loss of generality, let $n_1 > k$, then we must have another term, say, t_2 with $n_2 < k$. Since $n_1 > k > n_2$, this implies

$$n_1 > n_2 + 1.$$

From $n_1 > k \geq 1$ and the fact that each service has a distinct primary term, we have that among the services that use the term t_1, there exists a service d_i whose primary term is different from t_1. We will show that the provider i of service d_i can increase his payoff by changing from term t_1 to t_2 even when t_2 is a non-primary term for d_i. When t_2 is the primary term of d_i, clearly the increased payoff can be more.

Let π_1 denote the payoff of provider i when he uses the term t_1, and π_2 the payoff after he changes from t_1 to t_2. Note that after provider i changes to term t_2, the number of services that use term t_2 will be n_2+1. From the definition of payoff, we have

$$\pi_2 - \pi_1 = \frac{\epsilon}{n(n-1)(1-\delta)}\Big(\frac{1-\delta^{n_2+1}}{n_2+1} - \frac{1-\delta^{n_1}}{n_1}\Big).$$

It can be shown that the function $f(x) = \frac{1-\delta^x}{x}$ is decreasing for $x > 0$ and $0 \le \delta < 1$. Since $n_1 > n_2 + 1$, we have $\pi_2 - \pi_1 > 0$. Therefore, at a Nash equilibrium, the number of services that use any term must equal to k, or the requirement of Nash equilibrium is contradicted. □

Lemma 3.3. *The payoff of a provider at Nash equilibrium is* $\frac{1-\delta^k}{nk(1-\delta)}(1-\epsilon\frac{n-k}{n-1})$ *when it uses its primary term, or* $\frac{\epsilon(1-\delta^k)}{n(n-1)(1-\delta)}$ *when he does not.*

Proof. According to the definition of payoff (Eq. 4) and $P(d_i) = \frac{1}{n}$, for provider i and a strategy combination s, we have

$$\mu_i(s) = \frac{1}{n}\sum_j \mathrm{Q}(d_i|t_j)\ \mathrm{P}(t_j|d_i)\ .$$

From Lemma 3.2, we know that at a Nash equilibrium, the number of services that use term t_j is k. From the perspective of requesters or service discovery, the number of services that match with t_j is k. Therefore, from the definition of service visit possibility, we have

$$Q(d_i|t_j) = \frac{1-\delta^k}{k(1-\delta)}.$$

Therefore,

$$\mu_i(s) = \frac{1-\delta^k}{nk(1-\delta)}\sum_j \mathrm{P}(t_j|d_i).$$

Clearly, in symmetric single-mode games, a provider's strategy is either to use the primary term and some $k-1$ non-primary terms or to use some k non-primary terms. We call the former the *primary-term strategy*, and the latter the *non-primary-term strategy*. According to the definition of the symmetric single-mode game, we have $P(t_i|d_i) = 1-\epsilon$ when t_j is the primary term describing the service d_i, and $P(t_i|d_i) = \frac{\epsilon}{n-1}$ otherwise.

For the primary-term strategy we have

$$\sum_j P(t_j|d_i) = (1-\epsilon) + (k-1)\frac{\epsilon}{n-1} = 1-\epsilon\frac{n-k}{n-1}.$$

For the non-primary-term strategy we have:

$$\sum_j P(t_j|d_i) = k\frac{\epsilon}{n-1}.$$

In conclusion, the payoff of a provider at Nash equilibrium is

$$\mu_i(s) = \begin{cases} \frac{1-\delta^k}{nk(1-\delta)}(1-\epsilon\frac{n-k}{n-1}) & \text{if primary term is used,} \\ \frac{\epsilon(1-\delta^k)}{n(n-1)(1-\delta)} & \text{otherwise.} \end{cases}$$

□

Lemma 3.4. *When* $1-\epsilon > \frac{2}{n+1}$*, a strategy at Nash equilibrium must be a primary-term strategy.*

Proof. Suppose there exists a Nash equilibrium non-primary-term strategy for provider i. Then we want to show that provider i can increase its payoff by changing to a primary-term strategy. The payoff of a non-primary term strategy (denoted by π_1) is:

$$\pi_1 = \frac{\epsilon(1-\delta^k)}{n(n-1)(1-\delta)}.$$

When the service provider converts some non-primary term into the primary term, its new payoff is:

$$\pi_2 = \frac{1}{n}\Big(\frac{(1-\epsilon)(1-\delta^{k+1})}{(k+1)(1-\delta)} + \frac{\epsilon(k-1)(1-\delta^k)}{k(n-1)(1-\delta)}\Big).$$

Thus the payoff changes by

$$\pi_2 - \pi_1 = \frac{1-\delta^k}{kn(n-1)(1-\delta)} \times\Bigg((1-\epsilon)\Big(\frac{k+1+k(n-1)\frac{1-\delta^{k+1}}{1-\delta^k}}{k+1} - 1\Big)\Bigg).$$

Since $0 \le \delta < 1$, we know that expression $\frac{1-\delta^{k+1}}{1-\delta^k}$ in the equation above is greater than or equal to 1. Then we have:

$$\begin{aligned} 1-\epsilon &> \frac{2}{n+1} \\ &\ge \frac{k+1}{kn+1} \\ &= \frac{k+1}{k+1+k(n-1)} \\ &\ge \frac{k+1}{k+1+k(n-1)\frac{1-\delta^{k+1}}{1-\delta^k}}. \end{aligned}$$

It follows that $\pi_2 - \pi_1 > 0$, and this contradicts the conditions of Nash equilibrium. Thus at Nash equilibrium each provider must use a primary-term strategy. □

With the above lemmas, now it is obvious that the following theorem is true:

Theorem 3.5. *When* $1-\epsilon > \frac{2}{n+1}$*, the social payoff in the description games is unique:* $\frac{1-\delta^k}{k(1-\delta)}(1-\epsilon\frac{n-k}{n-1})$.

4 Worst-Case Cost of Selfish Description

In this section we give the details that demonstrate why the worst-case cost of selfish description is a function of k (the limit to the number of descriptive terms allowed) and the discount rate δ, namely $\frac{k(1-\delta)}{1-\delta^k}$. This worst-case cost is equal to k when $\delta = 0$. This is interesting because it means that limiting the number of terms an agent can use to describe or request a service can also limit the degree of ambiguity in the semantics of descriptions. This clearly can have the effect of improving the efficiency and scalability of ADD systems.

4.1 Optimal social payoff

Since we want to know how high the cost of selfish description at Nash equilibrium can be at worst, we need to compute the *optimal* social payoff. Though the social optimum may depend on the setting of parameters ϵ and δ, we will show that, with probability 1, the social optimum can be achieved when every provider describes his service only using his primary term, as long as the number of services or providers is large enough.

Let k_j be the number of services that use term t_j as descriptors. Then we can build a matrix $\boldsymbol{Q}$ by $Q_{ji} = Q(d_i|t_j)$ that may look like this (rows indicate the terms and columns the services):

$$\mathbf{Q} = \begin{pmatrix} \frac{1-\delta^{k_1}}{k_1(1-\delta)} \cdot s_{11} & \frac{1-\delta^{k_1}}{k_1(1-\delta)} \cdot s_{21} & \cdots \\ \frac{1-\delta^{k_2}}{k_2(1-\delta)} \cdot s_{12} & \frac{1-\delta^{k_2}}{k_2(1-\delta)} \cdot s_{22} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix},$$

where the binary vector $(s_{i1}, s_{i2}, \cdots, s_{in})$ is the strategy of provider i. Clearly, $\sum_i s_{ij} = k_j$, and the sum of the entries in the jth row is $\sum_i \frac{1-\delta^{k_j}}{k_j(1-\delta)} = \frac{1-\delta^{k_j}}{1-\delta}$.

Now, let $\boldsymbol{P}$ be the matrix $\boldsymbol{P}_{ij} = P(t_j|d_i)$ as defined in Table 1. Then the social or total payoff is:

$$\begin{aligned} \mu_{total} &= \frac{1}{n} TR(\boldsymbol{PQ}) \\ &= \frac{1}{n} \sum_j \frac{1-\delta^{k_j}}{k_j(1-\delta)} \Big(s_{jj}(1-\epsilon) + \sum_{i \neq j} s_{ij} \frac{\epsilon}{n-1} \Big) \end{aligned}$$

Given this definition, our task is to find the *maximum* social payoff, that is:

$$\mu_{opt} = \max_{s_{ij} \in \{0,1\},\, i,j=1,\cdots,n} \mu_{total}$$

subject to $1 \leq k_j \leq k$ and $k_j = \sum_i s_{ij}$.

Since $1 - \epsilon > \frac{\epsilon}{n-1}$, for μ_{total} to be maximized, we need to set $s_{jj} = 1$. In addition, given the symmetricity of the service description game model, we can reduce the above optimization task to this one:

$$\max_{1 \leq k^* \leq k} \frac{1-\delta^{k^*}}{k^*(1-\delta)} \Big((1-\epsilon) + \frac{k^*-1}{n-1}\epsilon \Big),$$

which is equivalent to

$$\max_{1 \leq k^* \leq k} \frac{1-\delta^{k^*}}{k^*(1-\delta)} \big(1 - \epsilon \frac{n-k^*}{n-1} \big).$$

Next, we will show when n is large enough, $k^* = 1$ is the solution for the above function to be optimal. Clearly, when $k^* = 1$, the social payoff is: $1 - \epsilon$. But before we give the proof, we give an example with $n = 3$ providers to show that the social optimum can be reached at $k^* = 1$, $k^* = 2$, and $k^* = 3$ respectively, given different values of δ.

Let $n = 3$ and $\epsilon = \frac{2}{5}$. Then, when $0 \leq \delta \leq \frac{1}{2}$, the social optimum is achieved at $k^* = 1$. When $\frac{1}{2} \leq \delta \leq \frac{1+2\sqrt{5}}{10}$, the social optimum is achieved at $k^* = 2$. When $\frac{1+2\sqrt{5}}{10} \leq \delta < 1$, the social optimum is achieved at $k^* = 3$. (See Fig. 1 for illustration.)

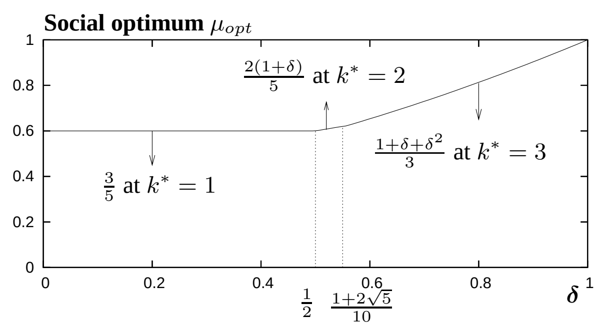

Figure 1. Given different δ, the social optimum is reached at different k^*.

Lemma 4.1. *Let* $f(k) = \frac{1-\delta^k}{k(1-\delta)}(1 - \epsilon\frac{n-k}{n-1})$ *with* $0 \leq \delta < 1$ *and* $0 \leq \epsilon \leq 1$. *If* $f(1) \geq f(2)$, *then for* $k \geq 2$ *we have* $f(1) \geq f(k)$.

Proof. From $f(1) \geq f(2)$, we have

$$\epsilon \leq \frac{(n-1)(1-\delta)}{n - n\delta + 2\delta}.$$

Now we can prove that $f(1) \geq f(k)$ by computing

$$\begin{aligned} f(1) - f(k) &= (1-\epsilon) - \frac{1-\delta^k}{k(1-\delta)} \Big(1 - \epsilon \frac{n-k}{n-1} \Big) \\ &= 1 - \frac{1-\delta^k}{k(1-\delta)} - \epsilon \Big(1 - \frac{n-k}{n-1} \cdot \frac{1-\delta^k}{k(1-\delta)} \Big) \\ &\geq \frac{2\delta}{n - n\delta + 2\delta} \Big(\frac{1+\delta^{k-1}}{2} - \frac{1+\delta+\cdots+\delta^{k-1}}{k} \Big) \\ &\geq 0. \end{aligned} \tag{7}$$

The inequality (7) comes from

$$\frac{1+\delta^{k-1}}{2} \geq \frac{1+\delta+\cdots+\delta^{k-1}}{k}, \text{for } k \geq 2 \text{ and } \delta \geq 0,$$

which can be proved using mathematical induction and the arithmetic mean-geometric mean inequality.

Therefore, we have proved that $f(1) \geq f(k)$ for $k \geq 2$. □

Lemma 4.2. *Given function $f(k)$ defined above in the Lemma 4.1, suppose δ is uniformly distributed in $[0, 1)$. Then, with probability 1,*

$$f(1) \geq f(2) \quad \textit{as} \quad n \to \infty^+.$$

Proof. Since $f(1) \geq f(2)$ is equivalent to $\epsilon \leq \frac{(n-1)(1-\delta)}{n-n\delta+2\delta}$, and the probability density of δ is 1, we have

$$\begin{aligned} Pr\Big(f(1) \geq f(2)\Big) &= Pr\Big(\epsilon \leq \frac{(n-1)(1-\delta)}{n-n\delta+2\delta}\Big) \\ &= \int_0^1 \frac{(n-1)(1-\delta)}{n-n\delta+2\delta}\mathrm{d}\delta \\ &= 1 - \frac{2\ln\frac{n}{2}-1}{n-2} - \frac{2\ln\frac{n}{2}}{(n-2)^2}. \end{aligned}$$

Therefore,

$$\lim_{n\to\infty^+} Pr\Big(f(1) \geq f(2)\Big) = 1,$$

completing the proof. □

Now, with the above preparation, we have the following theorem:

Theorem 4.3. *With probability 1, the social optimum for our service description game is $1-\epsilon$, as n is large enough.*

4.2 Worst-case cost of selfish description

As the number of service providers n is large enough, we have proved that the social optimum is $\mu_{opt} = 1 - \epsilon$, which is achieved when every service provider only chooses his primary term to describe his service. We also know from the last section that, when $\epsilon > \frac{2}{n+1}$, the social payoff at Nash equilibrium is unique: $\mu_{nash} = \frac{1-\delta^k}{k(1-\delta)}(1 - \epsilon\frac{n-k}{n-1})$, where k is the largest number of terms that a service provider is allowed to choose for describing his service. Since the social payoff is unique, we have $\mu_{worst} = \mu_{nash}$.

Therefore, according to the definition of the cost of selfish description, we have

$$\begin{aligned} \rho &= \frac{\mu_{opt}}{\mu_{worst}} \\ &= \frac{1-\epsilon}{\frac{1-\delta^k}{k(1-\delta)}(1-\epsilon\frac{n-k}{n-1})} \qquad (8) \\ &\leq \frac{k(1-\delta)}{1-\delta^k} \qquad (9) \\ &\leq k. \qquad (10) \end{aligned}$$

The upper bound $\frac{k(1-\delta)}{1-\delta^k}$ in the above inequality (9) is tight and is achieved when $\epsilon \to 0$. And the upper bound k in the inequality (10) is achieved when $\delta = 0$.

Therefore, it turns out that, worst-case cost of selfish description is $\frac{k(1-\delta)}{1-\delta^k}$. The worst-case cost is equal to k when $\delta = 0$.

4.3 Discussions

From the above analysis, we see that the cost of selfish description can be very large if there is no restriction on the number of description terms. That is, we will have the worst case at $k = n$. Substituting $k = n$ to the above Eq. (8), we have $\rho = \frac{k(1-\epsilon)(1-\delta)}{1-\delta^n} \approx k(1-\epsilon)(1-\delta)$. As an example, consider setting $\delta = \frac{1}{2}$ and $1 - \epsilon = \frac{1}{2}$, for which the cost will be $\frac{n}{4}$. This makes the effectiveness of service discovery extremely low. In fact, during the early period of web search engine use, information seekers often suffered from this kind of selfish description when page providers notoriously included many irrelevant terms to gain hits.

Since the cost of selfish description depends on k, we can get the best cost $\rho = 1$ when setting $k = 1$. That is, every provider can only choose one description term. At first this looks very impractical. But notice that in our symmetric single-mode model, we assumed that one service has only one primary term (single mode distribution). This assumption may play the key role in establishing that $k^* = 1$ guarantees the social optimum. If we relax this assumption and consider an extension which takes multiple modes instead of one, we hypothesize that the optimal setting of k might be the number of modes.

5 Related Work

Matchmaking. Middle-agents, matchmaking and service descriptions are active and exciting study areas. Existing work generally doesn't yet cover collective semantics of description languages, possibilities for dynamic descriptions, description equilibria, or social costs and benefits of description choices. See Sycara et al.'s numerous papers on middle-agents, matchmakers/brokers, and LARKS, e.g. [10].

Congestion games. The game model described in this paper is closely related to congestion games (CG), first proposed by Rosenthal[8] thirty years ago, and more recently studied in game theory and computer science [5, 9, 12]. A typical CG model is a road network, where the time taken (negative payoff) by a person on a road depends on traffic on that road—more traffic, more delay. It has been shown that any CG has at least one pure-strategy Nash equilibrium [8, 5]. Most CG work assumes that the payoff of a player only depends on the number of players taking the same strategy (this helps scalability of the analysis). There is little work on the case where payoffs do depend on specific providers (called subjective CGs). Milchtaich showed [4] that a special case of subjective CGs—called *subjective parallel edge congestion games*, which roughly means that each provider chooses only one road—possess at least one pure-strategy Nash equilibrium.

Service description games in our work can be viewed as subjective congestion games. In description games, providers compete for query terms, but usually providers have their own (subjective) preferences on different terms. We remark that description games go beyond the parallel edge assumption. We yet don't know if we have pure-strategy Nash equilibrium for general description games, but we have shown that for our special case of symmetric single-mode games, Nash equilibrium does exist.

Worst-case coordination ratio. Our work is part of a growing literature on measuring how selfish behaviors affect efficiencies of whole systems, which originated in Koutsoupias and Papadimitriou's work on worst-case equilibria [3], followed by extensive studies on selfish routing[9] and other work, e.g. [12].

Others. Interesting related work modeled distributed web search as a stochastic game[2]. The motivation was that topic-specific search engines compete for queries by choosing which documents or topics to index, and the research developed reinforcement learning approaches for individual search engines to manage index contents. This work can be seen as an application of learning Nash equilibrium in non-subjective congestion games, though the authors did not mention congestion games.

Readers familar with language games might find that our payoff function resembles an extension (with a consideration of discount rate) of the one used in [11]. However, language games research mainly examines evolution of social payoff, not individual payoff, using a population dynamics approach. In contrast, we assume that the language model of a service, represented as a probability distribution on the term set T, is fixed rather than dynamically evolving. A promising direction for our work is to consider dynamic properties of language models.

6 Conclusion and Future Work

We proposed a new game called the *service description/discovery game*, which models how decentralized service providers compete for requesters' interest in an agent description and discovery (ADD) system. We are especially interested in a specific game with symmetric single-mode distribution. We have proved several properties of Nash equilibrium of this game, and determined the total payoff of all providers at Nash equilibrium. We also analyzed the globally optimal social payoff, and compute the worst-case cost of selfish descriptions. Our results show that the best cost is achieved when the protocol of the ADD system limits all providers' descriptions to one term.

A more accurate model (left for future research) might be a more hierarchical one that provides several categories for describing a service, and a distribution on the terms defined in each category. Other assumptions in the symmetric single-mode distribution model could be relaxed, including symmetricity and the single-mode nature of the model. We also intend to explore the case in which service requests can be contain more than one term, making the ranking mechanism much more complex because of the introduction of partial matching. In this case, we need to use more knowledge from language modelling approach—requesters generate requests according to the description language model.

References

[1] R. A. Cole, J. Mariani, H. Uszkoreit, A. Zaenen, and V. Zue, editors. *Survey of the State of the Art in Human Language Technology*. Cambridge University Press, Cambridge, UK, 1997.

[2] R. Khoussainov and N. Kushmerick. Distributed web search as a stochastic game. In J. Callan, F. Crestani, and M. Sanderson, editors, *Proceedings of the SIGIR 2003 Workshop on Distributed Information Retrieval*, 2003.

[3] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.

[4] I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13:111–124, 1996.

[5] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.

[6] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[7] C. H. Papadimitriou, P. Raghavan, H. Tamaki, and S. Vempala. Latent semantic indexing: A probabilistic analysis. In *Proceedings of the Seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems - PODS 1998*, pages 159–168, 1998.

[8] R. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[9] T. Roughgarden. *Selfish Routing*. PhD thesis, Cornell University, 2002.

[10] K. Sycara, J. Lu, M. Klusch, and S. Widoff. Matchmaking among heterogeneous agents on the internet. In *Proceedings AAAI Spring Symposium on Intelligent Agents in Cyberspace*, 1999.

[11] P. E. Trapa and M. A. Nowak. Nash equilibria for an evolutionary language game. *Journal of Mathematical Biology*, 41(2):172–188, 2000.

[12] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *IEEE 43rd Symposium on Foundations of Computer Science (FOCS 2002)*, pages 416–425, 2002.