Agents' Strategies for the Dual Parallel Search in Partnership Formation Applications ⋆

David Sarne[1] and Sarit Kraus[1,2]

[1] Department of Computer Science, Bar-Ilan University, Ramat-Gan, 52900 Israel
[2] Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742

Abstract. In many two-sided search applications, autonomous agents can enjoy the advantage of parallel search, powered by their ability to handle an enormous amount of information, in a short time, and the capability to maintain interaction with several other agents in parallel. The adoption of the new technique by an agent suggests a reduction in the average cost per interaction with other agents, resulting in an improved overall utility. Nevertheless, when all agents use parallel search in Multi-Agent Systems (MAS) applications, the analysis must take into consideration mainly equilibrium dynamics which shape their strategies. In this paper we introduce a dual parallel two-sided search model and supply the appropriate analysis for finding the agents' equilibrium strategies. As a framework application for our analysis we suggest and utilize the classic voice communication partnerships application in an electronic marketplace. By identifying the specific characteristics of the equilibria, we manage to supply efficient means for the agents to calculate their distributed equilibrium strategies. We show that in some cases equilibrium dynamics might eventually drive the agents into strategies by which all of them end up with a smaller expected utility. Nonetheless, in most environments the technique has many advantages in improving the agents expected utility.

1 Introduction

Agents' search for partners is an inherent process in many MAS applications [15]. A common scenario in partnerships models is where the searching agent is satisfied with only one partner for forming partnership. Typical applications of this type include buyer-seller [2], peer-to-peer media exchange, dual-backup services [13], etc. In these applications each potential partnership suggests a different utility for the agent. The agent can't a-priori evaluate the expected utility from a partnership with any of the other agents, though it can evaluate the overall utility distribution. Learning about the expected utility from a partnership with a specific agent is possible by interacting with this agent. This process involves the consumption of some of the agent's resources (e.g. cost of search). Thus the agent's main challenge in such an application is to find a strategy for determining, at each stage of its search, whether to try and partner with one of the formerly

⋆ This research was supported in part by NSF under grant #IIS-0208608.

interacted agents, or resume its search. Partnering with an agent will yield an immediate benefit, while resuming the search might result in a better partnership (though inducing a further search cost). The search process is considered to be two-sided, as all agents are engaged in search, and a partnership between two agents will be formed only if both agents commit to it.

Traditionally, two-sided search models [3] consider humans engaged in dual search activities (marriage market, for instance) with pure sequential search strategies (where each party samples and evaluates one potential partner at a time). However, the advantages of agents in filtering and processing information, as well as their improved parallel interaction capabilities (in comparison to humans), suggests an improved two-sided search technique, by which the agents sample and assess several potential partner agents, during a search stage, in parallel. This is mostly beneficial when parts of the search costs do not necessarily depend on the number of potential partner agents sampled. In this case the parallel search reduces the average search cost per sampled potential partner.

Our goal in this paper is to present and analyze a dual parallel two-sided search model, where all agents in the electronic marketplace (or any other MAS environment) use the parallel search. The transition from one agent using the new technique into a dual usage requires a complete understanding of the dynamics which drive the agents' strategies towards a stable equilibrium. An important output of the analysis is an algorithm that significantly simplifies the process of extracting the equilibrium strategies. This type of algorithm is extremely valuable for both agents and their developers. As part of the discussion we show that in many cases the dual parallel search may yield a better utility for the agents, though in some cases the effect of the proposed search technology is an equilibrium, with lower expected utility for each agent. These are important inputs for market makers, when considering the integration of the parallel search techniques in the agents they supply to their customers.

As a framework application for our analysis we use an eCommerce environment where agents represent telephony service providers. Consider, for example, a service provider, serving a specific geographic location, having long term formal partners (other service providers) in various geographical destinations, with whom it has termination agreements[3]. At any time the service provider can produce a short term forecast for its unused bandwidth. In most cases, the amount of such bandwidth, and the relatively short period the service provider can commit to, makes the option of selling this remaining bandwidth non-beneficial. Alternatively, the service provider can use this bandwidth in an exchange process for reducing its termination costs. In the latter case, the service provider will create an agent operating in a designated market, where the other agents represents different service providers, serving different locations. The agent will seek to create ad-hoc short term partnerships with these agents, to make use of the unexploited bandwidth on both sides. Once a partnership is formed, each service provider will route some of its traffic to the remote destinations via the other service provider's

[3] An agreement defining a service provided by an interconnection provider, whereby it connects a call from a point of interconnection to a network termination point.

network, instead of using its costly formal long term termination partners. Similar application of exchanging unexploited resources, can be found in [13, 1].

In the telephony communication domain, the benefit that can be obtained from a given partnership is a function of the Quality of Service (QoS) that can be guaranteed by such a connection - service providers offer their customers different tariffs for different levels of services, as defined in a Service Level Agreement (SLA). The key measure for evaluating the Quality of Service is speech quality. While traditional methods of determining speech quality were based on subjective tests with panels of human listeners (ITU-T P.800/P.830, [9]), recent ITU standards [10] suggest automatic prediction of voice quality that would be given in typical subjective tests. This is done by intrusive end-to-end test calls (i.e. generating test traffic) and passive monitoring of traffic in strategic locations along the call route. It is notable, however, that speech quality is not linear, and as the call traverses the two networks, the quality can not be evaluated simply as a function of both networks' separate performances. Thus for the purpose of evaluating the perceived utility from each potential partnership, an agent needs to perform a full set of tests (and can't rely on former tests performed for the connection of its network with other networks). The costs of these tests represents the search cost in our model.

At this point the model diverges into two important variants, differing in the expected utility each partnership member obtains from the partnership. In the first variant, the utility each of the two agents forming a partnership obtain is equal. The second variant suggests different utilities for both agents, in a way that the utility gained by each of the agents will be considered as randomly drawn from a general distribution (see for example [3]). In the suggested telephony application there are many factors supporting each of the two variants, such as the SLAs structure (committing to the quality of both or just one of the inbound and outbound channels), different tariffs, attractiveness of different geographical locations, etc. We support both variants throughout the paper.

As we add the ability to interact with several potential partner agents in parallel, on top of the traditional sequential search model, we need to consider the search cost structure. The integration of passive monitoring and intrusive test calls, suggests a fixed and variable components in the cost structure. The fixed component can be associated with passive testing devices, monitoring all traffic simultaneously. The variable component is associated with the intrusive end-to-end test calls. The number of test calls performed is derived from the number of potential partners evaluated over each search round. Similar applications in which such dual parallel search can be used include secondary markets for exchanging remaining resources in those cases where selling them is not the core business of the organization, or when the overhead for selling them makes it non-beneficial, and thus an exchange mechanism is used (see [13]).

In the next section we address relevant multi-agent and matching literature. In section 3 we present the model. An equilibrium analysis and an efficient algorithm for finding the equilibrium strategy are provided in section 4. The incentive to deviate from the sequential two-sided search towards the parallel search is presented in section 5. Section 6 compares the two model variants. We conclude and suggest directions for future research in section 7.

2 Related Work

The application of agents seeking coalitions with other agents is quite common in MAS environments and in electronic markets in particular [5, 6, 2]. A specific case of such coalitions is the partnerships model where each agent seeks a single partner [12, 15, 13]. Different mechanisms for partnering suggest different assumptions regarding the agent's scanning capabilities, ranging from the option to scan as many agents as needed, through making use of a central matcher or middle agents [4] and up to a complete distributed process, without the help of a predefined organization or a central facilitator [11]. Our model is of the latter type, and assumes a distributed environment where each agent needs to invest resources for interacting with other agents and evaluate the potential partnerships.

The basic concepts of search, can be found in the classical search theory. Here, various models of one-sided search (assuming no mutual search activities) and two-sided search were suggested [7, 14, 3]. While the concept of parallel search (also known as variable sample sizes) was suggested for the one-sided search [8], two-sided models always assume a sequential search. Thus the uniqueness of the suggested model is in allowing all agents to use the parallel method. It is notable that the transition from the one-sided to two-sided models, when considering parallel search techniques, involves many complexities, as the equilibrium considerations become the main issue of the analysis[4]. Additionally, in search theory, search "costs" are usually modelled by the discounting of the future flow of gains, while in MAS environments the total search period is relatively short and utilities are immediate.

In a recent paper [12] we present an initial analysis of equilibrium in a two sided search model, where there are two types of searchers (buyers and sellers), and only one type may use parallel search. We have shown that the adoption of parallel search by buyers in C2C markets, leads to new equilibrium strategies, which can significantly improve the utility of the buyers (thus reducing the utility of the sellers using the sequential search). Nevertheless, our analysis was limited to the usage of the technique by one side of the interaction, and for the case where the perceived utilities are different. When all agents can use the parallel search technique, the problem of evaluating the stability of suspected equilibrium strategies becomes significantly more complex. The main challenge is in handling the enormous expansion of the strategy space. In sequential two-sided search the strategy space is bi-dimensional (having each agent's reservation value[5] on each axis), thus each agent's incentive to deviate from a given strategy, can be checked

[4] In one-sided search, the problem can be seen as a simple optimization problem, with no equilibrium concerns, since the focus is on a single searcher's decision

[5] As in most search models, in the following sections we show that the agents will use a reservation value based strategy. The agent will accept all offers that yield a utility greater than or equal to a reservation value, and reject all those that yield a utility less than this value. Notice the reservation value of the search strategy is different than a reservation price usually associated with a buyer or a seller that are not involved in a search. While the reservation price denotes an agent's true evaluation of a specific potential partnership, the reservation value of a search strategy is mainly a lower bound for accepted partnership, derived from the expected utility optimization considerations.

along its own reservation value. In our dual parallel search model, each potential strategy must be compared with all possible combinations of reservation values and the number of partners sampled over a search round.

3 The Dual Parallel Two-Sided Search Model

Consider an environment populated with numerous agents, seeking to form partnerships for their benefit throughout random interactions with other agents. As suggested in the introduction, the perceived utility for an agent from any suggested partnership with any specific agent, denoted by U, can be seen as randomly drawn from a population with p.d.f. $f(U)$ and c.d.f. $F(U)$, $(0 \leq U < \infty)$. We assume that agents, while ignorant of the utility obtained by partnering with specific agents, are acquainted with the overall possible partnership's utility distribution. This assumption is common in search models (see [3, 14, 13, 11]).

Integrating the above into the service providers application, we consider each agent as suggesting a termination service for a standard unit of time and volume[6]. Any random interaction between two agents, may yield a partnership for terminating calls, for the benefit of the two represented service providers. The utility from such a partnership is expressed in monetary units, as the service providers' main concern is revenue.

When using the parallel search, at any stage of its search the agent encounters N potential partner agents, interested in forming a partnership. This is in comparison to the traditional sequential two-sided search, where the agent samples only one other agent at any stage of its search [3]. For each encounter both agents evaluate the utility from such a partnership (in the telephony application this will mean testing the perceived connection between the two networks). We assume this utility is randomly drawn from a similar distribution function for all agents (either with a similar value or as two different values, according to the model variant), due to the high number of potential partners and the similar operational cost structure.

We denote, α and β as the fixed and variable cost components of a search stage. In the telephony application, these costs are associated with intrusive test calls and passive monitoring for testing the perceived connections with the N networks represented by the potential partnering agents. Thus the total search cost per a search round is $\alpha + \beta N$. Notice the values α and β are standard for all agents, as testing the quality of a connection between two service providers is conducted in similar methods based on standard testing devices. After evaluating each of the potential partnerships, each agent will make a decision whether to commit to one of them. Obviously, each agent is interested in partnering with the agent with whom the potential partnership will yield the maximum utility. A partnership will eventually be formed only if both agents are willing to commit to it. Otherwise the agents will resume their search according to the same process as described above. Notice that since each new interaction suggest a new potential utility, and since the agent may commit only to the best in its sample, deadlocks will never occur. Since the agents are not limited by a decision horizon and can

[6] Notice the agent can consider numerous potential partners as these include service providers from any geographical location, IXCs, and ISPs supporting VoIP.

control the intensity of their search, and the interaction with other agents does not imply any new information about the market structure, their strategy is stationary - an agent will not accept an opportunity it has rejected beforehand. Thus agents will use a reservation value strategy. The fast parallel interactions between agents, ensure finding a partner within reasonable time. This, along with the fact that utilities can be seen as immediate (due to the short partnership duration) allow us to ignore the influence of a discounting factor when considering expected utilities. Lastly, notice that as all agents are subject to a similar search cost, and the perceived utility can be seen as randomly drawn from the same population, all agents share the same reservation value.

As the search process is two-sided, the main challenge is in finding the equilibrium strategies. Any set of strategies that can't guarantee equilibrium stability will not hold. As suggested in the introduction, we distinguish between two variants of the above model. The first assumes both agents in a partnership will yield the same utility, while in the second each agent diversely evaluates the utility it can gain from a given partnership. Each variant will yield different equilibrium strategies and thus will differ with the expected utility gained by the agents.

For analysis purposes, we'll use several notations in the following sections. A strategy of sampling N other agents over each search round, and acting according to a reservation value x_N will be denoted (N, x_N). The expected utility of an agent when using strategy (N, x_N), will be denoted $V_N(x_N)$. As the agent is mainly concerned at each search round with the partnership offering the maximum utility in its sample, we will use the random variable U^N to denote the partnership with the maximal utility in an N size sample. The p.d.f. and c.d.f. of the variable U^N will be denoted $f_N(x)$ and $F_N(x)$, accordingly.

4 Agents Strategies and Equilibrium Dynamics

We start by formulating the expected utility for the agent when using a strategy (N, x_N), given the strategy (k, x_k) used by the other agents in the environment. In the variant where both agents obtain the same utility from a given partnership the expected future utility $V_N(x_N)$ is:[7]

$$V_N(x_N) = \frac{\int_{y=max(x_N,x_k)}^{\infty} y f_N(y) F(y)^{k-1} dy - \alpha - \beta N}{\int_{y=max(x_N,x_k)}^{\infty} f_N(y) F(y)^{k-1} dy} \tag{1}$$

This can be decomposed into two parts: the conditional expected utility from a partnership that will be eventually formed, and the aggregated search cost derived by the number of search cycles. The number of search cycles is geometric and the probability of success is $\int_{y=max(x_N,x_k)}^{\infty} f_N(y) F(y)^{k-1} dy$. In addition, notice that:

$$F_N(x) = F^N(x) \qquad f_N(x) = \frac{dF_N(x)}{dx} = N f(x) F(x)^{N-1} \tag{2}$$

[7] Due to space considerations, we include the basic formulation and sketch of proofs. The extended version of the paper, including the detailed formulation, proofs and discussions can be found at www.cs.biu.ac.il/ sarne/FullPapers/Dualparallel.pdf.

and thus, substituting (2) in (1) we eliminate the sample maximum notations from the equation and obtain a function which depends only on the cost structure α and β and the general distribution function $F(x)$:

$$V_N(x_N) = \frac{N\int_{y=max(x_N,x_k)}^{\infty} yf(y)F(y)^{N+k-2}dy - \alpha - \beta N}{N\int_{y=max(x_N,x_k)}^{\infty} f(y)F(y)^{N+k-2}dy} \tag{3}$$

Notice that any usage of $x_N < x_k$ will yield the same utility as using x_k. This is simply because any partnership suggesting a utility lower than this value will always result in a rejection from the other agent and consequently the continuance of the search. Similarly, for the variant where both agents obtain different utilities:

$$V_N(x_N) = \frac{(1-F_k(x_k))\int_{y=x_N}^{\infty} yf_N(y)dy - \alpha k - \beta Nk}{(1-F_k(x_k))(1-F_N(x_N))} \tag{4}$$

and by substituting (2) in the above equation we obtain a simpler function. Notice that unlike the expected utility in (1), here, any change in the reservation value, x_N, including a reduction to a value beyond x_k will affect the expected utility. This is because of the absence of any correlation between the utility an agent gains from a potential partnership and the probability of being accepted by the other agent. Thus:

$$\lim_{x_N \to 0} V_N(x_N) = E[U^N] - \frac{\alpha k + \beta Nk}{(1-F_k(x_k))} \tag{5}$$

Except for the above difference, the expected utility curve is quite similar in its structure for both variants of the model. As the agent increases its reservation value, and becomes more selective, the utility improvement gained by partnering with a better partner significantly decreases in comparison to the additional cost incurred by the increase in the number of search rounds. Thus both variants satisfy:

$$\lim_{x_N \to \infty} V_N(x_N) = -\infty \tag{6}$$

4.1 Reservation Value and Expected Utility Function

Having established the behavior of the expected utility function for upper and lower values of the reservation value, we now show that for both variants of the model, the expected utility function has a unique maxima.

Theorem 1. *When all other agents use strategy (k, x_k), an agent's expected utility function, $V_N(x_N)$, when using strategy (N, x_N) is quasi concave in x_N with a unique maxima satisfying:*

$$V_N(x_N) = x_N \tag{7}$$

Sketch of Proof: We present the proof for the variant where both agents obtain equal utilities from a partnership. The proof for the second variant is similar. Deriving (3) we obtain:

$$\frac{dV_N(x_N)}{dx_N} = \frac{f_N(x_N)F^{N+k-2}(x_N)\big(V_N(x_N) - x_N\big)}{N\int_{y=max(x_N,x_k)}^{\infty} f(y)F(y)^{N+k-2}dy} \equiv r(x_N)(V_N(x_N) - x_N) \;\;,\;\; x_N \geq x_k \tag{8}$$

A solution for (8) will require $V_N(x_N) = x_N$. Note that $f_N(x_N) > 0$ implies $r(x_N) > 0$, hence for x_N satisfying $V_N(x_N) = x_N$:

$$\frac{d^2 V_N(x_N)}{d{x_N}^2} = r^{'}(x_N)(V_N(x_N) - x_N) + r(x_N)({V_N}^{'}(x_N) - 1) < 0 \quad (9)$$

Thus $V_N(x_N)$ is quasi concave with a unique maxima. □

From (1-6) and theorem 1 we can sketch the graph $V_N(x_N)$. The basic structure of the curve is given in Figure 1, using the uniform distribution function and the parameters: $N = 3$, $k = 2$, $\alpha = 0.05$, $\beta = 0.005$ and $x_k = 0.55$ (substituted in equations (3) and (4)). Notice the difference between the two variants for reservation values $x_N \leq x_k$, as described above.

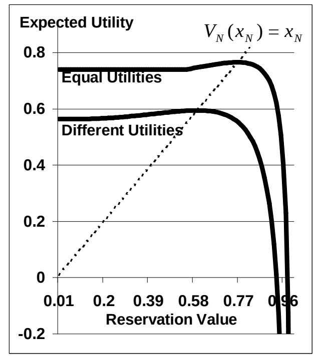

Fig. 1. Agent's expected utility

Once we have established the expected utility structure, we can suggest a simple algorithm for calculating the agent's optimal reservation value (and thus its expected utility, according to (7)) given the number of potential partners sampled, N, and the other agents' strategy (k, x_k). The algorithm makes use of a binary search and will always reach the agent's optimal reservation value, in a finite number of steps. Since the algorithm is very similar to the one we have suggested in [12] for the case where only one of the agents uses the parallel search technique, it will not be detailed in the current context.

4.2 Dual parallel Two-Sided Search Strategy

Having at hand an efficient means for finding the agent's optimal reservation value, given the strategy of the other agents (k, x_k), we move on to explore the dynamics affecting agents' strategies in the dual parallel search. We start by introducing two important equations that can be used for finding an agent's optimal reservation value, given the number of partners it samples over a search round, N, and the other agents' strategy (k, x_k).

Theorem 2. *When an agent samples N potential partners over a search round and the other agents use strategy (k, x_k): (a) The agent's optimal reservation value x_N in the variant with equal utilities, satisfies:*

$$(N{+}k{-}1)(\alpha{+}\beta N){=}N((max(x_k, x_N){-}x_N)(F^{k+N-1}(x_k){-}1){+}\int_{y=max(x_N,x_k)}^{\infty}(1{-}F^{k+N-1}(y))dy) \quad (10)$$

(b) The agent's optimal reservation value x_N in the different utilities variant, satisfies:

$$\alpha k + \beta N k = (1 - F_k(x_k)) \int_{y=x_N}^{\infty} (1 - F_N(y))dy \quad (11)$$

Sketch of Proof:

(a) Set the first derivative of (3) to 0 and use integration by parts, substituting: $dv = f(y)F(y)^{N+k-2}$ and $u = y$ in (3). Manipulating and rearranging the result we obtain (10).

(b) Use a similar methodology as in (a) to obtain (11). □

The suggested equalities in theorem 2 are useful computational methods for calculating x_N. This can be particularly important when checking the incentive to deviate from a given strategy. The following theorem 3 captures another important characteristic of the agent's strategy that will aid us in finding the equilibrium.

Theorem 3. *(a) In both variants of the model, when all other agents sample k potential partners over a search round, if an agent's expected utility of sampling* $k+1$ *potential partners,* $V_{k+1}(x_{k+1})$ *is smaller than* $V_k(x_k)$*, then the expected utility when sampling N potential partners,* $V_N(x_N)$*, where* $N > k+1$*, is also smaller than* $V_k(x_k)$*. (b) Similarly, when all other agents sample k potential partners over a search round, if an agent's expected utility of using* $k-1$ *potential partners,* $V_{k-1}(x_{k-1})$*, is smaller than the expected utility when using k potential partners,* $V_k(x_k)$*, then the expected utility when using N potential partners, where* $N < k-1$ *is also smaller than* $V_k(x_k)$*.*

Sketch of Proof:

(a.1) For the variant with equal utilities - assume otherwise (e.g. the expected utility when using N, when $N > k+1$ and $V_k(x_k) \leq V_N(x_N)$, is greater than $V_k(x_k)$). Then subtract two instances of equation (10) using N and k. Substituting β as obtained from manipulating the inequality $x_N - V_{N+1}(x_{N+1}) > 0$, using (1), and analytically exploring both sides of the inequality expression, we obtain a contradiction, given the incorrectness assumption. The detailed proof, depicted on several pages, can be found in the full version of the paper.

(a.2) For the variant with different utilities - assume otherwise ($x_N > x_k$). Create 3 instances of (11) for k, $k+1$ and N. Then subtract the first two and the last two equations. Set all integrals' lower bounds to x_N and notice the obtained inequality is strictly negative (by deriving the term inside the integral), which leads to a contradiction, given the incorrectness assumption.

(b) In a similar methodology as used for the proof of part (a), making use of the appropriate modifications of the inequalities. □

The above theorem is an important milestone towards the formation of an algorithm for finding the equilibrium, as it allows us to limit the number of potential partners that needs to be considered, when checking the stability of a potential equilibrium strategy (only the former and next subsequent numbers need to be considered).

4.3 Finding the Equilibrium

Having the results and proofs given in previous subsections, we now suggest an efficient method for finding the equilibrium in the dual parallel two-sided search model. From theorem 3, we conclude that in order to check the stability of strategy (N, x_N), one only needs to check the expected utility of an agent when deviating to strategy $(N+1, x_{N+1})$ and $(N-1, x_{N-1})$. This can be simply calculated using equations (10) and (11). If the expected utility for strategy (N, x_N) is greater than the other two, then this is an equilibrium.

As the agents are identical, they will all use the same equilibrium strategy (if an equilibrium exists), thus their expected utility will be identical. This resolves the uncertainty in case of a multiple equilibria scenario - all agents will use the equilibrium strategy with the highest expected utility.

An important consideration is the upper bound for N, when seeking the equilibrium strategy. An equilibrium doesn't necessarily exist, and while using the proposed method one may wonder when to stop as N grows and the equilibrium conditions are not satisfied. We propose a simple upper bound that can be used with both variants of the model.

Theorem 4. *(a) An upper bound value for the number of partners to be considered over a search round, in the variant with equal utilities, is the solution* $N = N_{max}$ *of the equation:*

$$E[U^N] = \alpha + \beta N \tag{12}$$

(b)An upper bound value for the number partners to be considered over a search round, in the variant with different utilities, is the solution $N = N_{max}$ *of the equation:*

$$E[U^N] = \alpha N + \beta N^2 \tag{13}$$

Sketch of Proof:

(a) + (b) - by substituting (12-13) in (3-4), we attain a negative expected utility. The expected utility will remain negative for any $k \geq N_{max}$. Though the agents will unavoidably abandon search activity for these k values, and if no equilibrium was found up to this point then the problem with the current search cost structure (α and β values) doesn't have a pure equilibrium solution. Such an N_{max} value can always be found as the left hand side of equations (12-13) is concave and the right hand side is convex (except for the case where the agents would have initially abandoned the search, e.g. where the left hand side of the equations is smaller than the right hand term for $N = 1$).

To summarize the methodology for finding the equilibrium (if any exists), we suggest the following algorithm.

Algorithm 1 *An algorithm for finding the equilibrium strategy* (N, x_N) *for the dual parallel search model.*

Input: α, β *- cost structure coefficients; F(x) - the utility c.d.f.*
Output: *(*$V_N(x_N), x_N, N$*) - Equilibrium strategy, if one exists, otherwise a "no equilibrium" message.*

```
01. Set N_max according to equation (12).
02. Set List[]=null;
03. for (N=1;N≤ N_max;N++) {
04.     if (N>1) calculate V_{N-1}(x_{N-1}) using (10), where k = N;
05.     calculate V_N(x_N) and V_{N+1}(x_{N+1}) using (10), where k = N;
06.     if (V_N(x_N) ≥ V_{N+1}(x_{N+1}))
07.         if (N=1) or (V_N(x_N) ≥ V_{N-1}(x_{N-1}))
08.             add (V_N(x_N), x_N, N) to List[].
09. }
10. If (List[]!=null) return member with highest V_N(x_N)
11 else return("no equilibrium");
```

The above algorithm is applicable for the variant with equal utilities. In order to use it with the other variant, one needs to replace equations (12) and (10) with equations (13) and (11).

If there is an equilibrium solution to the problem, the algorithm will find the equilibrium strategy in $o(N_{max})$ stages, where N_{max} is the upper bound, calculated according to (12-13), and its value is mostly influenced by the utility distribution function. The innovation of the proposed algorithm is in bounding the space of possible strategies which needs to be compared for any suspected equilibrium strategy. The complexity of the solution, in the absence of such a bound is discussed in the next section.

5 The Incentive to Use Parallel Search

Considering equations (1) and (4), it is clear that the traditional sequential search model is a specific case of the general dual parallel search as described in the model section. For example, by substituting $N = 1$ in (4), we obtain a similar expected utility function for the sequential search model as described in [3]. Nevertheless, the sequential two-sided search will not be stable in many cases, since single agents have an incentive to deviate from the sequential strategy for many plausible combinations of α and β values. Figure 2 demonstrates this phenomena for the uniform distribution function. As the utility varies from 0 to 1, the bottom triangular area represents all plausible α and β combinations where the agents will consider a sequential two-sided search (e.g. where the expected utility for the agents in a sequential equilibrium strategy is positive). Out of this area, we have isolated (on the left side) all combinations of α and β where an agent can increase its expected utility by deviating from such a sequential strategy (assuming all other agents' strategies are sequential). For calculation purposes we used equations (10) and (11). Notice that for a large portion of the cost structures, any single agent has an incentive to deviate from its sequential strategy. Furthermore, the incentive to deviate from the sequential strategy is mainly for the combinations of α and β with small values (in comparison to the average utility from a partnership), which characterizes most MAS applications.

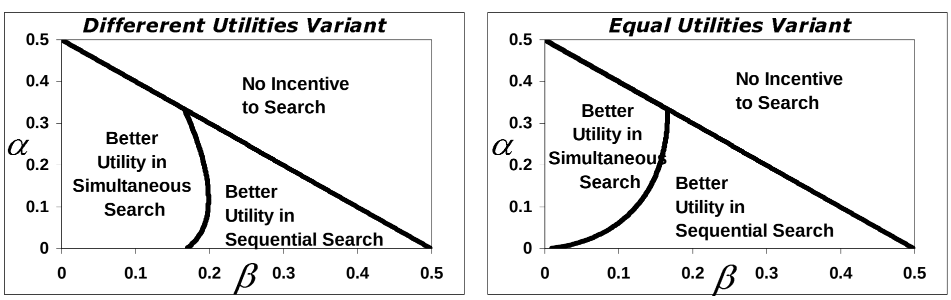

Fig. 2. The incentive for using parallel search

As each agent has the incentive to use the parallel search technique, the traditional sequential two-sided search model transforms into a dual parallel search. This also suggests a possible improvement in the expected utility of the agents, in comparison to the traditional sequential two-sided search models. Again, the improvement is mostly noticeable for values where α and β are relatively small in comparison to the utilities gained from the partnerships.

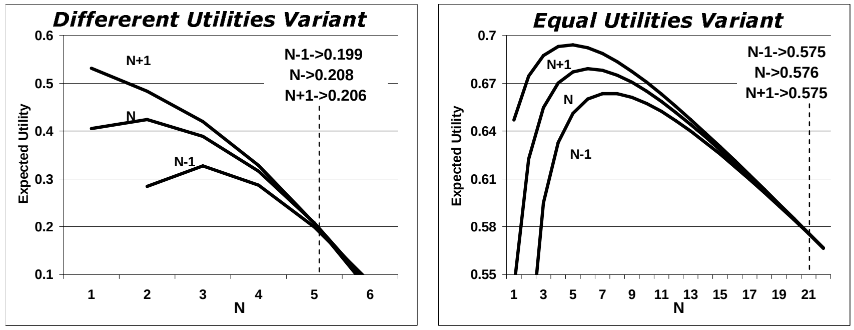

Fig. 3. Deviating from equilibrium

The transition into a dual parallel search model doesn't necessarily guarantee an expected utility improvement. In some cases, the dual parallel usage may result, inevitably, in an equilibrium where the agents worsen their expected utility. In such a scenario, all agents could gain more by using a sequential strategy, but each, separately has an incentive to deviate towards a parallel strategy, resulting eventually in a non-optimal result. This is demonstrated in figure 3, again for the uniform distribution function, with $\alpha = 0.1$ and $\beta = 0.05$. The middle curve represents the expected utility when sampling N potential partners over a search round according to the horizontal scale. The upper and lower curves represent the expected utility when a single agent deviates to $N+1$ and $N-1$, respectively. Here, the equilibrium utility is obtained when using $N = 5$ over a search stage (taking the variant of different utilities). Even though the sequential two-sided search utility is greater, none of the agents will maintain such a strategy as both agents have an incentive to deviate towards a higher number of sampled partners in a search round. The example given for the variant with equal utilities demonstrates a scenario where the expected equilibrium utility when using a parallel ($N = 21$) search is higher than when using a sequential search, though it is far below the possible expected utility that could have been obtained by using $N = 6$.

Thus, extra care should be taken in the analysis of the dual parallel search equilibrium. This can be extremely important for market makers for understanding the consequences of allowing the agents to sample more than a single potential partner over each search stage, or even for actually limiting the number of partners that can be sampled by the agents at each turn.

6 Model Variants Comparison

Throughout the examples given in the previous section, one might notice from figure 2 that there is a stronger incentive to use the parallel search technique in the variant where the utilities for both parties are different. This doesn't necessarily mean the expected utility in this variant is greater. In fact, as notable from figures 1 and 3, the agents' expected utility is greater in the variant where both agents gain the same utility. The explanation for this phenomena can be found in the strong correlation between both agents acceptance decision, when utilities are equal, in comparison to no correlation at all in the second variant. If both agents gain a similar utility from a given partnership, then the probability that each of them is the highest in the other agent's sample is relatively high. On the other hand, when the expected utility from a given partnership is random, the probability of being the agent with the highest utility to a given potential partner is $1/N$. This insight can be formally proven, as the following theorem states.

Theorem 5. *When all agents use a dual parallel search with N potential partners over a search round, the variant with equal utilities will yield the agents using the equilibrium strategy a higher expected utility in comparison to the equilibrium utility that can be gained in the other variant.*

Sketch of Proof:
Set $k = N$ in equations (10) and (11) to obtain equilibrium reservation values, and isolate the term $\alpha + \beta N$. Then subtract the two equations, to obtain:

$$\frac{N}{2N-1}\int_{y=x_N^I}^{\infty}(1-F^{2N-1}(y))dy = \frac{1}{N}\int_{y=x_N^{II}}^{\infty}(1-F^N(x_N^{II}))(1-F^N(y))dy \quad (14)$$

Where x_N^I is the reservation value of the equal utilities variant and x_N^{II} is the reservation value of the different utilities variant. Obviously the equation can hold only if $x_N^I \geq x_N^{II}$, and from theorem 2 we obtain $V_N(x_N^I) \geq V_N(x_N^{II})$. □

7 Conclusions

As demonstrated throughout this paper, in many cases, an agent engaged in search has an incentive to adopt the parallel search technique. Nowadays, as agents' technology is a reality and traditional processing and communication limitations were removed, it is high time to consider the dual parallel search model in MAS domains and in particular the two-sided search application for the electronic marketplace. We manage to present a complete equilibrium analysis, and suggest an efficient algorithm for calculating the agents' equilibrium strategies, given the environment parameters (utilities distribution and search cost coefficients). Deriving the equilibrium strategy is a complex task, as all agents can control both the number of partners they sample and their acceptance criteria; thus the challenge of finding a stable set of strategies becomes significantly complex. The novelty of the proposed algorithm is in the capability to bound the relevant strategy space and quickly eliminate non-equilibrium strategies. The adoption of the method can significantly improve the expected utility either when used one-sidedly or simultaneously by all agents. Nevertheless, as part of the discussion, we show that in some cases equilibrium dynamics might drive the agents into a strategy where the number of partners sampled results in a non-optimal expected utility. In some rare cases this could even worsen the expected utility in comparison to the sequential search. The later scenario further emphasizes the importance of the analysis given and the proposed algorithm, as market makers can use the results to understand and evaluate the influence of the decision to allow agents in their marketplace to use parallel search. The proposed analysis was followed by a plausible eCommerce application from the telephony call termination partnering domain. The division into two specific variants of the model, differing by the perceived utility from a given partnership, extends the variety of applications this model can be integrated in.

Notice that throughout the paper we assumed that the agent commits only to the potential partnership with the agent associated with the highest utility in the sample (assuming it is above its reservation value). Nevertheless, in a given sample, there might be several agents with a utility that might be greater than the reservation value being used. Thus the agent can improve its expected utility by considering committing also to the next best agent in the sample, upon receiving a rejection from the current potential partner agent. It is notable, however, that this technique has some setbacks. First, all agents need to wait for a rejection/acceptance message from their best sampled agent before considering committing to their next best candidate. This creates many constraints, and necessitate a protocol in which all the agents conduct each search round simultaneously. This might also require the introduction of a discounting factor into the model. Second, because of the significant amount of time that needs to be allocated for each search round (because of the expected bottlenecks), a failure or malfunction of one of the agents might, in extreme cases, drive the entire system into a "hold" position. Thus, prior to considering such a model for future research, a substantial research effort should be made to build the infrastructure and protocols for handling the additional dependencies and constraints involved.

References

1. R. Azoulay-Schwartz and S. Kraus. Stable strategies for collaborative behavior of agents in answering queries. *Artificial Intelligence journal*, 154(1-2):43–93, 2004.
2. S. Breban and J. Vassileva. Long-term coalitions for the electronic marketplace. In *B. Spencer, ed., Proc. of E-Commerce Applications Workshop*, 2001.
3. K. Burdett and R. Wright. Two-sided search with nontransferable utility. *Review of Economic Dynamics*, 1:220–245, 1998.
4. K. Decker, K. Sycara, and M. Williamson. Middle-agents for the internet. In *Proc. of IJCAI*, pages 578–583, 1997.
5. T. Ito, H. Ochi, and T. Shintani. A group-buy protocol based on coalition formation for agent-mediated e-commerce. *IJCIS*, 3(1):11–20, 2002.
6. K. Lermann and O. Shehory. Coalition formation for large scale electronic markets. In *Proc. of ICMAS'2000*, pages 216–222, Boston, 2000.
7. J. McMillan and M. Rothschild. Search. In Robert J. Aumann and Amsterdam Sergiu Hart, editors, *Handbook of Game Theory with Economic Applications*, pages 905–927. 1994.
8. P. Morgan. Search and optimal sample size. *Review of Economic Studies*, 50(4):659–675, 1983.
9. ITU-T Recommendation P.800. Methods for subjective determination of transmission quality. August 1996.
10. ITU-T Recommendation P.862. Perceptual evaluation of speech quality (pesq), an objective method for end-to-end speech quality assessment. February 2001.
11. D. Sarne and S. Kraus. The search for coalition formation in costly environments. In *Proc. of CIA2003*, pages 117–136, 2003.
12. D. Sarne and S. Kraus. Agent's multiple inquiries for enhancing the partnership formation process. In *Proc. of CIA2004*, pages 286–302, 2004.
13. D. Sarne and S. Kraus. Time-variant distributed agent matching applications. In *Proc. of AAMAS-2004*, pages 168–175, 2004.
14. R. Shimer and L. Smith. Assortative matching and search. *Econometrica*, 68(2):343–370, 2000.
15. K. Sycara, S. Widoff, M. Klusch, and J. Lu. Larks: Dynamic matchmaking among heterogeneous software agents in cyberspace. *Autonomous Agents and Multi-Agent Systems*, 5:173–203, 2002.