

Trajectories of Goods in Distributed Allocation

Yann Chevaleyre
Université Paris-Dauphine
yann.chevaleyre@dauphine.fr

Ulle Endriss
University of Amsterdam
ulle@illc.uva.nl

Nicolas Maudet
Université Paris-Dauphine
maudet@lamsade.dauphine.fr

ABSTRACT

Distributed allocation mechanisms rely on the agents' autonomous (and supposedly rational) behaviour: states evolve as a result of agents contracting deals and exchanging resources. It is no surprise that restrictions on potential deals also restrict the reachability of some desirable states, for instance states where goods are efficiently allocated. In particular *topological* restrictions make any attempt to guarantee asymptotic convergence to an optimal allocation impossible in most cases. In this paper, we concentrate on the *dynamics* of such systems; more precisely we study the trajectories of goods in such iterative reallocative processes. Our first contribution is to propose an upper bound on the length of the trajectories of goods, when agent utility functions are modular. The second innovative aspect of the paper is then to discuss how this affects, on average, the quality of the states that are reached. Finally, a preliminary study of the non-modular case is proposed, examining how synergistic effects between items can affect their trajectories.

Categories and Subject Descriptors

I.2 [Artificial Intelligence]: Distributed Artificial Intelligence

General Terms

Algorithms, Economics

Keywords

Resource Allocation, Dynamics of Complex Systems

1. INTRODUCTION

Resource allocation problems play an important role in multiagent systems research [2, 6, 4]. Firstly, multiagent systems are often the technology of choice for tackling the resource allocation problems occurring in various applications. Examples include logistics, grid computing, and electronic commerce. Secondly, for almost *any* application modelled as a multiagent system, resource allocation issues will typically arise even when this is not the primary purpose for

Cite as: Trajectories of Goods in Distributed Allocation, Yann Chevaleyre, Ulle Endriss and Nicolas Maudet, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp.1111-1118.

Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

building the system. For instance, agents may require certain resources to be able to carry out a specific task, and those resources could reside with other agents at the time, which requires negotiation techniques for identifying situations where those resources are most usefully reallocated.

Research on multiagent resource allocation has tackled a variety of issues in the past. The following are some of the most prominent lines of research:

- Resource allocation and negotiation require suitable *communication protocols*. If (combinatorial) auctions are used to identify a good allocation, then the issue of the interaction protocol is fairly simple, but for distributed approaches to multiagent resource allocation, where individual agents can freely search for potential trading partners this can become a major research challenge. Much research in this area builds on the well-known *contract net protocol* [19].
- A second line of research is concerned with developing successful *negotiation strategies*. The focus here is on building negotiating agents that will perform well in practice [14, 17, 20].
- Research of a more fundamental nature is concerned with questions in *mechanism design* [5]. How can we design rules for resource allocation mechanisms that will give incentives to the participating agents to behave in certain desired ways (e.g. reveal their true valuations when making deals)?
- Another interesting topic in multiagent resource allocation concerns the quality of the allocations of resources that we can expect to achieve if a certain mechanism is used. For instance, it is sometimes possible to prove that a particular negotiation regime will force a system to converge to a state that is socially optimal according to various efficiency or fairness criteria borrowed from social choice theory [10].

In this paper, we propose a new perspective. We are interested in the *dynamics of multiagent resource allocation*. What can be said about the movements of an individual good when autonomous agents negotiate with each other in a distributed manner? And secondly, once we have gained a basic understanding of these dynamics, what can be said about the influence of these movements on the quality of negotiation outcomes?

We are going to address this type of problem in the context of a particular resource allocation framework studied by

a number of authors [18, 10]. In this framework groups of agents can agree on a deal regarding the reallocation of some of the goods held by them amongst themselves (possibly including side payments). The assumption is that agents are individually rational (IR) in the sense of only taking part in deals that result in an immediate improvement of their utility. We can then ask under what circumstances a sequence of such locally negotiated deals would converge towards an optimal allocation (for different optimality criteria). While the global optimality of such distributed processes can only be guaranteed under ideal circumstances, different natural restrictions can be considered:

- The protocol may restrict the trading opportunities between agents, by restricting the number of agents and/or the number of resources possibly involved in a deal. In particular, the class of deals involving a single resource at a time have been widely used, for obvious practical reasons.
- The domain may impose some restriction by excluding specific allocations (for instance, you may want to assign at least one resource to each agent).
- The domain may impose a topological structure on the negotiation system that defines which agents can actually interact with which other agents [3].

Among these, the restrictions pertaining to the topological structure that may affect the framework are particularly difficult to grasp. This is so because the behaviour of such a system will eventually depend on the nature of the graph, on the way resources are allocated to start with, and on the distribution of the agents' preferences.

In this paper, we analyse the dynamics of such systems. More precisely, we shall analyse the *trajectories of goods*, restricting initially our attention to the modular case where no synergy exist between resources. We first obtain a result showing that, if utilities are drawn randomly from some distribution, then the expected length of such a trajectory can be (upper) bounded by the square of the degree of the graph. Remarkably, this result is fairly general. In particular, it does not depend on the distribution. Also, we will see that the bound holds for a wide class of possible strategies that agents may employ to decide upon what agent to deal with, as long as the exchanges agreed upon remain rational.

Having established this result, we investigate to what extent the utilitarian notion of optimality (requiring the sum of the utilities of the agents in the system to be maximal) can be derived from this result. Finally, we pave the way for a study of the non-modular case.

The remainder of this paper is as follows. In the next section we set up the resource allocation framework we shall use in this paper. In particular, we discuss the different classes of strategies agents may employ when confronted with more than one rational deal option. Section 3 then studies the length of a good trajectory, and we propose an upper bound that is quadratic in the max degree of the graph. The result is general enough to apply to any strategies, and to any distribution. Then, we show how this relates to the expected average social welfare. Finally, we begin an investigation of the non-modular case, by studying how the goods interact with one another, and relate this to the expected trajectory length. After a discussion on related work, Section 7 concludes.

2. ALLOCATION FRAMEWORK

We consider the setting of a finite set of *agents* $\mathcal{A} = \{0, \dots, n-1\}$ negotiating over the allocation of a finite set of m indivisible *resources* \mathcal{R} (or *goods*, or *items*). Not every agent may be able to “see” all of the other agents. A *negotiation topology* is an undirected graph $G = (\mathcal{A}, E)$, the vertices of which are the agents in \mathcal{A} . Two agents i and j stand in the relation E iff they can see each other. This means that i and j may engage in negotiation and exchange resources. Note that our visibility relation E is symmetric (the graph is undirected); this is important to be able to define negotiation along graphs in a meaningful manner.

An *allocation* $A : \mathcal{A} \rightarrow 2^{\mathcal{R}}$ specifies for each agent i the bundle of resources owned by i , such that each and every resource gets assigned to exactly one agent. A *deal* can be described as a pair of (distinct) allocations $\delta = (A, A')$, fixing the situation before and after the exchange. Existing work within this framework, e.g. [18, 10, 8], either allows for any such deals to take place or considers structural restrictions in terms of the number of goods moving between agents. For example, a “swap deal” [18] is any deal between two agents whereby exactly one good is moving from the first to the second agent, and *vice versa*. Instead, in this paper, we only consider deals involving a single resource.

Each agent $i \in \mathcal{A}$ is equipped with a *valuation function* $v_i : 2^{\mathcal{R}} \rightarrow \mathbb{Q}$, to express their preferences over alternative bundles of resources. In this paper, we shall mostly restrict our attention to domains where no synergy (either positive or negative) between resources can occur. Valuation functions are then simply *modular*, that is $v_i(B) = \sum_{r \in B} v_i(\{r\})$, where B represents a possible bundle of resources owned by agent i . For the sake of simplicity, we will write α_r^i as a shorthand for $v_i(\{r\})$, the value agent i assigns to the resource r .

2.1 Rationality

Deals may be coupled with monetary *side payments*. These are given in terms of a *payment function* $p : \mathcal{A} \rightarrow \mathbb{Q}$ satisfying $\sum_{i \in \mathcal{A}} p(i) = 0$. A positive value $p(i)$ means that agent i is paying; a negative value $p(i)$ means that agent i is receiving the respective amount. A deal $\delta = (A, A')$ is called *individually rational* (IR) iff there exists a payment function p such that $v_i(A'(i)) - v_i(A(i)) > p(i)$ for all agents $i \in \mathcal{A}$, except possibly $p(i) = 0$ for agents i with $A(i) = A'(i)$. We assume that agents will only agree on such IR deals, *i.e.* deals that benefit everyone involved. The sum of all previous payments $p(i)$ made by agent i is given by its *payment balance* $\pi(i)$.

2.2 Strategies

Another assumption that we will make concerns the agents' *strategies*, that is, how will they select their trading partner when faced with several (rational) deal options in their neighbourhood. It is possible to identify several classes of strategies, depending on the input required to compute the partner agent with whom the deal will be contracted.

- *blind* strategies—agents are only allowed to check which of their neighbours are proposing rational deals to make their decision. For instance, on the basis of this information, we can assume that the agent will select the first partner proposing a rational deal following a predefined lexical order; or alternatively just pick one at random until a rational deal is found.

- *heuristic* strategies—agents select partner on the basis of heuristics regarding the potential utility profit among all the neighbours. In the context of myopic agents, certainly a natural strategy is to seek to maximize its immediate potential utility gain, that is, to pass the resource to the neighbour valuing it the most (we shall from now on call this the *max* strategy).

It is noteworthy that blind strategies are *not* based on the value that agents assign to resources. On the other hand, an agent using heuristic strategies must first get to know the utility all their neighbours assign to a specific resource to contract a deal.

2.3 Social Welfare

A *negotiation state* (A, π) is a pair of an allocation and a function giving the payment balance for each agent. Finally, each agent i is equipped with a quasi-linear *utility function* u_i mapping pairs of resource bundles and past payments to a utility scale: $u_i(B, x) = v_i(B) - x$. For example, $u_i(A(i), \pi(i))$ is the utility of agent i in state (A, π) , while $u_i(A(j), \pi(j))$ is the utility that i would experience if it were to swap places with j (in terms of both the bundle owned, and the sum of payments made so far).

The *social welfare* $sw(A)$ of an allocation A is defined as

$$sw(A) = \sum_{i \in \mathcal{A}} v_i(A).$$

This is the *utilitarian* definition of social welfare, the one mostly used in the multiagent systems literature [21]. But several of the other notions of social welfare developed in the social sciences are also of interest. For instance, if egalitarian social welfare [15] is used then we should aim at maximizing the welfare of the poorest agent of the society.

It is known in general [10] that deals are rational iff they increase the utilitarian social welfare, and that sequences of deals involving only a single resource are guaranteed to converge to an optimal outcome with respect to that measure, as long as the agents' utility functions are modular and the negotiation topology is a fully connected graph. As soon as some dealing opportunities are restricted by a negotiation topology, convergence properties cannot be guaranteed any longer (because it is easy to construct examples that require exactly that deal to reach the optimum). In this context, the next big thing would be to gain some insight on the average behaviour of the system. This calls for an analysis of the dynamics of multiagent resource allocation.

3. LENGTH OF AN ITEM TRAJECTORY

In this section, we study the expected length of an item trajectory in the case of modular utilities. Indeed, each resource r is initially owned by an agent, and then is sold to another neighbour agent valuing r more, and so on. Thus, r can be viewed as taking a *walk* along G , moving only towards agents valuing it more. This is precisely the walk of this resource along the graph that we call the *trajectory* of r . The main theorem of this section (upper) bounds the *expected* value of the length of such a trajectory, in any type of graph, as a function of its maximum degree Δ . To obtain this, we will assume that there exists for each resource r a distribution \mathcal{D}_r from which the coefficients α_r^i have been drawn independently, for each agent i . We also assume that for a given resource r , all values α_r^i differ.




Figure 1: Example of topology G (left). Suppose $\alpha_r^0 = 2, \alpha_r^1 = 1, \alpha_r^2 = 3, \alpha_r^3 = 5, \alpha_r^4 = 0.5, \alpha_r^5 = 9$. Then the digraph G^* (right) shows the possible moves of r .

Before going further, note that for each resource r the values of α_r^i induce a total ordering on the agents $0 \dots n-1$. We will denote by $i \succ j$ the fact that $\alpha_r^i > \alpha_r^j$.

Note that an edge (i, j) of the graph G can be crossed by resource r owned by i iff $i \prec j$, otherwise the deal would not be individually rational. Thus, the graph G hides a directed graph G^* in which each edge is one-way, as illustrated by Figure 1. The orientation of the edges shows how the resource can move.

Because coefficients are drawn independently from a random distribution, all the $n!$ orderings are equiprobable. However, this does *not* imply that all possible digraphs induced from G are equiprobable. Consider Figure 2. From the graph G shown on top, the four possible digraphs G^* are displayed. Note that G_a^* is induced by the ordering $0 \prec 1 \prec 2$, whereas G_b^* can be induced by both $0 \prec 2 \prec 1$ and $2 \prec 0 \prec 1$. Thus, $Pr(G_b^*) = 2 \times Pr(G_a^*)$. This means we cannot simply assume that all digraphs are equiprobable.




Figure 2: Digraphs are not equiprobable

Let us now show how the expected length of trajectories can be upper-bounded. In the following lemma, we consider the simpler case where graphs are trees, and obtain a bound. Then, in Theorem 1, we will show that this bound still holds for any type of graph.

LEMMA 1 (LENGTH OF TRAJECTORIES IN TREES).

Let G be a tree of max degree Δ . Suppose that coefficients composing the agents' utility functions have been drawn independently from distribution \mathcal{D}_r . Then the expected length of a resource trajectory in the graph is bounded by Δ^2 .

PROOF. We will consider a *single* resource r , the associated distribution \mathcal{D}_r , and the values $\alpha_r^0 \dots \alpha_r^{n-1}$ (which we will assume to be all different) drawn from it. Because utilities are modular, resources can be treated independently, and the present result still holds if we have more than one

resource. Without loss of generality, let vertex 0 be the agent initially holding r . Our first goal is to find the probability that a resource starting at 0 makes a walk over exactly k arcs, which we will note $Pr[len = k]$.

Formally, the walk of a resource is an infinite sequence of vertices x_0, x_1, \dots . At time 0, we have $x_0 = 0$. Once a vertex with no out-going arc has been reached (meaning no more rational deal is possible), the resource will remain on the same vertex forever. Thus, if the walk occurs over k arcs exactly, then $x_k = x_{k+1}$. Clearly, each variable x_t is a *random variable* whose distribution depends on the strategy used, the topology of the graph, and on the distribution \mathcal{D}_r . We can thus rewrite the above probability:

$$Pr[len = k] = Pr[x_0 \neq x_1, \dots, x_{k-1} \neq x_k, x_k = x_{k+1}]$$

Using the standard “chain of events” formula $Pr[\bigwedge A_i] = \prod [A_i | \bigwedge_{j < i} A_j]$, combined with the fact that $x_{i+1} \neq x_i$ always implies $x_i \neq x_{i-1}$, we get:

$$\begin{aligned} Pr[len = k] &= Pr[x_0 \neq x_1] \times \dots \\ &\quad \times Pr[x_1 \neq x_2 | x_0 \neq x_1] \\ &\quad \times \dots \\ &\quad \times Pr[x_k = x_{k+1} | x_{k-1} \neq x_k] \end{aligned}$$

We now compute the different members of this probability. Let d_i refer to the degree (in and out) of vertex x_i .

First, $Pr[x_0 = x_1]$ is the probability that the out-going degree of vertex 0 is null. This is equal to the fraction of orderings in which vertex 0 is ranked higher than its immediate neighbors, among all possible orderings. In other words, this is related to the number of orderings in which $0 \succ 1, 0 \succ 2, \dots, 0 \succ d_0$ divided by the total number of orderings among $d_0 + 1$ agents, which makes $Pr[x_0 = x_1] = \frac{d_0!}{(d_0+1)!} = \frac{1}{d_0+1}$, and thus $Pr[x_0 \neq x_1] = \frac{d_0}{d_0+1}$.

Next, let us show how to evaluate $Pr[x_k = x_{k+1} | x_{k-1} \neq x_k]$. This probability is related to the ranking of x_k compared to that of its children. First note that the ranking of these children among all vertices is *independent* from the knowledge we have concerning x_{k-1} . To illustrate this, consider Figure 3, and suppose $x_0 = 0$ and $x_1 = 2$. Clearly, the ranking of the children of 2 is *independent* from the fact that $0 < 2$ (as materialized by the dotted line on the figure). Also notice that $x_{k-1} \neq x_k$ implies x_k is ranked higher than previous vertices $x_0 \dots x_{k-1}$, which increases the probability that all of the $d_k - 1$ remaining arcs of x_k are incoming. Applied to this example, we would write $Pr[2 < 3 \wedge 2 < 4 | 0 < 2] \geq Pr[2 < 3 \wedge 2 < 4]$. We can now generalize the following way: $Pr[x_k = x_{k+1} | x_k \neq x_{k-1}] \geq \frac{(d_k-1)!}{d_k!} = \frac{1}{d_k}$.




Figure 3: r moves from 0 to 2

We can now compute the overall probability, for $k > 0$.

$$\begin{aligned} Pr[len = k] &\leq \frac{d_0}{d_0+1} \times \frac{d_1-1}{d_1} \times \dots \times \frac{d_{k-1}-1}{d_{k-1}} \\ &\quad \times Pr[x_k = x_{k+1} | x_{k-1} \neq x_k] \\ &\leq \frac{\Delta}{\Delta+1} \times \left(\frac{\Delta-1}{\Delta}\right)^{k-1} \end{aligned}$$

In order to derive a formula valid for any walk, we have introduced in the last inequation Δ , which refers to the max degree of the graph. Now we can compute the average number of moves made by r :

$$\begin{aligned} E[len] &= \sum_{k=0}^{max\ walk\ length} k \times P(len = k) \\ &\leq \frac{\Delta}{\Delta+1} \times \sum_{k=1}^{\infty} k \left(\frac{\Delta-1}{\Delta}\right)^k \\ &= \frac{\Delta}{\Delta+1} \times \Delta \times (\Delta-1) \\ &\leq \Delta^2 \end{aligned} \quad \square$$

Now we will investigate how the bound derived for the case of trees can be used in the general case. The line of reasoning that we will follow is best explained by means of an example.




Figure 4: r moves from 0 to 2: we know that $1 < 2$

Consider a resource moving across arcs of the graph on Figure 3, using the lexical fixed order strategy (i.e. choosing accessible neighbors with lowest index). The first step of the trajectory has led the resource from 0 to 2. What does this tell us? Knowing that $x_1 = 2$ is very informative in this case (this is where this crucially differs from the case of trees), for the resource would have gone from 0 to 1 in the previous step, were it possible. In other words, this means that $1 < 0$, hence it must be the case by transitivity that $1 < 2$. (But this does not say anything about vertex 5). This means that the probability to stop at this step is *increased*, and our bound applies. The following proof of Theorem 1 generalizes this argument and shows that it applies to all blind and max heuristic strategies.

THEOREM 1 (LENGTH OF TRAJECTORIES IN GRAPHS). *Let G be a graph of max degree Δ . Suppose that coefficients composing the agents’ utility functions have been drawn independently from \mathcal{D}_r . Then the expected length $E[len]$ of a resource trajectory in the graph is bounded by Δ^2 .*

PROOF. We know from Lemma 1 that Δ^2 constitutes an upper bound on the length of trajectories for trees. Thus, we just need to show that it remains valid in the case of graphs, whatever strategy is being used. Suppose resource r has reached vertex x_j , and that it possess a common neighbour, say k , with a vertex x_i that the resource previously

passed (that is, $i < j$). Clearly, as x_i occurs before x_j in the trajectory, $x_i \prec x_j$. Also suppose the trajectory of r did not go through k . How does this knowledge affect the probability that the trajectory stops at x_j ?

Let us first consider *blind* strategies. Remember that blind strategies do not necessarily imply that an agent knows all its neighbours valuations before passing the resource on. Then we have two cases: either (i) agent x_i checked whether a rational deal was possible with k , which was not the case. This entails that $k \prec x_i$. Since $x_i \prec x_j$, it follows that $k \prec x_j$ by transitivity. Thus, x_j has at least two incoming arcs (x_{j-1}, x_j) and (k, x_j) . In this case the probability that $x_j = x_{j+1}$ becomes at least $\frac{1}{d_j-1}$, which is higher than the $\frac{1}{d_j}$ obtained in the lemma. Or (ii) k was not picked before, then x_j does not say anything about x_i . In both cases, the bound of Lemma 1 still applies.

Now we consider the *max* strategy. Then in this case, x_i necessarily checked k before selling the resource to x_{i+1} . As it was not selected, the agent went for another vertex, x_{i+1} such that $k \prec x_{i+1}$. Because x_j occurs later in the trajectory, it must be that $x_{i+1} \prec x_j$, and then $k \prec x_j$. The case is similar to case (i) studied before, and the bound applies. \square

This concludes this section on the length of goods' trajectories. A couple of remarks are in place here. First, the bound obtained is only dependent on the (max) degree of the graph (and not on the number of agents). Secondly, the result is fairly general in that it holds even if the distribution from which the utilities are drawn is different for each resource. The third important point to stress is that the bound obtained in the case of trees continues to hold in the general case. This is an important element, especially if when we consider that there are good hopes to improve the value of this bound. Finally, this upper bound for the expected length holds for all blind strategies, and the natural "utility maximizing" strategy. Note that the argument developed in the proof of the theorem may not hold for some very specific strategies (e.g. one may imagine a heuristic consisting instead in taking the *min* among potential buyers).

4. EXPECTED OUTCOMES

From the result obtained in the previous section, we are now going to derive a result regarding the quality of final allocations. Surely, the final state of the system (the final allocation) must be influenced by the length of the goods' trajectories. But this may be less obvious than it may seem at first sight. First, the result only really says resources are expected to run a given number of steps, but we do not know anything about their repartition (where the resource will end up). Secondly, it is not straightforward to see how the chosen social welfare measure will vary as a function of the length of these trajectories.

To put things more dryly, while the previous section gives $E[len]$, we want to determine $E[sw]$. Now we discuss what can be concluded for different social welfare measures. The following result, due to Jensen [13], will be crucial in the discussion.

PROPOSITION 1. *For any convex (resp. concave) function $\varphi(x)$, it is the case that $\varphi(E[x]) \leq E[\varphi(x)]$ (resp. $\varphi(E[x]) \geq E[\varphi(x)]$).*

4.1 Utilitarian Social Welfare

We start with the case of the utilitarian measure of social welfare—the sum of the agents' utilities. So we need to determine the distribution of the agents' utilities at the end of the process, when each resource gets to the end of its trajectory. Intuitively, the longer the trajectory, the higher the value should be. We make this idea more precise in the following lemma. Recall that \mathcal{D}_r refers to the distribution from which the coefficients associated with resource r have been drawn.

LEMMA 2. *Consider a single resource r , whose trajectory length is len . Let $\{X_1, \dots, X_k\}$ be a set of k independent random variables, each distributed among \mathcal{D}_r , where $k = \Delta \times len$. The expected final social welfare is bounded by the expected value of the *max* of these variables. More formally:*

$$E[sw|len] \leq E_{X_1 \dots X_k \sim \mathcal{D}_r} [max(X_1, \dots, X_k)]$$

PROOF. (sketch) The proof argument is best explained by considering each resource as an autonomous entity. Under this perspective, each resource can be seen as "probing" a number of vertices, seeking to get to a vertex offering an increase of its utility. Now at each step, the resource will probe in the worst case Δ neighbour vertices. At the end of its trajectory of len steps, the resource has probed at most $len \times \Delta$ vertices, and the value assigned to it by the final vertex is clearly bounded by the *max* of all probed vertices. It is important to notice that the value of each newly probed vertex is independent from the ones probed previously. \square

The expected value of the *max* over k variable depends on the underlying distribution of these variables. For example, it is known that if these variables are uniformly distributed over the interval $[0, 1]$, then the expected *max* is $\frac{k}{k+1}$. Formally, the following corollary holds:

COROLLARY 1. *Given a single resource crossing len arcs, given that agents' utility coefficients are drawn from a uniform distribution on the interval $[0, \alpha_{max}]$, the expected final social welfare $E[sw | len]$ is bounded by $\alpha_{max} \times \frac{len \times \Delta}{len \times \Delta + 1}$*

Depending on the class of distribution, other corollaries could be derived. *Extremal value theory*, a subfield of probability theory studying the distribution of the *max*, would be highly useful for this task. In particular, a seminal result from Fisher and Tippet [11] states that if $k \rightarrow \infty$, the *max* of k random variables drawn from *any* non-degenerate distribution is itself distributed over a particular parametrized distribution, called the *Generalized Extreme Value Distribution* (GEV). A deeper discussion of this point is beyond the scope of this paper.

We can now plug together Lemma 2 and Theorem 1, using Jensen's inequality, in order to bound the expected social welfare.

THEOREM 2. *Suppose we restrict ourselves to having a single resource. Suppose also that $\varphi(len)$ is an upper-bound on $E[sw | len]$, the expected social welfare considering only that resource, which has moved across len arcs. If φ is non-decreasing and concave, then $E[sw] \leq \varphi(\Delta^2)$.*

PROOF. We will show now that Jensen's inequality allows us to upper-bound the final expected social welfare $E[sw]$. First recall that, by the law of total expectation, we

	a_1	a_2
$\{r_1, r_2\}$	1	5
$\{r_1\}$	1	0
$\{r_2\}$	0	1

Table 1: Non-modular utilities can induce loops

can write $E[sw] = E[E[sw | len]]$. Applying the $E[\cdot]$ operator on each side of the inequality $E[sw | len] \leq \varphi(len)$, we get: $E[E[sw | len]] \leq E[\varphi(len)]$. Thus, $E[sw] \leq E[\varphi(len)]$. Because φ is concave and by application of Jensen’s inequality, we get $E[sw] \leq E[\varphi(len)] \leq \varphi(E[len])$. Because φ is non-decreasing, we can apply the upper-bound on $E[len]$, which becomes: $E[sw] \leq \varphi(\Delta^2)$. \square

It is straightforward to generalize this last theorem to the case where $m > 1$ resources are available. Clearly, resources are non interacting, and the bound presented in this section applies for each of them independently. Thus, given m resources, we have $E[sw] \leq m \times \varphi(\Delta^2)$

Suppose all coefficients are drawn uniformly from $[0, 1]$. Then, applying the above bounds, we get $E[sw] \leq \frac{\Delta^3}{\Delta^3+1}$. As an example, suppose the graph is a chain of any length of max degree 2 where 10 resources are negotiated. In that case, $E[sw] \leq 10 \times \frac{8}{9} \simeq 8.8$.

To conclude, note that not any measure of social welfare can be related in this way to the result of Theorem 1. Take the case of egalitarian social welfare. The length of the trajectory alone is certainly too simple a notion to give interesting results: what we need to know is exactly the repartition of the items at the end of the process, even maybe which agents are crossed along the item trajectory. This requires more refined notions of the system dynamics.

5. THE NON-MODULAR CASE

In this section, we propose a preliminary discussion of the non-modular case. Remember the assumption of modularity has been crucial so far, for it allowed us to take for granted that resource trajectories were completely independent. When synergies between resources occur, either in a positive (super-modularity), or in a negative way (sub-modularity), the trajectories of goods may influence each other. More precisely, trajectories will be influenced by the current allocation of resources. For instance, a resource may pursue its walk across an agent only if this agent holds a particular resource. These interactions between trajectories have multiple consequences. First, it may now be possible to find *loops* in the trajectories, in the sense that goods may cross several times the same agent, even if only deals involving a single resource are permitted.

Let us illustrate the existence of loops by an example. Let two agents a_1 and a_2 negotiate over two resources. Table 1 describes their utility functions. Note that a_2 uses a non-modular utility function: the value assigned to bundle $\{r_1, r_2\}$ is higher than the sum of values assigned separately to $\{r_1\}$ and $\{r_2\}$. Suppose the initial state is the allocation which gives resource r_1 to agent a_2 , and resource r_2 to agent a_1 . Now suppose for instance that we use a blind fixed order strategy: a_2 is designated to start the process, and gives resource r_1 to a_1 . In the next round resource r_2 makes the opposite journey, from a_1 to a_2 . But then r_1 is bound to

come back to a_2 since this is the only remaining rational deal. Hence it is possible that loops occur in the trajectories of goods when valuations are not modular. Note that allocations still cannot occur twice. This naturally entails that trajectories are always finite.

5.1 Non-Interacting Trajectories

In the previous section, we obtained a bound on the expected length of the walk of a resource, for the case of modular utilities. The goal of this section is to generalize this result to the non-modular case. It has been shown in [12] that *any* utility function v can be uniquely represented as a sum over a set of coefficients $\{\alpha_X | X \subseteq \mathcal{R}\}$ in the following way: $v(R) = \sum_{X \subseteq R} \alpha_X$. This entails that utility functions can be uniquely decomposed the following way: $v(R) = \sum_{r \in R} \alpha_r + \sum_{X \subseteq R, |X| > 1} \alpha_X$, where the first part is a modular utility which we will call v^{mod} , and the second is a non-modular function valuing single resources at zero, which we will call v^{syn} .

Let us first identify in which situation the synergetic part of the utilities plays a role. It obviously does when several resources are with the same agent, but it also plays a role when two agents connected by an edge each own a resource. Indeed one of these two may consider contracting a deal with the other one which would result in the possession of several resources, thus activating the synergetic part of its utility. To summarize, if during the whole negotiation process, resources never happen to be on the same agent or on the same edge, v^{syn} plays absolutely no role. In that case, we say that trajectories are *non-interacting*. Note that if the number of resources is low ($m \ll n$), we might expect most trajectories to be non-interacting.

More formally, let us consider a distribution \mathcal{D} among the class of *all* possible utility functions. Suppose the n utility functions are drawn from \mathcal{D} . If trajectories are non-interacting, for each agent i , only the modular part of the utility functions v_i^{mod} plays a role (as $v_i^{syn} = 0$) and the behavior of the system would be exactly that of the modular case, with the distribution \mathcal{D}_{mod} induced by \mathcal{D} .

5.2 Interacting Trajectories

Now we turn our attention to the case where some trajectories interact. We start with a discussion involving only two resources, represented by a square (\square) and a triangle (\triangle), respectively, in Figure 5. In all four cases, we assume that: utilities have been drawn randomly, vertex 0 holds the square, and the arcs represent the possible moves of \square only.




Figure 5: Disturbances in Items Trajectories when Utilities are not Modular

Both cases (a) and (a’) represent the modular cases. As discussed before, the fact that vertex 1 holds the triangle

in case (a') has no influence on the possible moves of the square (arcs in (a) and (a') are oriented the same way).

Case (b) represents the super-modular case: a positive synergetic value is added on vertex 1 when it possesses both goods. Here, vertex 1 will act as an attractor: it increases the utility gain (hence the thicker arrow), but once the square is on vertex 1, it makes it less likely to continue its path, and may even provoke the end of the trajectory, as illustrated in the figure.

Case (c) represents the sub-modular case, meaning that a negative synergetic value is added on vertex 1 when it possess both goods. The analysis is the opposite of the previous case: vertex 1 sort of pushes away the resource. It is less likely that the square will reach it (and may not, as pictured). But once it is on vertex 1, it has more chances to continue its route towards vertex 2.

The following result establishes more formally the consequences on the path length of the disturbances resulting from synergetic effects among goods. It concerns only interactions between two trajectories of items.

LEMMA 3 (LENGTH OF INTERACTING TRAJECTORIES). *Let G be a graph of max degree Δ . Suppose the agents' utility functions have been drawn independently from a distribution \mathcal{D} over the class of all utility functions (not necessarily modular). Then, the expected length of any resource trajectory is bounded by $\Delta^2 \times \left(\frac{\Delta}{\Delta-1}\right)^c$, where c is the number of interactions occurring on this trajectory.*

PROOF. Consider a resource r , for whose trajectory we want to bound the expected length. As in Lemma 1, let x_0, x_1, \dots represent the sequence of vertices visited by r . Let us suppose that, if r goes through more than i arcs, then it will cross another resource r' located on vertex x_i . Formally, this can be represented as an event $C = (len < i) \vee (r' \text{ is on } x_i)$. How does this affect the probability $Pr[len]$? Recall the formula of Lemma 1: For any $k > i$, we can write

$$\begin{aligned} Pr[len = k | C] &= Pr[x_0 \neq x_1 | C] \times \dots \\ &\quad \times Pr[x_{i-1} \neq x_i | x_{i-1} \neq x_{i-2} \wedge C] \quad (1) \\ &\quad \times Pr[x_i \neq x_{i+1} | x_i \neq x_{i-1} \wedge C] \quad (2) \\ &\quad \times \dots \\ &\quad \times Pr[x_k \neq x_{k+1} | x_k \neq x_{k-1} \wedge C] \end{aligned}$$

Note only terms (1) and (2) are affected by the event C . There are then two cases to consider, depending on the type of utility of agent i .

- **(1st case)** suppose the utility function v_i is super-modular. Then, $v_i(\{r, r'\})$ is higher than $v_i(\{r\})$. Thus, the probability that $v_i(\{r, r'\}) > v_{i-1}(\{r\})$ is higher than that of $v_i(\{r\}) > v_{i-1}(\{r\})$. Thus, we will simply bound $Pr[x_{i-1} \neq x_i | x_{i-1} \neq x_{i-2} \wedge C]$ by 1 from above. On the contrary, the probability that $v_{i+1}(\{r\}) > v_i(\{r, r'\})$ will be decreased, compared to $v_i(\{r\}) > v_{i-1}(\{r\})$. We will thus bound it by $\frac{d_i-1}{d_i}$. Thus, we will have:

$$(1) \times (2) \leq \frac{d_i-1}{d_i} \leq \frac{\Delta-1}{\Delta}$$

- **(2nd case)** suppose now v_i is sub-modular. Then, $v_i(\{r, r'\})$ is lower than $v_i(\{r\})$. Thus, the probability that $v_i(\{r, r'\}) > v_{i-1}(\{r\})$ is lower than that of $v_i(\{r\}) > v_{i-1}(\{r\})$. We will bound it by $\frac{d_i-1-1}{d_i}$. On the contrary, the probability that $v_{i+1}(\{r\}) >$

$v_i(\{r, r'\})$, which will be increased, will be bounded by 1. Finally, we have:

$$(1) \times (2) \leq \frac{d_{i-1}-1}{d_{i-1}} \leq \frac{\Delta-1}{\Delta}$$

Thus, in both cases, we can write (by using the bound of Lemma 1).

$$\begin{aligned} Pr[len = k | C] &\leq \frac{\Delta}{\Delta+1} \times \left(\frac{\Delta-1}{\Delta}\right)^{k-1-1} \\ &= \left(\frac{\Delta-1}{\Delta}\right) \times Pr[len = k] \end{aligned}$$

Now consider the general case, where in fact c interactions may occur with r at different locations of its trajectory. Repeating the above argument many times, we can infer the following bound:

$$Pr[len = k | c \text{ interactions}] \leq \left(\frac{\Delta}{\Delta-1}\right)^c \times Pr[len = k]$$

Finally, an upper bound on the expected path is

$$\begin{aligned} E[len | c \text{ interactions}] &\leq \left(\frac{\Delta}{\Delta-1}\right)^c \times E[len] \\ &\leq \Delta^2 \times \left(\frac{\Delta}{\Delta-1}\right)^c \quad \square \end{aligned}$$

Note that the higher the degree, the less influential the disturbance resulting from the synergetic effect among goods on the expected path length. Equipped with this result, the next step to make it practically useful will be to evaluate the probability that trajectories are interacting. This will be dependent on different graph properties, as well as on the expected path length. We leave this for future research.

6. RELATED WORK

The problem of computing the length of allocation processes has been studied in different papers [7, 9], but in situations where no topological constraints affect agents interactions (fully connected graphs). More importantly, the question addressed is not that of an (average case) expected length, but more typically a worst-case analysis of the longest negotiation sequences.

The question of determining the expected length of an *increasing subsequence* in a random permutation has been intensely studied in combinatorics [1]. A spectacular result shows that $2\sqrt{n}$ is the expected length of such an increasing subsequence. This is interestingly connected to our problem, but differs in many respects. As increasing subsequences must go from left to right, this induces a specific topology where each agent is simply connected to any agent that follows him in the lexical order. Furthermore, the sequence is really the best you could find. In our case, this may correspond to the expected length of a given resource when guided by an oracle guiding it through the longest trajectory.

More generally, a huge amount of work bounding the dynamics of complex systems as been produced in the past few years. Mainly two types of analysis have been developed: one based on dynamical system theory (for e.g. Lyapunov stability, pole-zero analysis), and another one on stochastic modeling (random walks in graphs, random graphs, statistical physics) [16]. Our approach is more related to the

latter, as the trajectories we are studying have some close connection to the concept of random walks. Nevertheless, the specificity of our framework did not allow us to exploit ready-to-use results from this field.

7. CONCLUSION

This paper initiates the analysis of the dynamics of multi-agent resource allocation. The main technical contributions presented here are: (i) an analysis of the length of good trajectories in a widely studied abstract negotiation framework, both in the modular and (although more tentatively) in the non-modular case, resulting in a general upper bound result in the modular setting; and (ii) a study of expected utilitarian outcomes based on the previous result in the modular setting, and on a function bounding the expected social welfare given a length of trajectory. We emphasize that the approach advocated here is not necessarily asymptotic: the constraint on the path length can be exogenously given (for instance you may know that for practical reasons resources would not move along more than 2 vertices). Simply plug it into your formula and the bound on the expected social welfare remains valid. This illustrates that the methodology proposed is of great modularity. In fact, beyond the technical results presented here, we regard this general methodology as the most valuable contribution of this paper.

We believe there are numerous directions in which this work may be pursued further.

The first obvious technical challenge will be to improve the proposed bound on the path length. In particular, while the generality of the result is one of its main advantages, it may be that the study of more specific graph structures, with well identified properties, could bring about interesting results. One such class are small-world-like graphs, whose properties have been widely studied in recent years.

A second, perhaps more fundamental question, is to characterize what notions of optimality can be approached with the tools presented here. As was briefly discussed earlier in the paper, not all social welfare measures can be straightforwardly related to the length of trajectories. In most cases, a more refined notion of the actual trajectories of goods will be required, where the precise repartition of goods at the end of (or throughout) the process can be appreciated. Note that this is also the same kind of notions that are needed when, in the non-modular case, the number of interacting trajectories is required to complete the analysis.

8. ACKNOWLEDGMENTS

The first and third author of this paper are partially supported by the ANR-funded project “Preference Handling in Combinatorial Domains”.

9. REFERENCES

- [1] D. Aldous and P. Diaconis. Longest increasing subsequences: from patience sorting to the Baik-Deift-Johansson theorem. *Bull. Amer. Math. Soc.*, 36:413–432, 1999.
- [2] Y. Chevaleyre, P. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30, 2006.
- [3] Y. Chevaleyre, U. Endriss, and N. Maudet. Allocating goods on a graph to eliminate envy. In *Proc. 22nd AAAI Conference on Artificial Intelligence (AAAI-2007)*. AAAI Press, 2007.
- [4] S. Clearwater. *Market-based control: A paradigm for Distributed Resource Allocation*. World Scientific Publishing Co., 1996.
- [5] R. Dash, N. Jennings, and D. C. Parkes. Computational mechanism design: A call to arms. *IEEE Intelligent Systems*, pages 40–47, 2003.
- [6] D. Dolgov and E. Durfee. Resource allocation among agents with MDP-induced preferences. *Journal of Artificial Intelligence Research*, 27:505–549, 2006.
- [7] P. E. Dunne. Extremal behaviour in multiagent contract negotiation. *Journal of Artif. Intell. Research*, 23:41–78, 2005.
- [8] P. E. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 164(1–2):23–46, 2005.
- [9] U. Endriss and N. Maudet. On the communication complexity of multilateral trading. In *Proc. AAMAS-2004*. ACM Press, 2004.
- [10] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating socially optimal allocations of resources. *Journal of Artif. Intell. Research*, 25:315–348, 2006.
- [11] R. A. Fisher and L. H. Tippett. On the estimation of the frequency distributions of the largest or smallest member of a sample. *Proc. Cambridge Phil. Soc.*, 24:180–190, 1928.
- [12] M. Grabisch. k -order additive discrete fuzzy measures and their representation. *Fuzzy Sets and Systems*, 92:167–189, 1997.
- [13] J. Jensen. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Math.*, 30:175–193, 1906.
- [14] S. Kraus. *Strategic Negotiation in Multiagent Environments*. MIT Press, 2001.
- [15] H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.
- [16] H. V. D. Parunak and S. Brueckner. Entropy and self-organization in multi-agent systems. In *Proc. AGENTS-2001*. ACM Press, 2001.
- [17] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.
- [18] T. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proc. AAAI Spring Symposium: Satisficing Models*, 1998.
- [19] R. G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Trans. Comput.*, C-29(12):1104–1113, 1980.
- [20] M. P. Wellman, A. Greenwald, and P. Stone. *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*. MIT Press, 2007.
- [21] M. Wooldridge. *An Introduction to MultiAgent Systems*. John Wiley and Sons, 2002.