Infinitesimal Nash Transfers for Resource Allocation in Strong Social Alliances

Paul-Amaury Matt and Francesca Toni
Department of Computing, Imperial College London
South Kensington Campus, London SW7 2AZ, UK
{pmatt,ft}@imperial.ac.uk

ABSTRACT

We introduce a distributed and tractable mechanism for the allocation of continuously divisible resources to agents, that locally maximises the (Nash) product of their individual welfare. The mechanism involves specific m-resources-at-a-time multilateral deals over bits of resources, termed *infinitesimal Nash transfers*. It provides an effective way of building "strong social alliances", where in a *social alliance* agents fully cooperate for the global interest of society, and a *strong* social alliance has near-optimal utilitarian and egalitarian social welfare, as understood in social choice and welfare economics. The mechanism is scalable, can be distributed amongst agents and can be used to support, e.g., fair trade.

1. INTRODUCTION

Negotiation over the distribution of resources amongst agents is an important area of research in multi-agent systems [2]. We consider the problem of allocating (fractions of) $m \geq 1$ continuously divisible resources $r_1, \ldots, r_m$ to an alliance of $n \geq 1$ agents $a_1, \ldots, a_n$. We define an *allocation* as a matrix $((A_{i,j}))$ of n lines and m columns, with coefficients $A_{i,j} \in [0,1]$ and such that $\forall j \in \{1,\ldots,m\}, \sum_{i=1}^{n} A_{i,j} = 1$. Intuitively, $A_{i,j}$ represents the fraction of resource r_j allocated to agent a_i. Also, in an allocation all resources are fully used and there are no leftovers.

We define the welfare $w_i(A)$ of agent a_i for an allocation A as: $w_i(A) = c_i + \sum_{j=1}^{m} A_{i,j} u_{i,j}$, where $u_{i,j} \in \mathbb{R}^+$ is the utility of resource r_j for a_i and $c_i \in \mathbb{R}^+$ is welfare of agent a_i prior to any allocation of resources.

Given an allocation A, the *egalitarian social welfare* of A is given by $sw_e(A) = \min_{i=1}^{n} w_i(A)$ and the *utilitarian social welfare* is given by $sw_u(A) = \sum_{i=1}^{n} w_i(A)$. Egalitarianism is traditionally associated with the values of fairness and equity [7] and utilitarianism with happiness (cf. J. Bentham). We aim at computing *strong social alliances*, whereby agents share continuously divisible resources according to an allocation quasi-maximising the utilitarian and egalitarian social welfare. Strong social alliances combine the advantages of utilitarianism and egalitarianism and thus attempt to reconcile these traditionally opposed models.

We provide a mechanism for computing strong social alliances within multi-agent systems by maximising their *Nash product*, defined, for a given allocation A, as $N(A) = \Pi_{i=1}^{n} w_i(A)$. Maximisation of the Nash product is a good candidate for providing strong social alliances as, intuitively, when any agent's welfare is decreased, then so is the Nash product, and when all the agents' welfare is increased, then so is the Nash product. Concretely, we provide a rigorous method for computing, by means of a form of negotiation among agents, allocations (locally) maximising the Nash product. These allocations are computed from scratch (with agents having control of no resources at the beginning of negotiation). The computation starts with a first tentative allocation whereby all agents get an equal share of each resource. Then, iteratively, portions of resources are exchanged amongst (groups of) agents so as to progressively increase the overall Nash product. We term these exchanges *infinitesimal Nash transfers*, and prove that they allow to converge to local maximisers of the Nash product.

The paper is organised as follows. Section 2 introduces our solution to the problem of computing strong social alliances based on infinitesimal Nash transfer. Section 3 develops a technique for the negotiation of these transfers among agents. Section 4 provides an experimental evaluation of the method. Section 5 gives a realistic application of our method, to support fair trade. Section 6 concludes.

2. INFINITESIMAL NASH TRANSFERS

Formally, infinitesimal Nash transfers are defined in terms of "perturbation matrices" with specific properties. Given $\epsilon > 0$, a *perturbation matrix* Ψ^ϵ is a real matrix of n lines and m columns with norm $\|\Psi^\epsilon\|_1 \leq \epsilon$, where $\|\Psi^\epsilon\|_1 = Max\{\sum_{i=1}^{n} |\Psi^\epsilon_{i,j}| : j = 1,\ldots,m\}$ (thus $\sum_{i=1}^{n} |\Psi^\epsilon_{i,j}|$ are all bounded by ϵ). The quantity ϵ provides a measure of the *amplitude* of the perturbation represented by the matrix.

Given a perturbation matrix Ψ^ϵ and an allocation A, $A + \Psi^\epsilon$ is the *perturbed allocation*. The first-order development of $N(A + \Psi^\epsilon)$ can be written: $N(A + \Psi^\epsilon) = N(A) + \nabla N(A) \cdot \Psi^\epsilon + o(\epsilon)$ where $o(\epsilon)$ is a function that converges to 0 when ϵ converges to 0. The computation of the partial first-order derivatives of $N(A)$ (needed to compute the gradient) is straightforward: $\frac{\partial N(A)}{\partial A_{i,j}} = N(A)\frac{u_{i,j}}{w_i(A)}$. The quantities $\frac{u_{i,j}}{w_i(A)}$ determine the potential for an agent to maximise the Nash product and thus the priority according to which agents should receive bits of resources in a transfer. Given an allocation A, the *allocation priority* for a_i with respect to r_j, is defined as $\tau_{i,j}(A) = \frac{u_{i,j}}{w_i(A)}$.

The Nash product $N(A + \Psi^\epsilon)$ of a perturbed allocation can be rewritten in terms of allocation priorities as follows: $N(A + \Psi^\epsilon) = N(A)(1 + \sum_{i=1}^{n}\sum_{j=1}^{m} \tau_{i,j} \cdot \Psi^\epsilon_{i,j}) + o(\epsilon)$. Then, an *infinitesimal Nash transfer* is a $n \times m$ matrix $\Psi^\epsilon(A)$ maximising

AAMAS'07, May 14–18, 2007, Honolulu, Hawai'i, USA.

978-81-904262-7-5 (RPS) ©2007 IFAAMAS

$\sum_{i=1}^{n}\sum_{j=1}^{m}\tau_{i,j}(A)\Psi^{\epsilon}_{i,j}$ and such that $\|\Psi^{\epsilon}(A)\|_1 \leq \epsilon$, and $A + \Psi^{\epsilon}(A)$ is an allocation.

The Nash product of an allocation can be iteratively increased using infinitesimal Nash transfers of decreasing amplitude, as outlined in algorithm 1.

Algorithm 1 Approximation A_K of a local maximiser of the Nash product. Input: $\epsilon_{initial}$, number of steps K.

```
$A_K \leftarrow ((\frac{1}{n}))_{n\times m}$; $N_K \leftarrow N(A_K)$ ; $\epsilon_K \leftarrow \epsilon_{initial}$ ; $k \leftarrow 0$
repeat
    $A \leftarrow A_K + \Psi^{\epsilon_K}(A_K)$
    if $N(A) > N_K$ then
        $A_K \leftarrow A$
        $N_K \leftarrow N(A)$
    else
        $\epsilon_K \leftarrow \epsilon_K/2$
    end if
    $k \leftarrow k + 1$
until $k = K$
return $(A_K, \epsilon_K)$
```

THEOREM 1 (LOCAL MAXIMUM). *$(N(A_K))_{K\geq 1}$ converges to a local maximiser of N.*

PROOF. Simply note that the gradient of the Nash product is $\nabla N(A) = N(A).((\tau_{i,j}))_{n\times m}$. Hence, the algorithm performs a gradient ascent of the Nash product in the space of allocations. ☐

When the coefficients c_i and $u_{i,j}$ are generated by a uniform distribution over the interval $[0, 1]$, a good heuristic consists in taking $\epsilon_{initial} = 5/n$ and $K = 20$. For these values, we observed experimentally that the Nash product obtained is usually above 98% of its limit value, even when n and m are larger than 100.

3. COMPUTATION OF TRANSFERS

The columns of $((\Psi^{\epsilon}_{i,j}(A)))_{n\times m}$ can be computed independently of each other, i.e. the agents may reason about one resource (r_j) at a time. We will focus on the computation of column j.

For resource r_j, we rank the n agents in decreasing order of the allocation priority $\tau_{i,j}(A)$: let $G_1, G_2, \ldots, G_P$, denote the unique partition of $\{a_1, a_2, \ldots, a_n\}$ such that all the agents in the same element G_k of the partition have exactly the allocation priority $\tau_{i,j}(A) = \tau_k$ and such that $\tau_1 > \tau_2 > \ldots > \tau_P$. $G_1, G_2, \ldots, G_P$ are referred to as the *priority groups*, where G_1 is the highest and G_P the lowest priority group.

Bits of resources must be transferred from the groups with lowest priority to the ones with highest priority, starting from $\Psi \leftarrow ((0))_{n\times m}$, accessible to all agents in all groups. We use variables p^+ and p^- to denote the indexes of the groups that respectively currently receive (highest priority) and give (lowest priority) bits of r_j. Initially, the agents in G_1 receive and the agents in G_P give (bits of) r_j. Hence, we initialise p^+ with value 1 and p^- with value P. Then, while $p^+ < p^-$, we execute a loop in which we transfer as much of the resource r_j as we can from agents in the group G_{p^-} to the agents in the group G_{p^+}, as follows.

At each step, the quantity (of r_j) transferred is equal to the minimum between two quantities q^+ and q^- defined as:

- $q^+ = \min(\frac{\epsilon}{2} - \sum_{i=1}^{n} pos(\Psi_{i,j}), \sum_{i:a_i\in G_{p^+}}(1 - A_{i,j}))$
- $q^- = \min(\frac{\epsilon}{2} - \sum_{i=1}^{n} neg(\Psi_{i,j}), \sum_{i:a_i\in G_{p^-}} A_{i,j})$

where $pos(x)$ and $neg(x)$ are equal respectively to x if $x > 0$ and 0 otherwise, and $-x$ if $x < 0$ and 0 otherwise. Intuitively, q^+/q^- is the maximum quantity group G_{p^+}/G_{p^-} can take/give (without violating any constraint). Three different scenarios may occur:

1. If $q^+ < q^-$, then G_{p^-} transfers q^+ to G_{p^+} and p^+ is increased by 1.
2. If $q^- < q^+$, then G_{p^-} transfers q^- to G_{p^+} and p^- is decreased by 1.
3. If $q^- = q^+$, then G_{p^-} transfers $q^- = q^+$ to G_{p^+}, p^- is decreased by 1 and p^+ is increased by 1.

The process terminates when $p^- \leq p^+$, after $n-1$ steps at most.

The total quantity of a resource received by or taken from a group must by split amongst the agents that compose it. Bits of resource should be taken from or given to the agents in a group with equity, as these agents all have the same allocation priority. This can be done iteratively. At each step, we try to give to or take from each agent in the group the remaining quantity to give or take divided by their number. Each agent stops giving when it does not have anything left and stop receiving when it owns a resource entirely. At the end, the remaining quantity to give or take is equal to 0. This process terminates after n steps at most.

THEOREM 2 (TRACTABILITY). *The overall complexity for computing a transfer is $O(m \cdot n^2)$.*

This makes the mechanism tractable. On a 1GHz processor for example, an instance of the problem with 50 agents and 100 resources is solved within roughly 10 minutes.

4. EMPIRICAL EVALUATION

Here, we restrict attention to the case of *binary* allocations (where $A_{i,j}$ is either 0 or 1 and thus resources are treated as indivisible). Our aim is to build *strong* social alliances, with quasi-optimal utilitarian and egalitarian social welfare. The computation of egalitarian allocations (with optimal egalitarian social welfare) is computationally demanding in general, and in the case of indivisible resources has an exponential complexity (it is in general NP-hard [1]). However, egalitarian allocations (and the optimal value sw^*_e of the egalitarian social welfare) can be computed effectively [5, 4] in small dimensions. The computation of utilitarian allocations (with optimal utilitarian social welfare) can be simply obtained by allocating each resource r_j to the agent a_i that maximises $u_{i,j}$ (thus also giving binary allocations).

Figure 1 shows the result of an experimental evaluation of the egalitarian (second column) and utilitarian (fourth column) social welfare of the allocation A computed by our technique, compared with the optimal egalitarian (third column) and utilitarian (last column) social welfare of binary allocations ($sw^{*\{0,1\}}_e$ and $sw^{*\{0,1\}}_u$, respectively), for various values of (n, m) (first column). Limited by the complexity in the egalitarian case, we restricted our simulations up to $n = m = 8$ and started with $n = m = 2$. We have generated 30 problems with uniform (0,1) distributions and average the performances obtained. The results in figure 1 prove the strength of our method: both measures of welfare are very close to their optimal value.

5. APPLICATION TO FAIR TRADE

Consider a small size European import organisation that wants to develop fairtrade in partnership with producers from developing countries and supermarkets. In practise, producer may be agricultural cooperatives of farmers. The organisation seeks to organise the production of agricultural products in a way that is fair to the producing countries. The countries selected are Ethiopia, Brazil,

(n,m)	$sw_e(A)$	$sw_e^{*\{0,1\}}$	$sw_u(A)$	$sw_u^{*\{0,1\}}$
$(2,2)$	0.811	0.798	1.947	2.003
$(3,3)$	0.965	0.995	3.881	4.020
$(4,4)$	0.932	0.932	4.919	5.161
$(5,5)$	0.913	0.976	6.430	6.707
$(6,6)$	0.934	0.939	7.699	8.060
$(7,7)$	0.978	0.949	8.757	9.106
$(8,8)$	1.098	1.024	11.003	11.318

Figure 1: The egalitarian $sw_e(A)$ and utilitarian $sw_u(A)$ social welfare of A are quasi optimal: $sw_e(A) \simeq sw_e^{*\{0,1\}} \simeq sw_e^*$ and $sw_u(A) \simeq sw_u^{\{0,1\}*}$. The social alliances are strong.

$Origin/Prod.$ $Qty.(tons)$	$cof.$ 1400	$ban.$ 3000	$coc.$ 500	$man.$ 1300	tea 600	$cot.$ 900
$Ethiopia$	1100	0	0	0	0	0
$Brazil$	900	650	1600	750	0	0
$Vietnam$	1000	0	0	500	1000	0
$Indonesia$	1200	500	1900	800	1100	0
$CostaRica$	950	800	0	0	0	0
$India$	0	0	0	900	1500	1200
$China$	0	400	0	850	1700	1350
$Thai.$	0	550	0	800	1300	0
$Mexico$	0	550	0	750	0	0

Figure 2: Fair trade products (coffee, bananas, cocoa, mangoes, tea, cotton) and origin of growers: (fictive) benefits in $ per ton.

Vietnam, Indonesia, Costa Rica, India, China, Thailand and Mexico. Products in which European consumers are interested are coffee, bananas, cocoa, mangoes, tea and cotton.

The organisation needs to plan every year the production of the different products in accordance with the quantities needed by the supermarkets. Traditional marketing techniques can be used for assessing the popularity and need of new products, as well as prices at which these can be sold. The organisation thus has to determine how much of each product each producer has to supply.

The resources of our allocation problem correspond to markets. There are m markets, each corresponding to one product. Markets are characterised by a quantity q_j of product r_j needed by the supermarkets and a fixed or expected buying price per unit (dollars per ton) $p_{i,j}$. The prices $p_{i,j}$ depend on the producers because these typically offer distinguishable products. All agricultural products are divisible and thus suitable for our application.

We view producers as agents, aiming at forming a social alliance. Their initial welfare is the sum over the past years of their benefits, assuming all agents have joined the social alliance at the same time with benefit zero. Each supplier has its own preferences (utilities) concerning the production of different products, because they do not all yield the same benefits. Quite simply, the benefits made are equal to the difference between the selling price and the cost of production. Hence, we adopt the formula $u_{i,j} = q_j.(p_{i,j} - c_{i,j})$ to construct the utilities of the problem (see table 2).

The allocation mechanism tells the organisation how to best plan the production of the product. The solution will satisfy the total supermarkets' demand, be fair from the producers perspective and guarantee quasi-maximal overall benefits for the developing countries. The solution obtained is shown in tables 3, assuming that all the agents start at year 0 with no prior welfare.

O/P	$cof.$	$ban.$	$coc.$	$man.$	tea	$cot.$	W
$Eth.$	0.50	0	0	0	0	0	773.2k\$
$Bra.$	0	0.42	0	0	0	0	817.2k\$
$Viet.$	0.50	0	0	0	0	0	697.1k\$
$Indo.$	0	0	1	0	0	0	950k\$
CR	0	0.42	0	0	0	0	1009.1k\$
$India$	0	0	0	0.10	0.11	0.55	817k\$
$China$	0	0	0	0	0.50	0.45	1058.6k\$
$Thai.$	0	0	0	0.48	0.39	0	795.3k\$
$Mex.$	0	0.16	0	0.42	0	0	674.7k\$

Figure 3: Optimal production plan and expected welfare.

6. CONCLUSION

We have introduced the notion of strong social alliance and argued that it could be realised by computing allocations maximising its Nash product. We have presented a mechanism for computing such allocations in the case of semi-linear welfare functions and continuously divisible resources based on the novel concept of infinitesimal Nash transfer. We have proved that the mechanism is tractable, and illustrated its application to support fair trade.

Maximisation of the Nash product can be defined axiomatically as the solution of a two-agents cooperative game called Nash bargaining [6]. However, Nash bargaining is not concerned with the actual process of finding a solution, but rather, it focuses on axiomatising desirable outcomes. Maximisation of the Nash product in the space of "individually rational" allocations can be obtained using the Zeuthen strategy [3, 8]. We have proposed a mechanism for computing (local) maxima of the Nash product in the space of *all* allocations, including ones that may not be "individually rational" but will nonetheless be beneficial to the overall society. This is needed in applications such as fair-trade.

7. ACKNOWLEDGEMENTS

This work was partially funded by the EC-funded ARGUGRID project. The second author was also supported by a UK Royal Academy of Engineering/Leverhulme Trust Senior Fellowship.

8. REFERENCES

[1] S. Bouveret, M. Lemaître, H. Fargier, and J. Lang. Allocation of indivisible goods: a general model and some complexity results. In *AAMAS*, pages 1309–1310, 2005.

[2] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating socially optimal allocations of resources. *Journal of Artificial Intelligence Research*, 25:315–348, 2006.

[3] J. Harsanyi. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge Univ. Press, 1977.

[4] P.-A. Matt and F. Toni. The distributed negotiation of egalitarian allocations. In *1st International Workshop on Computational Social Choice*, 2006.

[5] P.-A. Matt and F. Toni. Egalitarian allocations of indivisible resources: theory and computation. In *Cooperative Information Agents*, 2006.

[6] J. Nash. Two-person cooperative games. *Econometrica*, 21:1:128–140, 1953.

[7] J. Rawls. *A Theory of Justice*. Harvard University Press, 1971.

[8] G. Zlotkin and J. S. Rosenschein. Negotiation and task sharing among autonomous agents in cooperative domains. In *IJCAI*, pages 912–917, 1989.