Economic Coordination Mechanisms for Holonic Multi Agent Systems

Heimo H. Adelsberger
Wirtschaftsinformatik der Produktionsunternehmen
Universität Essen
45117 Essen
heimo@wi-inf.uni-essen.de

Wolfram Conen
XONAR IT Gmbh
Velbert, Germany
conen@gmx.de

Abstract

This paper discusses the importance of coordination mechanisms for holonic multi agent systems. Specifically, the design of mechanisms for worth-oriented domains with transferable utility is discussed. The specific focus is on how to allocate a set of resources to a set of agents. Agents may express preferences for the complementary use of resources. Agents are assumed to be self-interested. An example demonstrates relevant concepts, such as economic efficiency, equilibrium, and surplus distribution. The importance of a rigid quality assessment is emphasized.

1. Coordination in Holonic Multi Agent Systems

> "Future manufacturing systems need to cope with frequent changes and disturbances. As such, their control requires constant adaptation and high flexibility. Holonic manufacturing is a highly distributed control paradigm that promises to handle these problems successfully. It is based on the concept of autonomous co-operating agents, called *holons*." [18]

A holon is an autonomous and co-operative building block of a manufacturing system. The concept of autonomity emphasises the capability of a holon to create and control the execution of its own plans and strategies. It is, however, embedded in a holarchy, that is a system of holons that can co-operate to achieve goals or objectives (compare [18, 17]).

A holon may consists of a physical and an informational component. The informational component contains information, has decision-making capabilities, and manages the physical component and the cooperation with other holons [7]. We will concentrate on the informational component in the following and we will frequently call it *agent*.

Gou et al. emphasise the importance of localised information and decision-making authority. However, they also postulate that individual elements (holons) should not operate with absolute autonomy, that is, they should cooperate with others, observing the system constraints and acting according to coordination information [7].

We will view *rational cooperation* as joint acting with the goal to improve the situation of every cooperating entity. If the outcome of the joint acting will worsen the situation of some of the "cooperating" entities, we will call this *altruistic cooperation* [1].

The exchange of information (matter, energy) is interpreted as *coordination* if it enables and/or performs a joint action (a sequence of activities). Note that coordinating the efforts of actors does not necessarily mean that the result is mutually beneficial. It is clear that coordination is a constituting element of when cooperating.

The holons act according to strategies and derived goals. The consequences of strategies, be it in cooperative or non-cooperative situations, depend on the coordination mechanism deployed. From a technical perspective, a coordination mechanism is necessary to transport the information to and from the participating holons. From the perspective of cooperation, a certain protocol is required to allow for the coordination of activities or the adaptation of goals. Sometimes a dedicated holon is introduced to solve a coordination problem. This institution will be part of the coordination mechanism.

With *coordination mechanism* we will denote the message protocol (abstract and instantiated), the individual computations (decisions, goal, and strategy formation), and the coordination institution (if any is deployed).

The design of a coordination mechanism is a crucial task for any designer of cooperative systems: the quality of

[1] Only a cooperation that will predictably lead to an improvement of every participant's situation under any future circumstances is guaranteed to be *ex post rational*. The same is true for cooperations that will always and predictably affect the situation of one or more participants negatively. Only such cooperations are guaranteed to be *ex post altruistic*

Proceedings of the 11th International Workshop on Database and Expert Systems Applications (DEXA'00)
0-7695-0680-1/00 $10.00 © 2000 IEEE

the solutions obtained does mainly depend on the deployed mechanism.

In the following, we will concentrated on a specific class of coordination problems, namely problems with self-interested, rational agents that act in a worth oriented domain [14] and want to decide upon the distribution of resources. We think that such situations are rather general and frequent, especially if autonomity is a key aspect of an analyzed systems, such as it is the case with with holonic systems, where private information, economic autonomity, and individual goals are important (supply chains, segments, workshops, humans, groups of humans) and incentives have to be given to ensure cooperative behavior of the holons.

2. Economic Coordination Mechanisms

> "Cooperation in the economic tradition is mutual assistance between egoists." [12]

Viewing cooperation as assistance among egoists leads directly to the so called *efficiency postulate*: Efficiency must prevail among rational agents–they will cooperate if by doing so they both benefit (compare [12]).

Besides core analysis, which tries to analyse if stable allocations are reachable by mutual agreements among agents[2], equilibrium analysis is a key to understand the computation of efficient outcomes in rational cooperations. This concept is presented in the following subsection.

2.1. (Walrasian) Equilibria

Here, a mechanism is established that tries to distribute resources to agents by means of prices and based on individual decisions: given a certain price vector for a set of goods (or a set of bundles of goods), each agent has to decide on its own which trade will maximise his utility and has to act accordingly, that is to sell or buy the products he chooses to trade. If the economy is not closed, and thus trading can not rely solely on barter, the presence of a meaningful currency is required.

Again, the outcome of such a mechanism should exhibit certain properties: the outcome should be efficient (1), i.e., it should maximise the aggregated surplus; participation in trades should be voluntarily and the trades should be individually rational (2), that is each agent should be able to participate in the trades that maximise his individual utility with respect to the given prices.

[2] In short, stability requires that no coalition of agents can generate a higher benefit from trade among themselves. We say that an allocation is a *core allocation* if (1) the trade required to implement the allocation is voluntarily, (2) the allocation is Pareto efficient, and (3) no coalition of agents can find a better (individually judged!) trade of its own (compare [12]).

We explicitly distinguish between agents wanting to consume resources and agents offering resources. We assume to be in a discrete world with indivisible resources and we maintain the assumption that a transferable currency exists.

The consumers, $C := \{1, \ldots, n\}$, are interested in purchasing bundles of indivisible objects, $\Omega := \{\omega_1, \ldots, \omega_m\}$, from the suppliers, $S := \{n+1, \ldots, n+k\}$. Any subset $B \subset \Omega$ is called a *bundle*. The set of all bundles is $2^\Omega := \{B | B \subseteq \Omega\}$.

A map $u_i : 2^\Omega \to R$ is called the *reservation value function* of agent i on Ω; it assigns a value to each *bundle* of Ω. Assume $u_i(\emptyset) = 0$ and $u_i(\cdot)$ is weakly increasing. The agents have utility over bundles and (divisible) money, $m \in \mathcal{R}$. Buyer i's (quasi-linear) utility function is $U_i(B, m) = u_i(B) + m$, seller j's utility (or cost) function is $U_j(B, m) = m - u_j(B)$[3].

We say that an allocation of goods from suppliers to consumers is *efficient*, if it maximises the joint surplus of all participating buyers and sellers. This seems to be an agreeable objective for a group of cooperating holons. So, one task is to determine (one of) the efficient allocations.

However, since participation in each trade must be voluntarily, some further requirements have to be fulfilled: Each seller must receive an amount of money that is at least as high as his cost of providing the requested goods and no buyer can be forced to provide more money than the goods he receives are worth (valued according to his reservation value function). This leads to the notion of prices. Right now, we have determined the necessity to exchange money against goods in a buyer-seller pair. If our system is symmetric with respect to flow of information, each agent would want to ask all potential trade partner what transfer of money would be required – only this would allow to optimise the individual decision about the trade partner: a consumer would want to buy from the cheapest seller and each seller would want to sell to the most interested consumer. However, this would leave us in a situation completely analogous to considering core allocations – the individual decisions would be based upon heavy negotiations with all potentially interested partners. If an (anonymous) price vector is computed (by a trusted third party, e.g., a benevolent auctioneer), it can be used by every participating agent to optimise his own trade decision and, if it is ensured that all individual decisions can be realized afterwards, this would perfectly institutionalise the computation of efficient allocations.

This leads to the notion of *competitive equilibria*: Given a vector of prices for the goods to be traded, an outcome[4] is an *equilibrium*, if the bundles that are assigned to buyers and that are sold by sellers maximise the individual net utilities of buyers and sellers.

[3] With $u_j(B) = \infty$ for every bundle that j does not own.

[4] That is an allocation and transfers of money.

The interesting question now is, if, for every efficient allocation, a corresponding equilibrium price vector exists.

This question becomes somewhat hard to answer, if complementarities are possible, that is, if the utility functions are not strictly additive with respect to the bundling of goods.[5] This is regularly the case for holonic manufacturing systems in which the value of utilising a resource depends on the availability of other resources.

In a related paper, [4], we have outlined the concept of consistent equilibria that allow to guarantee the existence of equilibria, if the gross-substitutes condition hold for the bundles to be sold. This can be extended to virtually all problems if a controlled form of collusion is deployed.

In the remainder of the paper, we will consider an example of applying the concepts presented so far to a holonic MAS, present some pointers to related work and discuss the results and prospects of the economic coordination approach.

3. Example

We will analyse a situation with two supplying holons, A and B, and two requesting holon, 1 and 2. The two supplying holons have created a coordination holon, $\mathcal{C}$. The coordination holon collects bids from the requesting holons and reserve value information from the supplying holons. $\mathcal{C}$ searches for an (economically) efficient allocation and tries to compute an equilibrium price vector supporting it.

The following matrix gives the valuations of the requesting agents for the resources. Agent 1 requires two complementary resources to fulfil his needs (hence his utility is 0 for receiving only one resource and 3 for both of them). For agent 2, it does not matter if he receives resource A or B or both of them – he can use both alternatively.

	A	B	AB
Agent 1	0	0	3
Agent 2	2	2	2

It is clear that the efficient, that is the joint-utility maximising allocation, is to assign both resources to agent 1 (creating a joint surplus of 3).

In the first variant, we assume that the reservation values of the resources are 0. To ensure a straightforward distribution of the collected prices to each seller, $\mathcal{C}$ it designed to determine prices for all individual resources.

With this assumption, the following will hold: A price vector that would support the efficient allocation must exclude agent 2 from the allocation. To keep agent 2 satisfied, the resources he receives (the null resource, $\emptyset$, which will always have a price of 0) must optimise his local decision problem. Here, this means that his utility minus the price for the available ("buyable") resources must be equal to or lower than 0. Therefore, $p_A \geq 2$ and $p_B \geq 2$. This, however, will exclude agent 1 from the allocation too: $p_A + p_B = 4$ exceeds his valuation of 3 for the bundle AB. It follows that it is impossible to support the efficient allocation with prices for the individual goods only. A practical approach for the coordinator would be to determine the best supportable allocation. Here, this would simply be to assign arbitrarily one of the resources to agent 2. Let's assume that A will be assigned to 2. Now, the following conditions would have to hold: $p_A \leq 2$ and, to keep agent 1 satisfied with the outcome, $p_A + p_B \geq 3$. In this case, a range of equilibrium price vectors supporting the allocation exist. However, the value of this outcome is only 2.

¿From this analysis follows that it is not very reasonable to allow $\mathcal{C}$ to sell the resources individually only. Note that this simple example has a huge impact on the design of coordination mechanism, such as the FCC auctions [10], in the presence of complementarities, as in the example: If the resources are sold individually, it can not be guaranteed that the economically efficient allocation can be established.

In the second variant, the coordinating holon is allowed to bundle resources, even from different suppliers, and to sell the bundle instead of selling the resources (or goods) individually. In the example, the following conditions for equilibrium prices would have to hold: $p_{AB} \leq 3$ and $p_{AB} \geq 2$. While this is already sufficient to establish the efficient outcome (with the additional condition that only *priced* resources and bundles can be bought), the economic literature mostly determines prices for all bundles and all individual goods. In this case, the following additional conditions are required to keep agent 2 satisfied: $p_A \geq 2$ and $p_B \geq 2$.

However, now the *distribution of the revenue* becomes a little bit difficult. The coordinating holon receives a payment x for the bundle AB from agent 1. This has to be split. The split rule might be part of the design of the coordinating holon (e.g., always split revenue in equal shares) or it may always require individual negations between the supplying holons. Note, that it is rational (with reasonable assumptions) for the suppliers to allow bundling: The individual revenue that can be expected as an outcome of the first variant, which chooses arbitrarily among the resources, is $1/2$ of 1 (if we assume an equal probability distribution over the equilibrium price interval). For variant two, it is $splitfactor_i * 2.5$ for resource i. Applying the equal-share rule would lead to an expected value of 1.25.

A second and even more relevant problem for a group of cooperating holons is the *distribution of the surplus*. In the example, as we have assumed that the reservation value for the resources is 0, the surplus is 3. Depending on the price

[5] If there are no such dependencies, it is known from Kelso and Crawford [9] (see also [8]) that such a Walrasian equilibrium always exist – the gross-substitute condition has to hold.

vector that the coordinating holon chooses, this surplus is distributed differently among buyer and sellers. Due to the effect of competition between agent 1 and 2, the sellers can be sure of receiving at least a surplus share of 2. In fact, it might not be possible to realize any higher share. This can be seen as follows: Assume that agent 1 is informed about the valuations of agent 2 (e.g., by having observed agent 2's behavior in the last coordination procedures or by listening to the transmission of agent 2's bids). In this case, he may choose to submit a dishonest bid: my value for AB is $2 + \epsilon$. If the coordinating holon was designed to search for the efficient allocation, agent 1 would still receive the bundle AB. Additionally, the equilibrium price range is now reduced to $[2, 2 + \epsilon]$. If the pricing rule that is applied depends on the bid of agent 1, he will always have an incentive to lie in comparable situations. Note, that agent 2 can not improve his situation by being dishonest (if we rule out collusion, otherwise, agent 2 may offer agent 1 to bid significantly less and to share the additional surplus). In more complex problem settings, it may become difficult or impossible to find a dominant lying strategy. In such a case also the computation of the efficient allocation may become impossible. So, a prominent design objective for the coordination mechanism is to give the agents no incentive to lie (*incentive compatibility*). Generally, an incentive to behave strategically, to collect information about competitors, and to lie is present, if the pricing rule is designed so that the price for a resource or a bundle depends on the bid of the winner of this resource or bundle (agent 1 in the example above). A well-known incentive compatible pricing rule was developed independently by Vickrey, Clark, and Groves. Here, the winner has to pay (or is compensated) for the *effect of his participation* in the coordination procedure. But please note that in a market with self-interested buyers and sellers, incentive compatibility can generally not be guaranteed[6]: a supplier may use his reservation value to raise the price that he receives. If both suppliers would raise their reservation value from 0 to 1.2, the minimal equilibrium price would be 2.4 (and this would also be a bound for agent 1's lying).

In a purposefully formed cooperation of holons, it would be reasonable to use a pricing rule that distributes the surplus that is available between the minimal and the maximal equilibrium prices equally[7] among supplying and requesting holons. To reduce the incentives to lie, a control procedure can be established that allows a control holon to inspect from time to time the way in which reservation value and utility calculations are performed in arbitrarily chosen holons. Other control procedures are possible but it has to be kept in mind that certain procedures may violate autonomity too much.

Additional problems in designing such a system are related to security and trust. We have already mentioned the possibility to listen to transmissions of bids (encryption!). Also, the coordinating holon has a number of possibilities to manipulate – and by keeping all transmitted data encrypted and, thus, private to the sending holon and the coordinating holon, these possibilities are even enhanced. For example, the coordinating holon may compute different price vectors for sellers and buyers, and keep the difference. So, this coordinating holon, which might as well be realized based on an outsourced service, must be trustable.

This may suffice to introduce the key ingredients of economic and price-based coordination mechanisms: efficient allocations, equilibrium, revenue distribution, surplus distribution, strategic consideration and incentive compatibility, security and trust. Further details of combinatorial auctions can be found in [5, 22, 8, 3, 13].

4. Related Work

The considerations presented in this paper are related to a number of active research areas. We will give some brief pointers to ease orientation.

The problem of complementarities led to a number of research activities in economics and AI. An early approach to solve problems with complementarities is described in [20]. Wellman et al. discuss the problems of trying to apply Walrasian-like economic coordination models of competitive equilibrium in a non-convex domain. Wellman also gives an interesting motivational account of his experience with Market-oriented Programming in [19].

A recent report of Tuomas Sandholm [16] suggests an efficient algorithm for winner determination in combinatorial auctions and gives pointer to further relevant literature. Especially the problem of NP-hardness and weak approximability is discussed. A similar approach is presented in [6].

A comprehensive account of market-based coordination mechanisms along with instructing application examples can be found in [23]. The design space of auction-based coordination mechanisms is explored in [21]. Details of a coalition-oriented analysis are laid out in [15]. In the broader context of coordination mechanism (without requiring the use of a meaningful "currency" as a medium of coordination), much information and numerous pointers can be found in [14].

Previous work on economically motivated coordination mechanism in a manufacturing context is numerous. Mostly, the suggested models are either applied in a simplified context [2] or difficult to interpret in terms of economic analysis [1, 11], or both. However, the approaches often

[6] Only, if the procedure is not budget-balanced, e.g., if the coordinating holon would pay the maximum equilibrium price (3 in the example above) and would receive only the minimum equilibrium price (2).

[7] Or, more general: fair, compare [12]

suggest interesting heuristic approaches and thus may be useful to heuristically obtain a solution for the allocation or price determination problem.

5. Discussion

Applying economic analysis to holonic manufacturing systems has a number of advantages: the quality of coordination mechanisms is assessable within a well-formed framework of concepts and techniques, such as economic efficiency and equilibrium analysis. It emphasises the importance of autonomity and clarifies the notion of cooperation.

However, the applicability is also restricted. In the immediate scope of the analysis are rational cooperations. The analysis is confined to situations where it is reasonable to relate individual objectives to valuations. While this may be possible surprisingly often (ultimately, all activities are tied to cash-flows and the problem of relating cause and effect becomes easier to handle if one tries to make the analysis explicit), it would require appropriate budgeting, accounting and calculation systems of fine granularity [1].

There is a number of open topics. The interplay between distribution of power and the design of coordination mechanisms has to be made more explicit. The analysis of possible classes of problem structures has to be completed. The infrastructure requirements (security, information flow, computational capabilities) have to be analysed more detailed. The possibility for large instances of NP-hard problems (the example above may be interpreted as a job-shop problem, in fact, each (economic) job-shop problem can easily be transferred into the Resource/Valuation formulation used above) makes the use of heuristics imperative. The effect of heuristic allocation algorithms on the incentives for the holons to participate in the coordination procedure or to be honest has to be determined.

Nevertheless, we do believe that applying instruments of economic analysis to holonic multi agent systems will help to improve the design of coordination mechanisms and to promote the necessity of a rigid analysis of the quality of proposed coordination mechanisms.

References

[1] H. H. Adelsberger, W. Conen, and R. Krukis. Scheduling utilizing market models. In *Proc. of the ICCIM 95*, Singapore, 1995. World Scientific.

[2] A. D. Baker. A case study where agents bid with actual costs to schedule a factory. In S. H. Clearwater, editor, *Market-Based Control*, pages 184–223. World Scientific, Singapore, 1996.

[3] S. Bikhchandani and J. Ostroy. The package assignment model. UCLA Working Paper Series, mimeo, 1998.

[4] W. Conen. An efficient iterative sealed-bid combinatorial auction. Submitted for publication.

[5] W. Conen. Economic coordination, bundled goods, and the impact of complementarities. In Parsons and Wooldridge, editors, *Workshop on Decision theoretic and Game theoretic Agents*, London, 1999.

[6] Y. Fujishima, K. Leyton-Brown, and Y. Shoham. Taming the computational complexity of combinatorial auctions. In *Proc. IJCAI 99*, Stockholm, 1999.

[7] L. Gou, P. Luh, and Y. Kyoya. Holonic manufacturing scheduling: architecture, cooperation mechanism, and implementation. *Computers in Industry*, 37:213–231, 1998.

[8] F. Gul and E. Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87:95–124, 1999.

[9] A. S. Kelso and V. P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50:1483–1504, 1982.

[10] J. McMillan. Selling spectrum rights. *Journal of Economic Perspectives*, 8(3):145–162, 1994.

[11] T. E. Morton, S. R. Lawrence, S. Rajagopolan, and S. Kekre. Sched-star - a price-based shop scheduling module. *Journal of Mfg. and Operations Mgmt.*, pages 131–181, 1988.

[12] H. Moulin. *Cooperative Microeconomics*. Princeton University Press, 1995.

[13] D. C. Parkes and L. Ungar. Iterative combinatorial auctions: Theory and practice. In *Proc. of the AAAI-00*, 2000. to appear.

[14] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.

[15] T. Sandholm. *Negotiation among Self-Interested Computationally Limited Agents*. PhD thesis, Computer Science, University of Massachusetts, Amherst, 1996.

[16] T. Sandholm. An algorithm for optimal winner determination in combinatorial auctions. Technical Report WUCS-99-01, Washington University, St. Louis, 1999.

[17] P. Valckenaers, F. Bonneville, H. Van Brussel, L. Bongaerts, and J. Wyns. Results of the holonic control system benchmark at the k.u. leuven. In *Proc. of the CIMAT Conference*, pages 128–133, Troy, N.Y., 1994. Rensselaer Polytechnic Inst.

[18] H. Van Brussel, J. Wyns, P. Valckenaers, L. Bongaerts, and P. Peeters. Reference architecture for holonic manufacturing systems: Prosa. *Computers in Industry*, 37:255–274, 1998.

[19] M. Wellman. Market-oriented programming: Some early lessons. In S. Clearwater, editor, *Market-Based Control*. World Scientific, Singapore, 1996.

[20] M. P. Wellman. A computational market for distributed configuration design. In *Proc. of the 12th National Conference on Artificial Intelligence (AAAI-94)*, pages 401–407, Seattle, WA, July 1994. AAAI, AAAI Press.

[21] P. Wurman. *Market Structure and Multidimensional Auction Design for Cumputational Economies*. PhD thesis, University of Michigan, 1999.

[22] P. Wurman and M. Wellman. Equilibrium prices in bundle auctions. Submitted for publication. To appear as Santa Fe Institute Working Paper, 1999.

[23] F. Ygge. *Market-Oriented Programming and its Application to Power Load Management*. Phd thesis, Lund University, Lund, Schweden, 1998.

Proceedings of the 11th International Workshop on Database and Expert Systems Applications (DEXA'00)
0-7695-0680-1/00 $10.00 © 2000 IEEE