The society of mind requires an economy of mind

Ian Wright
The Cognition and Affect Group
School of Computer Science,
The University of Birmingham,
Birmingham, B15 2TT, ENGLAND
I.P.Wright@cs.bham.ac.uk

Michel Aubé
Faculté d'Éducation,
Université de Sherbrooke,
2500 Blvd Université,
Sherbrooke, Qc, CANADA, J1K 2R1
maube@courrier.usherb.ca

December 3, 1996

Abstract

A society of mind will require an economy of mind, that is multi-agent systems (MAS) that meet a requirement for the adaptive allocation and reallocation of scarce resources will need to use a quantitative universal representation of value that mirrors the flow of agent products, much as money is used in simple commodity economies. The money-commodity is shown to be an emergent exchange convention that serves both to constrain and allow the formation of commitments by functioning as an ability to buy processing power. MAS with both currency flow and minimally economic agents can adaptively allocate and reallocate control relations and scarce resources, in particular labour or processing power. The implications of these views are outlined for MAS research and cognitive science.

1 The society of mind

... a group of agencies inside the brain could exploit some "amount" to keep account of their transactions with one another. Indeed agencies need such techniques even more than people do, because they are less able to appreciate each other's concerns. But if agents had to "pay their way," what might they use for currency? One family of agents might evolve ways to exploit their access to some chemical that is available in limited quantities; another family of agents might contrive to use a quantity that doesn't actually exist at all, but whose amount is simply "computed".

M. Minsky, *The Society of Mind*, 'magnitude and marketplace', page 284.

Marvin Minksy's *The Society of Mind* (Minsky, 1987) is the best example of the social metaphor applied to the understanding and design of minds. It outlines a computational society of heterogenous agents that compete and cooperate to produce mental capabilities. The approach of decomposing a computational mind into a society of less intelligent agents is compelling because social systems and large, parallel computing systems share design features. (Simon, 1981) discusses how generic design solutions in one type of system can be usefully transferred to the design or analysis of another. For example, both kinds of system consist of a set of mutually connected, interacting subcomponents or agents that are able to perform work, for example computational units that process or people that labour. Such agents can function as both producers and consumers, for example the input and output of information or the consumption and production of commodities. The agents operate within a social division of labour, that is a functional specialisation of the subcomponents of the system, which affords fine-grained parallelism, that is the subcomponents function relatively autonomously and concurrently perhaps pursuing their own local goals. Both types of system need to be coordinated by mechanisms for the production, distribution and consumption of agent products, such as globally accessible databases or free market mechanisms. In addition, such systems must adaptively allocate scarce resources, be they limited labour resources or processing time or commodities or information in restricted supply.

It is these kinds of considerations that suggest that the social 'metaphor' is no metaphor at all, but is a partial identity between a class of complex systems at the information processing level of abstraction, just as the rules that govern genetic evolution have parallels in memetic evolution (Dawkins, 1990). However, as with all compelling parallels it is important to identify differences as well as similarities. There are many ways in which the evolution of ideas is different from the evolution of genes: for example, ideas are often subjected to rational criticism along dimensions such as probable truth and consequences for others and self. Human societies and computational societies will also differ in important ways and it would be a mistake to construct computational theories dominated by current ideas about social organisation without exploring the full range of possible designs for social organisation. Despite these warnings, the aim of this paper is to argue that *a society of mind will require an economy of mind*, and that economic theories, concepts and methods will have new applications in multi-agent systems (MAS) and the understanding of cognition. The paper, therefore, emphasises similarities not differences, and is primarily speculative, bearing on the epistemological foundations of MAS. The key idea is that the management of computational *resources* and computational *processes* needs to be combined. Such combination can occur through MAS that utilise a quantitative universal representation of value that mirrors the flow of agent products, much as money mirrors the flow of commodities in human social systems. Minsky anticipated such an idea, but did not develop it.

1.1 The coordination problem in multi-agent systems

A MAS can be thought of as a system that is composed of a collection of agents that normally have their own beliefs and goals, sharing a domain that allows actions to be performed, including communicative actions, such that the system meets some global requirements. The global requirements normally specify goals that can be met by agents acting cooperatively or competitively to discover solutions. Jennings (Jennings, 1996) discusses the *coordination problem* in MAS, which is the problem of ensuring that a society of agents interact in such a manner to achieve global goals given available resources. Coordination is required because 'there are dependencies between agent actions', 'there is a need to meet global constraints' and 'no one individual has sufficient competence, resources or information to solve the entire problem'. Without coordination the MAS would fail to produce useful global results. In this context, Jennings introduces the 'centrality of commitments and conventions hypothesis' which states that: *all coordination mechanisms can ultimately be reduced to commitments and their associated conventions.* It is important though to recognize that all commitments need not be generated in a conscious and deliberative manner: attachment structures in most bird and mammal species, for instance, certainly involve some kind of built-in commitments already 'installed' between certain individuals (selective mating, caring and protection of the young, territorial defense and so forth), without necessarily relying upon a conscious and explicit contractual basis.

A commitment is essentially a goal: an agent can make a commitment to itself (e.g., 'I will tidy my desk today') or to others, in which case it can be thought of as a pledge or promise (e.g., 'I will meet you at ten tomorrow'). As goals, commitments could result from many goal generactivators (Beaudoin, 1994), some very primitive, and some more deliberative. Joint commitments are possible (e.g., 'We will both move house') and are preconditions for cooperative action. Conventions, however, are rules that determine how an agent's commitments are to be formed, reconsidered, or rejected, and social conventions are rules that determine how agents should behave towards each other, for example if they change mutual commitments. For example, agent A may commit to meet agent B at ten because it is conventional for A to obey B because B has greater authority. Subsequently, however, A acquires a more pressing commitment and does not have sufficient time resources to honour the commitment to B. Hence, A informs B of the difficulty because it is a social convention to do so, allowing B to replan and ask another agent C, who can do the work of A, to meet at ten. This is an example of cooperation, communication of failure and replanning. Designing conventions and social conventions is difficult. It is likely that in natural systems, powerful mechanisms have evolved to generate, protect, manage and regulate conventions (Aube & Senteni, 1996a; Aube & Senteni, 1996b). Designing conventions amounts to designing a set of rules that can interact to produce coherent and useful emergent behaviour. Even simple rules, such as those that define Langton's ant, can produce emergent behaviours that are very difficult to deduce from the rules themselves (Cohen & Stewart, 1994). All coordination mechanisms may be reducible to, or expressible as, collections of commitments and conventions, but discovering useful coordination mechanisms is no easy task. Gasser (Gasser, 1991) has also written that commitments should not be seen as mainly initiated by individuals, but might rather be understood as emerging from the web of social interactions, which operate upon individuals as 'field forces' that constrain individual courses of action into joint behavior of cooperation or conflict.

1.2 Adaptive multi-agent systems

Adaptive MAS (AMAS) are a subclass of MAS that can continually reconfigure their activity to produce solutions that meet changing global constraints. For example, (Schaerf, Shoham & Tennenholtz, 1995) describe an AMAS that adaptively allocates jobs to different processing units under variable loads. The class of AMAS is sufficiently general to include many diverse kinds of system and mechanism, in much the same way as the class of adaptive agent architectures can include such mechanisms as reinforcement learning algorithms, artificial neural networks, genetic algorithms, and so forth. In the abstract, there are three distinct ways in which an AMAS can modify its global behaviour. It can (i) alter the behaviours of individual agents or (ii) alter the control relations between agents, for example dynamically defining groups of leader and follower agents. The former is a change of commitments, the latter a change in conventions and social conventions. Yet electing a leader or respecting some authority do involve commitments from the individuals concerned, and altering the kinds of control relations between agents may, in turn, alter the kinds of cooperative groups that can form. Alternatively, (iii) existing agents may be removed or qualitatively new agents may be introduced into the system. AMAS require coordination mechanisms that can cope with this kind of complexity. Such mechanisms need to allocate and reallocate agents to different tasks, alter social hierarchies, change individual agent behaviours to fit new circumstances, and provide means by which global constraints can direct local processing without the need for high bandwidth communication. In addition, there need to be natural ways in which global constraints can be defined within the system.

2 Money and exchange-value

Human designers of robots often turn to the natural world for design ideas. Similarly, human designers of coordination mechanisms for MAS can also turn to the natural world. The study of ant colonies, primate groups and human social interaction are all potential sources of inspiration. For example, (Aube & Senteni, 1996a; Aube & Senteni, 1996b) propose that the emotions arose to coordinate animal groups and therefore can serve as a foundation for coordination in MAS. They view commitments as a special kind of resource that insures access to basic commodities of survival value, and emotional structures as the control mechanism that manages these special resources. This section develops the contention that human economic activity provides an example of another important coordination mechanism – *currency flow* – that may be common to a certain class of *adaptive* MAS. We even think that such a view might help uncover the inner mechanics of motivations: that is, why and how some mental processes within the society of mind come to take precedence (be 'preferred') over others.

2.1 Basic requirements for the development of money

All human societies are in commerce with nature, extracting raw materials from the environment and returning human waste to the earth. Social organisation implies a division of labour amongst the individuals of the society, that is individuals perform different, socially useful functions. The total labour of society is shared amongst the different functions, and the products of this labour distributed according to some, usually implicit, scheme and through some collection of mechanisms. One very obvious requirement for a successful social system is that it reproduce its conditions of existence; that is, it must create conditions such that individuals survive and produce offspring so that the available labour of society is continually replenished. This requirement entails that what is produced, distributed and consumed should be so organised to satisfy those needs. This is one of the important coordination problems that social organisations are required to solve: *labour must be divided* and its *products distributed* so that at least a sufficient number of individuals' *basic needs are met*. It is the satisfaction of at least the basic needs of consumers that defines one of the major global constraints for successful human social systems. Money – the representation of the, as yet to be defined, value of a thing – arose at a certain point in human history to solve certain problems of production, consumption and exchange. Pure gold was first synthesised and coined as money in 625 BC in Greece (Boardman, Griffin & Murray, 1993). In a matter of fifty years trade had burgeoned, and there were banks, merchants, and money-lenders. A numerical representation of value had a revolutionising effect on the capabilities of human society. Subsequently, currency flow has been a common feature of human social organisation, surviving and developing through classical society, feudal arrangements, and industrial and modern finance capitalism. To understand the nature of money it is necessary to examine how and why it arose. The following account of the development of money is based on the opening analysis in Marx's *Capital* (Marx, 1970). It is a rough historical sketch of the *emergence of a social convention* in a human society. The account abstracts from the real historical development of money and uses simple stages and examples for the purpose of exposition. In addition, the emergence of money is examined in an idealised *simple commodity economy*, allowing complications such as price-fixing, cartels, monopolies, taxation, trade tariffs, travel costs, power relations, trade unions, and the legislative power of the state, to be ignored.

Stage one – simple exchange or swapping. Individual and relatively self-sufficient producers with a small surplus product, for example the peasant farmer whose chickens have lain too many eggs, exchange their goods for other goods. For example, 24 eggs may be exchanged for 2 loaves of bread. In this isolated act of exchange the equality relation (24 eggs = 2 loaves) is determined by the producers' respective opinions of the use-value of the other's goods. The term 'use-value' simply means that the good satisfies some desire or need. In other words, the respective values of the goods are determined locally and subjectively. The exchange of products has a precondition: each producer must have a surplus-product that the other desires. All exchange is performed with a view to obtaining another's surplus-product for the purposes of consumption. Money does not as yet exist.

Stage two – extended exchange or organised swapping. The development of better production techniques and increase in population size creates a greater surplus-product available for exchange. Instead of isolated acts of exchange there may be a definite geographical locale where trading takes place, that is the market. The peasant's 24 eggs now enter into potential relations with all the other commodities available. For example, the 24 eggs may now be exchanged for 2 loaves, or a pair of socks, or five candles, or a pound of butter and so forth. Importantly, an element of *competition* appears that was not present in stage one. Instead of a single peasant and consumer there is a social community of interconnected producers and consumers, for example peasants, bakers, and candlestick makers. Given the choice a baker will tend to exchange his bread for as many eggs as he can get from the community of peasants; conversely, a peasant will tend to exchange his eggs for as many loaves as he can get from the community of bakers. This systemic dynamic – colloquially, the notion of 'shopping around' – will, all other things being equal, have a tendency (which may not be fully realised) to force the equivalence relation between eggs and bread towards a particular ratio that holds for *all* such transactions. This equivalence relation will thus be *determined by the joint action* of the peasants and bakers in mutual competition. In other words, the respective values of the commodities are determined globally and socially (as opposed to locally and subjectively in stage one): local utility functions give way to global utility functions. An individual's local calculation increasingly becomes ineffective in the determination of the equivalence relation, which now tends to be fixed by the community as a whole.

Stage three – ubiquitous exchange. A community in which a good deal of exchange occurs soon finds it convenient to select a particular commodity to serve as the general form of value. A widely valued article would be the commodity to choose. For example, in a cattle community, cows would be the natural choice as medium of exchange (hence comes the derivation of the word 'pecuniary'). This special commodity then serves as a unit of comparison of value and is *directly* exchangeable with all other commodities. This overcomes the limitations of organised swapping, as all producers will now be willing to swap their goods for the general form of value. There need be no local coincidence of wants.

Stage four – money. As soon as a particular commodity is socially agreed upon to serve as the general form of value it becomes the money-commodity, that is it serves as a universal means of exchange. In most societies this commodity has been gold or silver, and not cows. For example, if 24 eggs = 1 measure of gold, and 1 measure of gold is coined as 10 pence, then 24 eggs have the price 10p. Gold serves as *the* embodiment of value, and can be exchanged for any other commodity. 'Although gold and silver are not by Nature money, money is by Nature gold and silver' (Marx, 1970). Precious metals were chosen because they exhibit uniform qualities but can be repeatedly divided and reunited at will to represent fine-grained differences in the numerical values of things. Also, they have a high value to weight ratio, which is useful if wealth is to be transported in pockets. There has been little computational, as opposed to historical, work on the development of universal means of exchange in MAS: Marimon et al. (Marimon, McGrattan & Sargent, 1990) describes investigations of the conditions in which money emerges in an artificial economy of adaptive, classifier system (Holland, 1986; Holland, Holyoak, Nisbett & Thagard, 1986) agents, although the chosen domain ontology bears only a superficial resemblance to real economies.

3 The function and properties of money

Money, therefore, is just like any other commodity except for a social convention that ensures it is the means of exchange in all transactions. The particular form of value, be it gold, silver, bronze, paper or virtual currency flows, is a secondary matter: it is function that counts. We will now examine the function and properties of money in greater detail.

Universal use-value. Stage two serves as the starting point. This stage presupposes a social division of labour that allows individual producers to specialise in the production of particular commodities. Specialisation leads to greater productivity of labour and a greater surplus-product. Such producers are therefore no longer self-sufficient and have a greater need to exchange. Exchange may only occur when there is a local coincidence of wants, which is a limiting factor. When this is the case, chains of exchange, or mediated exchanges involving "middle-men", will tend to occur. In fact, coincidence of wants can be only overcome through the use of "middle-men", and the probability of chains of coincidences of wants occurring decreases with the length of the chain; hence, the flow of commodities will be constrained. Exchange will tend to occur at definite locations of high connectivity, that is market places, which provide a higher probability of coincidences of wants. Also, due to the perishable nature of some commodities the chains of exchange may have to occur within definite time periods – for example, an agent will have difficulty accumulating enough apples in order to exchange them for a new boat. Although this stage of a social development is largely hypothetical it does highlight the limitations of bartering, and demonstrates that the eradication of the requirement for local coincidence of wants and commodities is one function of money, that is it serves as a *universal use-value*: it is a commodity that all find useful. Producers become willing to exchange for a representation of value which has the functional property of being able to buy the products of others' labour. One effect of the introduction of money, therefore, is to free up the flow of commodities.

Multiple instrumentality. In a developed money economy everything has a price. Money may be exchanged for any product of any labour.

Semantic determinacy. The exchange-value of commodities as represented by the money-commodity is expressed *quantitatively* and is compared to other quantities of value. Consequently, the meaning of money is globally determined in a society of numerate agents.

Constraints on possible exchanges. A loaf of bread may cost 50p but will not normally be exchanged for 49.5p because of the prevailing social convention. An agent with money can enter into many possible exchanges, whereas an agent without money has those possibilities severely constrained. The globally determined value of commodities defines what is and what is not a legal exchange, and serves as a kind of economic 'all-or-nothing' law.

Low communication costs. Consider the following thought experiment: instead of money there are a host of 'middle-men' exchanging lengthy notes listing individuals with their surplus-products and needs in an attempt to coordinate great chains of exchange mediated by coincidences of wants – a kind of global 'swap shop'. Such notes will entail high communication costs, due to the high information content of the notes, and high administration costs, such as matching up lists with lists. In direct contrast, money, being a number, is easily represented and removes the need for middle-men and their costly communications. (However, in some real markets, such as the housing market, chains of exchange occur frequently.)

Low storage costs. The quality of money does not change. It can be stored by adding up all the quantities into a bigger quantity – a larger denomination of note, for example. There need be no storage of many qualitatively different things, such as filing cabinets of 'coordination notes' in the above example.

Simple operators. Money requires only the very simplest operators: addition, subtraction and numerical comparison. No sophisticated local machinery is required to mediate the transaction. Money is quickly and easily parsed.

Accumulatability. Money, if it is metal, such as gold, does not perish. It can be stored indefinitely.

Distal connectivity. The coincidence of geographical location, time and wants for exchange to occur in a barter economy is overcome with the introduction of the money-commodity. Money can mediate wants, be easily transported from place to place, and be stored for future use, unlike perishables.

Domain independent representation. In exchange, value is compared with value. The value of a commodity does not represent anything external to the economy, nor does it represent any thing within the economy: it is internally relational, specifying an ordering over the set of commodities. Consequently, it would not make any difference to the functional role of money if the specific kind of labours within society changed, or if the external environment changed.

Coordination mechanism. Importantly, money introduces *supply and demand* dynamics that implement a distributed solution to a global coordination problem. The coordination problem is how private labour can be coordinated on a social scale so that individuals' basic and higher needs are met. Without a coordinating mechanism the social system would break apart; for example, basic goods might not be produced in sufficient quantities, or non-use-values (commodities that are not in demand) might be produced indefinitely.

Consider the following simplified scenario. An increase in productivity in one branch of production, say egg production, means that more eggs can now be produced by the same share of the total labour of society. As the rate of appearance of eggs on the market is higher then before, their value will decrease relative to other commodities, all other things being equal. Hence, egg producers will receive less money for their goods. This may force some egg producers 'out of business' as the money they get for their eggs cannot sustain their needs. Labour will then be freed to be employed elsewhere in other branches of the economy: some egg producers may turn to producing milk or bread, or any commodities that have sufficient social value to sustain their needs. In this way, the value of commodities and the market mechanism regulate the economy as a whole. The total labour of society is dynamically allocated and reallocated in definite proportions to reflect changes in production techniques and demand for products. 'It is only through the "value" of commodities that the working activity of separate, independent producers leads to the productive unity which is called a social economy, to the interconnections and mutual conditioning of the labour of individual members of society. Value is the transmission belt which transfers the working processes from one part of society to another, making that society a functioning whole' (Rubin, 1988). Currency flow *reinforces* social cooperation: for example, a particular agent will not be able to acquire a commodity without first expending labour that has sufficient value to other agents. The market mechanism of exchange-value, the social convention of money, and the local reasoning of autonomous economic agents serves to meet the basic requirements of economic organisation outlined at the beginning of section three.

4 Currency flow in multi-agent systems

Our analysis of the role of money in a commodity economy has implications for the design of adaptive multi-agent systems. In this section the particular form of value in economic systems is examined and compared to existing reinforcement learning algorithms, followed by a sketch of how currency flow could solve the problems of coordination in AMAS. Finally, a design hypothesis for AMAS coordination is proposed.

4.1 A universal, quantitative representation of value

All adaptive systems conform to the abstract schema of a *selective system* (Cziko, 1995), and all selective systems support concepts of *value* (Pepper, 1958; Wright, 1996a). A selective system has three components: (i) a *trial generator*, which is any mechanism that generates a variety of functions to produce outputs for particular inputs, (ii) an *evaluator*, which is a mechanism that evaluates the results of using particular functions to generate trials, where evaluation occurs through comparison to a norm, and (iii) a *process of selection*, which retains those functions associated with 'good' evaluations for future use, while discarding others. Selective systems implement the well-known generate, test, and select cycle. Specific examples of selective system improve their behaviour over time (cf. Darwinian evolution, genetic algorithms, classifier systems, neural networks, and so forth). In the abstract, economic systems are selective systems: the trials are the various concrete labours that produce commodities, the evaluatory mechanisms are the various needs and demands of individual consumers, and selection occurs through the buying and selling of commodities. Over time what is produced matches what is required given available resources. Economic systems have employed, from the beginnings of commodity production to the present day, a *quantitative representation of exchange-value* that is associated with every commodity and universally represented by a money-commodity. *Money mirrors the flow of commodities*, reinforcing those productive activities that meet the demands of consumers. Human economic systems are an existence proof that complex processing systems can be regulated by exchanging numerical quantities. Information-theoretic analogues of some of the properties of currency flow that were identified in section four may be useful for coordinating adaptive, largely parallel information processing systems composed of autonomous agents (e.g., multiple instrumentality, semantic determinacy, locally constraining, low communication and storage costs, simple operators, domain-independence and the ability to form a basis for coordination). In fact, recent work in artificial intelligence uses economic ideas for resource allocation problems (Wellman, 1995), including allocation of processing time, and reasoning about plans (Doyle, 1994).

4.2 Generalised reinforcement learning

Reinforcement learning (RL) algorithms are selective systems as defined above (see (Kaebling, Littman & Moore, 1995) for a review). RL is a type of trial and error learning, and holds out the promise of programming control programs for agents by reward and punishment without the need to specify how a task is to be achieved. The main design problem to be solved in reinforcement learning is the *credit assignment* problem, which is the problem of 'properly assigning credit or blame for overall outcomes to each of the learning system's internal decisions that contributed to those outcomes' (R. S. Sutton, quoted in (Cichosz, 1994)). More precisely, RL involves learning functions defined on the state and action space of a task, driven by a real-valued reinforcement signal. The details of how this is achieved depends on the particular function representation used. Examples of RL algorithms are Q-learning (Watkins & Dayan, 1992), classifier systems (Holland, 1975; Holland, Holyoak, Nisbett & Thagard, 1986; Wilson, 1995), and W-learning (Humphreys, 1996). Marvin Minsky's *Snarc* machine was an early reinforcement learner that encountered the credit-assignment problem (see section 7.6 of (Minsky, 1987)).

RL algorithms use a quantitative representation of value, the reinforcement signal, to select those behaviour-producing components that satisfy conditions of reward over and above those components that do not. Behaviour-producing components with high reward will be more likely to dispositionally determine the behaviour of the system in the future than those components with low reward. For example, the bucket-brigade algorithm used in early classifier systems was inspired by an economic metaphor, in which system rules are agents consuming and producing internal messages (commodities) who each possess a certain amount of strength (money) which they exchange for messages at a global blackboard (the market). Most RL algorithms are composed of rules. (Shoham & Tennenholtz, 1994) discuss a generalisation of RL to MAS called co-learning. Co-learning involves individual agents learning in an social environment that includes other agents. Co-learning agents must adapt to each other. (Kittock, 1995) describes some computational experiments on the emergence of social conventions through co-learning. Work of this kind is beginning to explore how MAS can adapt by reinforcement signals. The use of a universally recognised, quantitative representation of value is common to both RL algorithms, co-learning, and economic adaptation via currency flow. However, the latter may require MAS with substantially more sophisticated agents than those used currently. The theoretical relations at the information processing level of abstraction between reinforcement and payment for goods is an issue that can be fruitfully investigated by MAS research.

4.3 The ability to buy processing power

In economic systems and reinforcement learners, *possession of 'money' by an 'agent' is a dispositional ability to buy processing power* (Wright, 1996b). For example, a producer who makes a profit will have more money to employ more people (to buy processing power directly) and more raw materials (to buy the results of prior processing). Whether a thing is purchased or a person is purchased for a certain period of the day, an amount of labour power has been assigned to the purchaser. That the labour power has already been expended and is in the form of a commodity, or will be expended and is in the form of a commodity-maker, is a secondary matter. In both cases, processing resources have been bought. Individual profits and losses regulate this ability to commandeer and allocate social resources. Similarly, a rule in a classifier system uses its accumulated strength to bid against other rules for messages in the 'marketplace'. Rules with high strength are more likely to outbid rules of low strength, process the message, and dispositionally determine the behaviour of the system. The bucket-brigade adaptively alters the ability of rules to buy processing power. The same holds for the weights of policy functions in Q-learning. One of the most important scarce resources in MAS are the agents themselves. The total processing power of the MAS is limited, where processing power is ability to do work. Similarly, Marx, drawing on the classical tradition in economics, emphasised labour-power as a finite resource in economic systems, developing the labour theory of value based on this conception. Labour-power is also the ability to do work. Whether it is computational agents performing abstract operations, or real people performing concrete operations, a transformation is taking place that can be called work.

Adaptive MAS must search for solutions to, perhaps continuously changing, global constraints. Therefore, there needs to be an ordering over the various agents of the adaptive system: some agents will perform more useful work than others with respect to certain constraints. The computational resources of the system should be concentrated on useful agents, be it in terms of giving them greater social power or allowing them access to more social products. In other words, useful work within a society (or useful processing within a mind) should be reinforced. The design principle of a quantitative representation of value that functions as an ability to buy processing power can integrate processing (useful computational work) and resources (limited computational power) with relatively low communication costs. Agents with more money can employ other agents, buy the products of other agents' work, and have greater control over system behaviour. Given these abstract and general considerations it is possible to sketch how currency flow could serve as a basis for coordination in adaptive MAS.

4.4 Specifying global constraints

Economic systems suggest a natural way to specify the global constraints of an AMAS. In simple commodity economies it is the wants of consumers that determines what is and what is not a use-value. In just so happens that in real economies consumers are normally also producers, but in artificial AMAS the functions can be separated and assigned to different agents. A set of consumer agents that function as the *sole sources of payment* can define the goals of the system. Producer agents must satisfy consumers' wants if they are to receive value for their work. It is feedback from consumers to antecedent producers in the form of payment that selects those productive behaviours that satisfy the global goals of the system, much as conditions for reward select adaptive policies in RL algorithms. For example, an AMAS may be designed to find plans for successful operation in a microworld domain, such as blocks-world. A set of consumer agents can be defined whose various needs are information items declaring that the system has achieved certain objectives, such as stacking a tower of blocks or building certain shapes and so forth. These information items are analogous to desired commodities in economic systems: they are the use-values of the system. A set of producer agents may then attempt to produce the required information items by performing work in the domain, that is produce information items interpretable as actions by a scheduler. Only those agents or group of agents that produce the correct set of actions and corresponding results receive money from the set of consumers. Partial solutions may receive partial payment allowing hill-climbing and iterative trial and error search. Baum (Baum, 1996) describes the 'Hayek machine' that learns to solve blocks world planning problems using a free market of interacting agents and a simplified price mechanism. Weiss (Weiss, 1995) describes the 'Dissolution and Formation of Groups' algorithm that solves block world problems using a collection of agents that learn through reinforcement and form into cooperative groups with 'leaders'. The Contract Net Protocol (d'Inverno & Luck, 1996; Smith, 1980; Smith

& Davis, 1981) has, for many years in the field of DAI, also embodied some of these economics-flavored ideas. In a contract net, a manager agent broadcasts a task announcement message, and receives bids from contractor agents. The manager evaluates the bids, selects among them, and allocates the task, or part of it to the most interesting bidder. One innovative aspect of our proposal though, is to place emphasis on the currency flow itself as embodying global constraints, as opposed to examining the mechanics of a local announcement-bidding-allocation process.

4.5 Dynamic control relations

As stated, an AMAS may need to alter the control relations between agents in order to meet global goals. A relation of control exists between agent A and agent B if A can determine, or dispositionally determine, B's processing. For example, A may be able to command B to perform a particular task, or A may be only able to request that B perform a task in particular circumstances, and so forth. In human societies there is a wide variety of relations of control, some more benign than others. *Autonomous* agents will often have objectives that conflict with other autonomous agents. One way for agents to overcome conflicts of interest is through negotiation, a process by which a group of agents communicate with one another to arrive at a mutually acceptable course of action. For example, when a conflict is encountered the agents involved may generate proposals for joint commitments with associated explanations. The mooted proposals may then be evaluated, and various counter-proposals or compromises suggested. The Socratic dialogue continues until agreement is reached (Parsons & Jennings, 1996).

In order that local negotiations can meet global requirements there is need for local information, referring to those requirements, that can form a basis for controlling the negotiations. Without such information agents could negotiate commitments that led to globally incoherent behaviour or that required too many resources (i.e., the construction of unrealisable social plans). In human societies many negotiations occur within the context of financial costs. For example, much institutional behaviour consists of negotiating compromises constrained by available funding. The local possession of value limits the formation of commitments, which are essentially about resources (Bond, 1990; Gerson, 1976). By giving access to additional resources, commitments thus become valuable resource in themselves (Aube & Senteni, 1996a; Aube & Senteni, 1996b). However, local possession of value can allow in turn the formation of new commitments. For example, a new injection of funding can release prior constraints on planning: planners may now have sufficient power to employ other agents to do their bidding or buy the resources needed to complete their plans. Money, as the ability to buy processing power, is an ability to form control relations; and *the flow of money adaptively allocates and reallocates constraints on local commitment formation*[1]. Again, one reason for this rests on the fact that commitments themselves constitute a special kind of resource, and that money embodies the value that is computed for these resources through social transactions. It is the requirements for global problem solving that necessitate the imposition of limits on local problem solvers: *Hobbes chairs the Socratic dialogue*. 'Participation in any situation, therefore, is simultaneously *constraining*, in that people must make contributions to it, and be bound by its limitations, and yet enriching, in that participation provides resources and opportunities otherwise unavailable' (Gerson, 1976). Social agents commit to a social convention of money that simultaneously constrains and enriches possible local outcomes.

[1] Compare (Bond, 1990; Gerson, 1976) where money is viewed as just another kind of resource.

4.6 Dynamic reallocation of labour

As stated, an AMAS may need to reallocate agents to different tasks in order to meet global goals and maintain coherent behaviour. One possible solution is a global controller that has a wider picture of the whole system and directs the activities of others; however, keeping the agent informed could entail high communication costs, create a communication bottleneck, and render the other agents unusable if the controller failed (Jennings, 1996). The alternative is to distribute data and control, and economic systems suggest at least two possible mechanisms. One one side, a system composed of adaptive agents that attempt to maximise personal utility will exhibit distributed reorganisation of labour. Adaptive utility maximisers will search for rewarding tasks, allocating and reallocating themselves to different parts of the developing solution. For example, if a system constraint changes, such as a consumer agent requesting a qualitatively different result, then the agents that previously serviced the consumer will search for new forms of cooperation in order to produce the new result and regain gainful employment (c.f. rule discovery of rewarding areas of the pay-off landscape in classifier systems). On the other side, a system that allows agents to sell their processing power to employer agents will exhibit organisational control, that is 'centralised' reorganisation of labour. For example, sufficiently wealthy employers may direct and redirect the processing of large groups of agents, perhaps at the expense of relatively high communication costs within the organisation. In both cases, however, it is money that forms the basis of the allocation of labour, either as a universal want or an ability to buy processing power. Note also that areas of the search space may be redundantly assigned to multiple agents, much as competition occurs within branches of production in real economies.

4.7 The currency flow hypothesis

Given these considerations the following design hypothesis is proposed:

> **The currency flow hypothesis for adaptive multi-agent systems:** Currency flow, or circulation of value, is a common feature of adaptive multi-agent systems. Value serves as a basis for coordination; for example, it integrates computational resources and processing by constraining the formation of local commitments. Circulation of value involves (i) altering the dispositional ability of agents to gain access to limited processing resources, via (ii) exchanges of a quantitative domain-independent representation of value that mirrors the flow of agent products. The possession of value by an agent is an ability to buy processing power.

The design hypothesis is a hypothesis because it is a statement about designs that can be falsified. It states something about the *functional* organisation of AMAS at the level of information processing. If the MAS research community discovers designs that meet the requirements for AMAS but do not use a currency flow mechanism then the hypothesis is falsified: the design feature is not common to that set of requirements. It is more likely, however, that the hypothesis in its current form is too general and imprecise. Future research may show that currency flow cannot meet all possible requirements for adaptive MAS behaviour, or that currency flow is necessary but not sufficient, or it is simply one of a range of possible alternatives, or it works for only certain types of constituent agents, and so forth. Therefore, the hypothesis serves as a guide, pointing towards perhaps fruitful areas of AMAS design-space based on an analysis of an existing AMAS.

For a MAS to use currency flow mechanisms the constituent agents will require a minimal set of capabilities. A first pass requirements analysis suggests that minimally economic agents will need to be able to form mutual plans with other agents, possess planning capabilities to construct and choose between alternative possible options, handle money, reason about costs, negotiate, and take and give requests and commands. Without these capabilities the economic system may fail to use currency properly or fail to find solutions to global requirements and so forth. Minimally economic agents also require motivations to cooperate. A number of theories of the evolution of cooperative, altruistic behaviour in the natural world have been proposed, notably how selfish behaviour leads to widespread TIT-FOR-TAT strategies (Axelrod & Hamilton, 1981), or how genuinely altruistic behaviour can be inculcated in docile (those that are receptive to social influence) agents with bounded rationality (those unable to fully evaluate how acquired behaviours affect personal fitness) (Simon, 1990). The Baldwin Effect (Baldwin, 1896) predicts that adaptive, acquired traits tend to be genetically assimilated over time. Assimilated TIT-FOR-TAT may form the foundation for certain social emotions, such as gratitude and guilt, which facilitate cooperative behaviour in natural MAS. The existence of a universal 'norm of reciprocity' (Gouldner, 1960) in social behaviour has been known for some time in sociological theory. (Shoham & Tennenholtz, 1994) report experimental results that seem to show that MAS consisting of agents that use a 'highest cumulative reward' rule (i.e., agents that choose actions likely to yield the highest pay-off) are inefficient in developing social cooperation. Societies of pure personal utility maximisers may not cooperate effectively. As Aube and Senteni (Aube & Senteni, 1996a; Aube & Senteni, 1996b) have argued, MAS agents may need powerful 'emotional' control structures to support cooperative behaviour, perhaps even as a precondition of the social transactions from which the money-commodity 'device' and the commitments upon which it has to rely could emerge at all. Thus, whether minimally economic agents need such control structures is an open question. MAS design can explore what kinds of cooperative behaviours are required for what kinds of global behaviour.

5 Conclusion: implications for cognitive science

> ... another family of agents might contrive to use a quantity that doesn't actually exist at all, but whose amount is simply "computed". I suspect that what we call the pleasure of success may be, in effect, the currency of some such scheme.
>
> M. Minsky, *The Society of Mind*, p. 284 of (Minsky, 1987).

If the society of mind requires an economy of mind and the information processing level of the brain is organised in such a manner, then we would expect some evidence of 'currency flow' in our mental flora and fauna. Wright (Wright, 1996b; Wright, 1996a) presents a circulation of value theory of achievement pleasure and failure unpleasure that explains the valenced component of some emotional states, in particular those involved in attachment and loss scenarios. Very briefly, the monitoring of virtual 'currency flows' performing credit-assignment can account for some forms of mental pleasure and unpleasure. The theory is related to some old ideas of Freud, in particular his concept of 'psychical energy' or 'libido'; however, the circulation of value sheds the connotations of vitalism but retains and extends the functionality of 'libido'. This work builds on previous work with Aaron Sloman and Luc Beaudoin on cognitive modelling of the emotions (Wright, Sloman & Beaudoin, 1996; Sloman, Beaudoin & Wright, 1994). It is a recurring assertion that there is a relative neglect of motivation and emotion in cognitive science. For example, Simon's seminal paper (Simon, 1967) was an attempt to answer Neisser's criticisms that information processing theories of mind cannot account for feelings and emotions. More recently, (Newell, 1990) lists motivation and emotion as missing elements that need to be included in more comprehensive information processing theories of mind. (Shoham, 1996) argues that AI can and should benefit from economic ideas, for instance modelling the cost and value of information. If economic ideas are applicable to artificial intelligence then they should also be applicable to natural intelligence and therefore of relevance to cognitive science. The concepts of value, currency flow, and ability to buy processing power are a step towards this.

This paper has argued that a society of mind will require an economy of mind, in particular that adaptive multi-agent system design will benefit from using currency flow as a coordination mechanism. A hypothesis was proposed stating that currency flow mechanisms are likely to be a common feature of AMAS. The useful design properties of a money-commodity were analysed. One important feature is that currency flow adaptively allocates and reallocates the ability of agents to form local commitments. The social convention of money integrates both resources and processing by functioning as an ability to buy processing power.

Acknowledgments

Thanks to Aaron Sloman and Tim Kovacs for comments and criticism.

References

Aube, M. & Senteni, A. (1996a). Emotions as commitments operators: a foundation for control structure in multi-agent systems. In *Proceedings of the Seventh European Workshop on Modelling Autonomous Agents in a Multi-Agents World, MAAMAW '96, Lecture Notes in Artificial Intelligence*. Springer-Verlag.

Aube, M. & Senteni, A. (1996b). What are emotions for? commitments management and regulation within animals/animats encounters. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J., & Wilson, S. W. (Eds.), *From Animals to Animats IV, Proceedings of the Fourth International Conference on the Simulation of Adaptive Behavior*, pages 264–271. The MIT Press.

Axelrod, R. & Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211:1390–1396.

Baldwin, J. M. (1896). A new factor in evolution. *American Naturalist*, 30:441–451, 536–553. Reprinted in *Adaptive Individuals in Evolving Populations: Models and Algorithms*, edited by R. K. Belew and M. Mitchell (SFI Studies in the Sciences of Complexity, Proc. Vol. XXVI, Addison-Wesley, Reading, MA, 1996).

Baum, E. B. (1996). Toward a model of mind as a laissez-faire economy of idiots. In *Proceedings of the Thirteenth International Conference on Machine Learning.*

Beaudoin, L. P. (1994). *Goal processing in autonomous agents.* PhD thesis, School of Computer Science, The University of Birmingham.

Boardman, J., Griffin, J., & Murray, O. (1993). *The Oxford History of the Classical World.* Oxford: Oxford University Press.

Bond, A. H. (1990). A computational model for organization of cooperating intelligent agents. In *Proc. of the Conference on Office information Systems*, pages 21–30. Cambridge, MA.

Cichosz, P. (1994). Reinforcement learning algorithms based on the methods of temporal differences. Master's thesis, Institute of Computer Science, Warsaw University of Technology.

Cohen, J. & Stewart, I. (1994). *The collapse of chaos, discovering simplicity in a complex world*. Viking.

Cziko, G. (1995). *Without Miracles, universal selection and the second Darwinian revolution*. Cambridge, Massachusetts: The MIT Press.

Dawkins, R. (1990). *The selfish gene*. New York: Oxford Unviversity Press.

d'Inverno, M. & Luck, M. (1996). Formalizing the contract net protocol as a goal-directed system. In de Velde, W. V. & Perram, J. W. (Eds.), *Agents Breaking Away, Proceedings of the 7th European Workshop on MAAMAW*, Lecture Notes on Artificial Intelligence, No. 1308, pages 72–85, Berlin. Springer.

Doyle, J. (1994). A reasoning economy for planning and replanning. In *Technical papers of the ARPA Planning Initiative Workshop*.

Gasser, L. (1991). Social conceptions of knowledge and action: Dai foundations and open systems semantics. *Artificial Intelligence*, 47:107–138.

Gerson, E. M. (1976). On 'quality of life'. *American Sociological Review*, 41:793–806.

Gouldner, A. W. (1960). The norm of reciprocity: a preliminary statement. *American Sociological Review*, 25(2):161–178.

Holland, J. H. (1975). *Adaption in natural and artificial systems*. The MIT Press.

Holland, J. H. (1986). Escaping brittleness: the possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In Michalski, R. S., Carbonell, J. G., & Mitchell, T. M. (Eds.), *Machine learning, an artificial intelligence approach*. Los Altos, California: Morgan Kaufmann.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986). *Induction: processes of inference, learning and discovery*. The MIT Press.

Humphreys, M. (1996). Action selection methods using reinforcement learning. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J., & Wilson, S. W. (Eds.), *From Animals to Animats IV, Proceedings of the Fourth International Conference on the Simulation of Adaptive Behavior*. The MIT Press.

Jennings, N. (1996). Coordination techniques for distributed artificial intelligence. In O'Hare, G. & Jennings, N. (Eds.), *Foundations of distributed artificial intelligence*. John Wiley & Sons.

Kaebling, L. P., Littman, M. L., & Moore, A. W. (1995). Reinforcement learning: a survey. In *Practice and Future of Autonomous Agents, volume 1*.

Kittock, J. E. (1995). Emergent conventions and the structure of multiagent systems. In Nadel, L. & Stein, D. (Eds.), *1993 Lectures in Complex Systems: the proceedings of the 1993 Complex Systems Summer School*, Santa Fe Institute Studies in the Sciences of Complexity Lecture Volume VI. Santa Fe Institute, Addison-Wesley Publishing Co.

Marimon, R., McGrattan, E., & Sargent, T. J. (1990). Money as a medium of exchange in an economy with artificially intelligent agents. *Journal of Economic Dyanmics and Control*, (14):329–373.

Marx, K. (1970). *Capital, a critical analysis of capitalist production*, volume 1. Lawrence and Wishart. Originally published in 1887.

Minsky, M. L. (1987). *The Society of Mind.* London: William Heinemann Ltd.

Newell, A. (1990). *Unified Theories of Cognition.* Cambridge, MA: Harvard University Press.

Parsons, S. D. & Jennings, N. R. (1996). Negotiations through argumentation – a preliminary report. In *Proceedings of the Second International Conference on Multi-Agent Systems.*

Pepper, S. C. (1958). *The Sources of Value.* University of California Press.

Rubin, I. I. (1988). *Essays on Marx's Theory of Value.* Montreal: Black Rose Books. Originally published 1928.

Schaerf, A., Shoham, Y., & Tennenholtz, M. (1995). Adaptive load balancing: a study in multi-agent learning. *Journal of Artificial Intelligence Research*, 2:475–500.

Shoham, Y. (1996). The open scientific borders of ai, and the case for economics. Available at URL http://robotics.stanford.edu/users/shoham/ai-econ.html. Draft note written for ACM/CRA/NSF worksop on Strategic Directions for Computing Research, to be held at MIT in June 96.

Shoham, Y. & Tennenholtz, M. (1994). Co-learning and the evolution of social activity. Technical Report CS-TR-94-1511, Robotics Laboratory, Department of Computer Science, Stanford University.

Simon, H. A. (1967). Motivational and emotional controls of cognition. Reprinted in *Models of Thought,* Yale University Press, 29–38, 1979.

Simon, H. A. (1981). *The Sciences of the Artificial* (second ed.). The MIT Press.

Simon, H. A. (1990). A mechanism for social selection and successful altruism. *Science*, 250:1665–1668.

Sloman, A., Beaudoin, L. P., & Wright, I. P. (1994). Computational modeling of motive-management processes. In Frijda, N. (Ed.), *Proceedings of the Conference of the International Society for Research in Emotions*, Cambridge. ISRE Publications.

Smith, R. G. (1980). The contract net protocol: high-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, 29(12):1104–1113.

Smith, R. G. & Davis, R. (1981). Frameworks for cooperation in distributed problem solving. *IEEE Transactions on Systems, Man and Cybernetics*, 11(1):61–70.

Watkins, C. & Dayan, P. (1992). Technical note: Q-learning. In *Machine Learning 8*, pages 279–292.

Weiss, G. (1995). Distributed reinforcement learning. *Robotics and Autonomous Systems*, 15:135–142.

Wellman, M. (1995). Market-oriented programming: some early lessons. In Clearwater, S. (Ed.), *Market-Based Control: A Paradigm for Distributed Resource Allocation.* World Scientific.

Wilson, S. W. (1995). Classifier fitness based on accuracy. *Evolutionary Computation*, 3(2):149–185.

Wright, I. P. (1996a). Design requirements for a computational libidinal economy. Technical Report CSRP-96-11, School of Computer Science and Cognitive Science Research Centre, University of Birmingham. Submitted to *Cognition and Emotion*.

Wright, I. P. (1996b). Reinforcement learning and animat emotions. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J., & Wilson, S. W. (Eds.), *From Animals to Animats IV, Proceedings of the Fourth International Conference on the Simulation of Adaptive Behavior*, pages 272–281. The MIT Press.

Wright, I. P., Sloman, A., & Beaudoin, L. P. (1996). Towards a design based analysis of emotional episodes. *Philosophy Psychiatry and Psychology*, 3(2):101–137.