

UTILISING HAYEK'S CATALAXY IN A GLOBAL COMPUTER GRID

Colin Richardson

Internet Economist
William Penney Laboratory
Imperial College London

Abstract

Friedrich von Hayek's 'catallaxy' predicts the costless emergence of a 'spontaneous order' of stable prices in competitive markets, whose participants are free to make the price *bids* and *asks* that each thinks will maximise their own utility or profit. Catallaxy is the only serious alternative to the orthodox general equilibrium paradigm.

Since the mid-1990s, co-operating university, government and corporate research bodies have been pooling their high performance computing resources into computational grids, using central file servers to facilitate resource discovery and control resource allocation and job scheduling. Such inefficient centralised approaches represent a single point of failure and a bottleneck choking the flow of control information, thus ruling out Walrasian auctioneers.

Computer scientists now envisage merging these grids with billions of Internet-connected computers to birth a powerful World Virtual Machine (WVM) accessible by anyone, anywhere, anytime. This paper discusses research aimed at building a WVM using software architectures and fully-decentralised peer-to-peer networks based on the catallaxy paradigm of heterodox (Austrian) economics.

Introduction

The first section of this paper describes the catallaxy paradigm of Austrian economics, which predicts the costless emergence of a 'spontaneous order' of stable prices in a system of competitive markets. In such economies, providers and users of goods and services are free to make the price *bids* and *asks* that each thinks will maximise their own utility or profit. Friedrich von Hayek's notion of 'order for free' emerging from market transactions among these scattered holders of partial knowledge is the only serious alternative to Leon Walras' paradigm of general equilibrium. GE theory features an omniscient auctioneer who prohibits 'false trading', i.e. deals made *before* the economy's vector of equilibrium relative prices is disclosed via open outcry of bids and asks. Yet non-equilibrium false trading is the very process that drives the catallaxy to emerge in the real, not the theoretical, world.

In the second section it is noted that, for over a decade, 'islands' of co-operating university, government and corporate researchers have been pooling their expensive high performance computing and high data throughput resources into computational grids, using centralised file servers to facilitate resource discovery and control resource allocation and job scheduling. Also, since 1999, 'archipelagos' of personal computer owners have been participating in distributed grid computing projects like SETI@home. Most of the post-2000 'peer-to-peer' (P2P) filesharing networks also suffer from this need for central servers, which represent a single point of failure and present bottlenecks to the flow of control information.

The third section discloses that the number of digital devices connected to the Internet reached one trillion in 2006. Computer scientists now envisage merging the island and archipelago grids and networks with the 'Internet ocean' (the planet's billions of standalone

computers) to birth a powerful World Virtual Machine or Global Open Grid, accessible by anyone, anywhere, anytime. At Imperial College London and elsewhere¹, the boffins are turning to heterodox economics for elegant solutions to the problem of efficient resource discovery, resource allocation and job scheduling. This section describes leading-edge research presently underway to create software architectures and design fully-decentralised – hence infinitely scalable – P2P networks, whose novel design features are drawn directly from the catallaxy paradigm and the famous ‘small worlds principle’ of ‘six degrees of separation’. The Walrasian auctioneer of neoclassical economics is ill-equipped, in fact unable, to solve the WVM problem.

In the fourth section we report some early results of simulating a Global Open Grid based on Hayek’s catallaxy paradigm predicting spontaneous emergence of the vector of efficient market-clearing prices for commodities called ‘grid fabric services’. Although serious price evolution experiments are still being performed, these results already indicate that modelling grid computing markets as decentralised peer-to-peer networks leads to stability. This research has been undertaken by computer scientists in the Department of Computing, Imperial College London and is continuing.

The conclusion sums up and describes one possible future for the Internet, Web and Grid. A planet-wide *digital economy* is envisaged, one in which not only computing, storage, bandwidth, software, database, instrument, and sensor services are commercially traded (and often given away or cooperatively shared) on a global P2P network, but also *any* commodity whose characteristics can be inscribed on an XML message² resembling a simple structured email.

Hayekian Microeconomics

Friedrich Hayek published three contributions to a famous 1930s debate between left- and right-wing economists: *Socialist calculation I: the nature of the problem* (1935a), *Socialist calculation II: the state of the debate* (1935b) and *Socialist calculation III: the competitive solution* (1940). He also wrote three essays on the problem of knowledge in competitive markets: *Economics and knowledge* (1937), *The use of knowledge in society* (1945) and *The meaning of competition* (1948). In this body of work, Hayek developed and refined the notion of catallaxy as an alternative paradigm to Walrasian general equilibrium for explaining how a large set of interacting markets actually works.

The problem for Austrian School economists in the 1930s was that their left-wing colleagues emphasised how GE, as well as assuming perfect knowledge, was actually a *centralised* system in that it needed a single mythical ‘auctioneer’ to discover and publicise the market-clearing vector of equilibrium relative prices. Oskar Lange (1936) and others asked: Could this theoretical entity be replaced by a real-world Central Planning Bureau having access to all relevant scientific, technological, price, and quantity data? In principle it could, which troubled Hayek deeply and led to his frontal attack on the Walrasian system.

¹ See Appendix B for information on the CATNETS project; participating computer scientists are accredited to the Universities of Bayreuth, Politecnica de Catalunya, Karlsruhe, Mannheim, delle merche Ancona, Cardiff, and ITC-irst Trento (now Fondazione Bruno Kessler).

² According to Wikipedia (2007a) the Extensible Markup Language (XML) is classified as an extensible language because it allows its users to define their own tags. Its primary purpose is to facilitate the sharing of structured data across different information systems, particularly via the Internet. It is used both to encode documents and serialize data.

Hayek proposed that *decentralised* market systems feature local knowledge being discovered and learned via evolutionary processes, and economies successfully self-organising to utilise vast amounts of such scattered fragments of information. What the emergent catallaxy spontaneously delivers is ‘order for free’: a competitive economy, in fact, *coordinates itself*, establishing price vectors that ensure supply and demand come into approximate balance in every market. Prices are fairly stable, not straying too far from equilibrium over time. The catallaxy paradigm is an inherently dynamic and nonlinear complex adaptive system (CAS), whereas the GE paradigm is essentially a static and linear set of simultaneous excess demand equations.

John Holland was a pioneer in the CAS field. Along with genetic algorithms he invented a unique machine-learning software architecture: the ‘classifier system’ (Holland 1986). Rules trigger actions, and these actions will have good or bad consequences. When judged good, the system ‘pays’ the triggering rule. This then passes along some of the payment to the ‘contractor’ rules that triggered it and they, in turn, pay their own ‘subcontractors’. The flow of payments retraces, in reverse order, the chain(s) of causation that led from stimuli to useful actions. Over time those rules (cf. business firms) not contributing to useful system behaviour go extinct (cf. go bankrupt). This Lamarckian process calls to mind the fate of inefficient businesses in a market economy. Networks of market relationships – who contracts with whom, or even who knows whom – constantly shift, adapting to complexity in the world. When a classifier network alters its connectivity relationships because of payment flows, the system indirectly ‘learns’ facts never previously known to any agent within the system, just as human trading networks promote such learning.

Effectively, Holland’s spontaneously ordering and evolutionary learning system is the same one Hayek had earlier proposed. Eymann et al. (2005) state that catallaxy is derived from the Greek ‘katallatein’, meaning ‘to barter’ and ‘to join a community’. The catallaxy is a state of coordinated *individual* actions, brought about by the bartering, communicating and relationship-building activities of economic agents, leading them to achieve an unplanned *community* goal. It is Hayek’s take on Adam Smith’s ‘invisible hand’: prices are not only rates of exchange between goods, but also an efficient mechanism for communicating information between localities and communities. In a complex, uncertain environment, economic agents are never able to predict the precise consequences of their actions, and it is this existential *ignorance* that actually makes the price system work. The spontaneous order can never be designed by a perfectly informed planner (or auctioneer) who simply ‘gets the prices right’, because the pattern of pricing actually evolves as a result of *lack* of knowledge. People cannot know each and every circumstance that determines their actions; since this fog of unknowing is permanent and incurable, it is called by Hayek ‘constitutional ignorance’.

Competitive markets guarantee more ‘informational efficiency’ than centralised plans because they are better able to deal with a slew of facts that “never exists in concentrated or integrated form but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess” (Hayek 1945, p. 77). Hayek emphasised that while GE theory deals mostly with scientific or technological facts, the main question was how societies deal with personal, dispersed and fragmented knowledge (Hayek 1945, pp. 79-81). This implies that much information cannot be dealt with as if it were given *ex ante*. In fact, most of the knowledge is generated and used *during* the process of adjustment to equilibrium. This cannot occur in the GE mechanism, because the Walrasian auctioneer forbids ‘false trading’ (i.e. deals struck at non-equilibrium prices) until open outcry bids and asks reveal the economy’s equilibrium vector of relative prices. Yet false trading is the very process that drives the catallaxy to emerge. The market process is necessarily one of incessant trial-and-error, so it is not surprising to observe that people often are wrong in their decisions.

Participants in each single market are motivated and directed by their income expectations. Starting with the ruling relative prices, people create plans, determine their market behaviour, then revise both according to experiences learned during bouts of false trading. The economy's relative prices keep changing, providing new information and shaping motivations, until the spontaneous order emerges. Income expectations and changes in relative prices are the stabilisers of the system as a whole, without which coordination of the actions of individual agents would fail. All markets are connected and this 'interdependency of prices' defeats constitutional ignorance by bringing into play formerly unknown pieces of knowledge, widely distributed among anonymous individuals. The price system thus *distills information* about concrete market circumstances, regardless of human intentions. (The far more stringent requirements for establishing a Walrasian general equilibrium vector of relative prices are listed in Appendix A.)

Grid Present: Islands and Archipelagos in the Internet Ocean

From its inception, the Internet developed mainly along 'server-client' lines. Less intelligent 'client' computers would send requests – using slow dial-up data communications – to more powerful machines for files known to be held in storage on these 'server' computers. Later, the introduction of HTML/HTTP³ encouraged emergence of the Web, which soon developed a centre-periphery split. The 'core' of corporate, government and academic servers transmitted pages of information (often including offers to sell goods and services) to the 'edge' of far-more-numerous clients, operated by potential paying consumers. But this all began to change after the year 2000, with the spread of fast broadband data communications and peer-to-peer Web overlay filesharing networks like Gnutella, Freenet and Napster.

According to CacheLogic (2006), "Over the last five years P2P has established itself as ... one of the key drivers for consumer broadband uptake. This popularity has positioned P2P as the dominant protocol on the Internet, representing between 60% and 80% of total traffic on Service Provider networks." Peer-to-peer computing is a classic 'killer innovation', empowering the mainly client-only edge and levelling up the eCommerce playing field *vis à vis* the Web's mainly server-only core. Unfortunately, P2P has had a bad press because much of this activity constitutes filesharing of *copyright-protected* intellectual property – mainly by swapping tiny fragments of video, audio, games, software, and text files, with the recipients reassembling them into the original entire work. Significantly, the way the edge computers are challenging the hegemony of the core machines is by acting as combination server-clients, or 'servents', communicating directly with each other, i.e. in peer-to-peer fashion. (Until the owners cleaned up their act, Napster was the most notorious of these barely legal P2P filesharing systems.)

The incentive to encourage uploading from one's own stock of file fragments is the prospect of downloading desired file fragments from stocks held by others, in order to possess the complete work. P2P network designers have had to build in protection against free riders who download without providing any *quid pro quo*. These designers also must provide some means for participants to identify those servents on which their desired file fragments are held, so that request messages can reach them. This requires some kind of centralised registration component(s), meaning that even advanced P2P network designs (distributed hash table systems like Chord, Pastry and Tapestry) must utilise one – or more usually a series of –

³ HTML=Hypertext Markup Language and HTTP=Hypertext Transport Protocol.

‘central’ servers to supply the correct Internet addresses.⁴ Centralisation militates against the Holy Grail of *infinite scalability* of P2P networks. For this reason, P2P filesharing has been unable to colonise the immense ‘Internet ocean’ of billions of machines comprising the Web’s edge. Instead, hundreds or thousands of servers daily join, leave and rejoin. By analogy with a maritime chart, the filesharing overlay networks constitute ‘archipelagos’ of tiny islets, expanding and contracting in area as the tides ebb and flow.

The packet-switching Internet began as a TCP/IP⁵ overlay on the circuit-switching PSTN, or Public Switched Telephone Network. The World Wide Web is a HTTP overlay on the Internet and, in turn, the various P2P filesharing networks are overlays on the Web. However, what insiders call ‘The Grid’ is not an overlay network at all, nor an archipelago, but a set of large ‘islands’ of networked high performance computing (HPC) resource clusters that form part of the server-dominated core of the Web. Since the mid-1990s, academic, government and corporate research groups have been coalescing into separate ‘virtual organisations’ (VOs) to pool and interwork their resources of computing, storage, software, databases, bandwidth, instruments, and sensors. The enabling software is called ‘grid middleware’ because it sits between grid users and the HPC resources they call on by submitting jobs (called ‘workflows’) to whichever grid VO they are accredited to. One of the largest and best-known grids is the Worldwide LHC Computing Grid (WLCG), to which up to 6,000 scientists already have begun submitting workflows for accessing and analysing the 15 petabytes of particle collision data that emanate annually from the Large Hadron Collider (LHC) project at CERN in Switzerland.

Like most grids, the WLCG is built out of components drawn from the Globus Toolkit and Condor middleware solutions. But the hybrid CERN grid will strain to control more than a few thousand computers. This is due to reliance on key centralised components: the Metacomputing Directory Service for Globus and the Central Manager for Condor. In all these island grids there also is the problem of incentives to form a VO in the first place, and later to resolve conflicts between the partners over equitable prioritisation of access to the heterogeneous stock of common HPC resources.

The P2P *network* archipelagos discussed above are designed for filesharing, not grid computing. But since the SETI@home project got underway in 1999, the Web’s edge has contained *grid* archipelagos as well; they are called ‘distributed’ or ‘desktop’ grids. A team member of the Search for Extra-Terrestrial Intelligence had been struck by how often the screensavers would appear on SETI computers, indicating that their central processing units (CPUs) were idle. This waste of compute cycles would be a worldwide phenomenon, and he reasoned that software code incorporated within screensavers could scavenge these idle cycles and use them to perform the repetitive computations involved in analysing deep space radio sources for signs of artificial-looking signal patterns. The resulting software is called BOINC, the Berkeley Open Infrastructure for Network Computing.

Apart from SETI@home, Wikipedia (2007c) now lists 16 current BOINC-based projects and a further 36 under development. Some 430,000 edge clients, embedded in screensavers, are now performing repetitive tasks like protein folding, theorem solving and rendering film animations – ‘parallel computing’ problems with no complex cause-effect feedback loops. However, each project needs clusters of core servers to manage the clients (typically installed on personal computers owned by volunteers), which raises the familiar problems of

⁴ Wikipedia (2007b) states that the best examples are Gnutella, Kazaa or eMule with Kademia, whereby Kazaa has still a central server for logging in. eDonkey2000/Overnet, Gnutella, FastTrack and Ares Galaxy have summed up approximately 10.3 million users (as of April 2006, according to slyck.com).

⁵ TCP/IP=Transport Control Protocol/Internet Protocol.

centralisation and incentives. Wikipedia (2007c) also explains that the BOINC Credit System “... is designed to avoid cheating by validating results before granting credit. A credit management system helps to ensure that users are returning results which are both scientifically and statistically accurate. Online distributed computing is almost entirely a volunteer endeavor.” Thus there are likely to be ultimate limits on how extensive a BOINC grid can become.

Grid Future: A World Virtual Machine?

In 2006 the number of digital devices connected to the Internet reached one trillion.⁶ Computer scientists now are seriously considering the possibility of merging all the island and archipelago grids and networks with the vast ‘Internet ocean’, i.e. the planet’s billions of standalone computers. The boffins want to birth a powerful World Virtual Machine (WVM) that is accessible by anyone, anywhere, anytime. This requires an infinitely scalable Global Open Grid capable of creating and sustaining a market in ‘grid fabric services’ generated by the computing, storage, software, database, bandwidth, scientific instrument, and sensor array capital assets known as ‘grid fabric resources’. It will have to be a secure distributed peer-to-peer grid that is resistant to free-riding and features efficient resource discovery, resource allocation and job scheduling mechanisms. Above all, it will have to offer powerful *incentives* for users and providers to participate, e.g. expectations of increases in utility and/or profit.

Any future WVM is likely to emerge as The Grid to which most consumer-only users submit their workflows, each structured as a series of calls to various Web Services, as defined in Wikipedia (2007d). A smaller number of users would be both providers and consumers, and the minority would comprise large provider-only operators of standalone grids, data centres and commodity box clusters located at the Web’s core. The grid fabric resources would reside in, or behind, the providers’ front-end host devices. These devices would comprise Grid nodes on an unstructured global P2P Web overlay network. Such resources eventually could include all standalone grids, data centres and clusters, but also servers, PCs, storage device farms, libraries of Web Services, software archives, data repositories, sensor arrays, scientific instruments ... in fact, any resource that can be digitally represented and whose services are valued by users when offered *gratis*, cooperatively or (in most cases) for sale by their providers.

With almost all of the trillion-plus Internet ocean digital devices having one or more underutilised grid fabric resources – especially overnight, as their owners sleep – the prospect looms that every one of them potentially is a provider, not just a consumer, of grid fabric services. Given an efficient resource broking and job scheduling grid middleware, an infinitely scalable P2P Web overlay network, the right incentive structure, and a workable micropayments system, the ‘WVM problem’ could be solved in fairly short order.

Three disciplines currently claim to offer WVM problem-solving grid middleware designs:

- from mathematics, optimal control theory;
- from biology, self-organising systems; and

⁶ “Someday soon, more than one million businesses will be connected to more than one billion people by one trillion devices” is a famous prediction made by IBM’s Louis Gerstner in 1999. [<http://www.industryweek.com/ReadArticle.aspx?ArticleID=517&SectionID=38>] That day arrived some time during 2006, according to IBM Canada’s Dan Fortin. [<http://feeds.feedburner.com/TorontoBoardOfTrade>]

- from economics, Hayekian catallaxy.

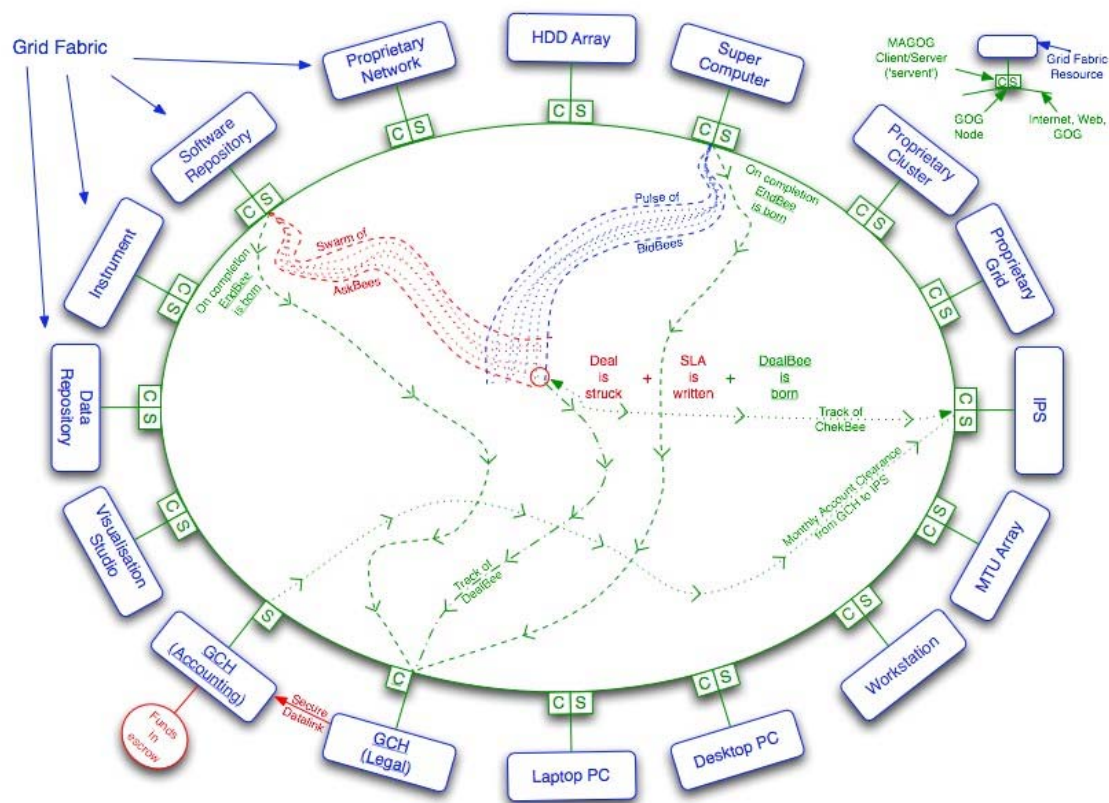
Optimal control theory is a dead end because it aims to maximise some metric that is a function of the system's global state ... which would be *unobservable* in a Global Open Grid because its topology and state variables will be so complex and dynamically changing that the information would be obsolete before very much of it could be collected by more or less centralised servers, no matter how strategically situated on the continents.

Self-organising systems are more promising because biology offers many useful metaphors. IBM's Autonomic Computing concept is a self configuring/healing/organising/protecting (CHOP) cyber-organism, whose architecture is based on the autonomic nervous system (IBM 2005). Then there are the adaptive multi-agent system (MAS) collectivities used in systems biology because they can mimic what happens in open dynamic environments. Biological systems involve numerous asymmetric interactions between autonomous and heterogeneous entities facing constraints that define a 'fitness landscape'. They act directly on the environment but can experience a collective emergent phenomenon that indirectly modifies the agent state. Ritchie & Levine (2003), for instance, offer a middleware component based on the foraging behaviour of an ant colony. Usually it is the lack of incentives that is fatal: neither WVM users nor providers will be (directly) driven by the standard biological imperatives of obtaining food, defending territory, attracting mates, repelling rivals, avoiding predators, etc.

That just leaves economics, with its vision of an economy as a dynamic system of interacting markets, in which relative prices automatically flex to equilibrate supplies and demands for goods and services. With the *flow prices* of grid fabric *services* being revealed by the catallactic 'higgling of the market', the *stock prices* can be inferred to also value (using interest and depreciation rates) the grid fabric *resources* that generated those services. Such changing capital asset valuations in the WVM economy, when compared with the capital costs of acquisition determined in the general economy, will signal providers to invest in extra capacity, to maintain or increase current levels, or to allow their surplus grid fabric resource stocks to depreciate until such time as the value/cost ratios rise, i.e. changes in the 'Tobin's q' measure that was introduced by Tobin (1969). So, by creating a popular *economic* resource broker/job scheduler for The Grid, middleware designers also will cause to emerge an *incentive structure* for grid fabric resource owners to alter their output of grid fabric services, in line with market demands.

For those computer scientists at the leading edge of Grid research, Hayek's catallaxy is their paradigm of choice. They recognise that, of all the unrealistic, even illogical, general equilibrium axioms, assumptions and corollaries listed in Appendix A, four are fatal. Markets for grid fabric services are not complete, perfect information and the omniscient auctioneer rule out full decentralisation of servers, and the 'strong gross substitutes' axiom cannot apply in a market where processing, memory, software, bandwidth, etc. are actually 'strong gross complements'. Appendix B shows that a large EU-funded research project called CATNETS has been investigating catallaxy in the grid context. Richardson (2007) and Internet Centre (2007) show that Imperial College London have their own catallaxy-based research project underway; it's called MAGOG: Middleware for Activating the Global Open Grid. This project is described below, and Appendix C lists nine substantive differences between MAGOG and CATNETS.

The following diagram is a schematic of the Global Open Grid (GOG) peer-to-peer network:



In reality, the GOG overlay on the Web will resemble a mesh, not a ring. The point is simply that any grid fabric resource device can send messages (called GridBees) to any other device on the ring, provided all peers have the MAGOG middleware [C,S] servent installed. Although only 16 nodes of GOG are shown, each peer represents a *class* of devices and the total number on the unstructured P2P overlay network can scale indefinitely. Significantly, two of the [C,S] combinations represent classes of business firms. At the far right is IPS=Internet Payment System and to the left of bottom is GCH=Grid Clearing House. Every market participant must register as a member with the IPS and GCH of their choice before commencing trading. For security, each GCH separates its [C=Client] software that registers legal Service Level Agreement (SLA) deals contracted by members from its [S=Server] software that notifies monthly accounting balance clearances to each member's nominated IPS. The entire network is suffused with competition from the start, and this will intensify as GOG grows along its sigmoid expansion curve.

The rationale for P2P schemas like Chord, Pastry, Tapestry, and Symphony is to allow peers (here, buyers and sellers of grid fabric services) to find each other, but GOG does not need them. An early search method – one peer flooding the network with identical messages until the right machine responds – is effective but takes a long time, clogs the communication channels, constitutes a breach of netiquette, and is now all but outlawed by the Internet Protocol authorities. Instead of this inefficient 'single message flooding' search process, the MAGOG middleware uses 'double message flooding', which multiplies the chances of users and providers getting together to do business. In other words, *both* seller and buyer peers launch swarms or pulses of messages (AskBees and BidBees, respectively) in the knowledge that two messages inscribed with compatible SLAs are highly likely to meet at *some* GOG node and later close a deal. In fact, this could happen at *any* GOG node because every MAGOG servent contains a software component (called a 'Pub') where wandering GridBees can meet and look each other over.

Typical ‘time-to-live’ (TTL) or ‘hopcount’ default values on the Internet are 30 hops between machines, but MAGOG further restricts the TTL of a message to only seven hops. This is made possible by the famous ‘small worlds hypothesis’ which states that, in any social network, there are only ‘six degrees of separation’ between the source and destination of a message (National Science Digital Library 2007). Near the centre of the diagram above, a Super Computer has sent a pulse of BidBees to inquire about some particular software not held locally. Eventually this ever-increasing number of buyer messages intersects with a growing swarm of AskBees originating from a seller’s Software Repository. (In fact, this action takes place *inside the Pubs* of various [C,S] MAGOG servents at nodes on the GOG network.) With high probability, two ‘deal-compatible’ GridBees *will* meet before expiring after seven hops – this takes place inside the Pub marked with a small red circle.

Each node’s Pub software component checks all arriving GridBees for deal-compatibility by ensuring that SLAs match, i.e. the grid fabric service(s) being sought and offered are identical. Also, the buyer’s bid price must equal or exceed the seller’s ask price, for otherwise no deal is possible. Compatible GridBee couples then set off on their ‘mating flight’ back to the seller’s servent, from which a ChekBee is sent off to the buyer’s IPS for identity and credit checks. If these come back marked ‘OK’, the deal is struck (midway between the bid and ask prices) and a DealBee flies off with the SLA and agreed price to each party’s GCH for legal registration. Upon successful completion of the contracted grid fabric services, both parties send an EndBee to GCH (Legal), which notes this on the SLA and authorises GCH (Accounting) to debit the buyer’s and credit the seller’s accounts. On a monthly basis, the GCH sweeps all its members’ account balances across to the IPS, which invoices the net debtors and pays out the net creditors.

There are no ‘mobile intelligent agents’ in this unique Grid middleware architecture; all processing is done by the MAGOG servents, which can be installed as the ‘front-end’ of *any* Web-enabled device, cluster, array, or grid. In all there are just five ‘species’ of GridBee. The BidBee pulses and AskBee swarms are legion, but only a few are destined to mate. Their single offspring (a successful SLA) generates just one ChekBee, one DealBee and two EndBees. The efficient 7-hop double message flooding procedure is enabled by every GOG network node maintaining its own (dynamically changing) 7-member ‘leafset’. This lists the seven other MAGOG servents from which GridBees have been most frequently arriving and finding a matching partner.

To understand how the messages propagate, consider the buyer’s actions: she simply loads MAGOG, composes her workflow, then waits for the final result(s). During that short delay, her servent deconstructs the series of commands into its corresponding set of calls for various grid fabric services to be performed. Any not available on her own local grid fabric resources will have to be bid for, so the middleware inscribes certain characteristics of each missing service onto a BidBee message. (At this point, the buyer may have specified that she be interrupted to manually approve or alter the price bids.) Then one copy of each BidBee is dispatched to the seven GOG nodes on her servent’s leafset, in whose Pubs all the remote MAGOGs try to find a match. If unsuccessful, each Pub (temporarily) retains one copy of her BidBee in case a deal-compatible AskBee happens to arrive later. Meanwhile, each of the seven nodes has sent copies of the BidBee to all seven nodes on their own leafsets. To this point the BidBee has visited $7 \times 7 = 49$ Pubs in two hops and, in the unlikely event that no match is ever found, the message will in total have been exposed in $7^7 = 823,543$ Pubs during the seven hops before its TTL expires.

Finally, what is it that causes the catallaxy to emerge? A stable vector of prices covering every conceivable grid fabric service (Hayek’s spontaneous order) is most likely to appear

because disappointed buyers and sellers (those whose flights of GridBees have not yet found willing sellers/buyers) will *experiment* – manually or via a software plugin – by slightly raising their latest bid prices and tentatively lowering their current ask prices, respectively. Of course, such experimentation (Hayek’s learning via false trading) will be constrained by the buyers’ computing budgets and the sellers’ cost structures (Hayek’s local knowledge). After some price gyrations, each community of grid participants eventually will hit upon the approximate ‘ruling’ or ‘equilibrium’ or ‘market-clearing’ prices of those particular grid fabric resources in which they are most interested (Hayek’s catallaxy).

In this endeavour, they (and any bid/ask negotiating strategy software they plug into their open-source MAGOG servents) will be assisted by the availability of live grid market information XML news and data feeds. This confident expectation is based on the success of market data, news and information sources like the Bloomberg Professional Service, Reuters Global Markets Reports Service and the Chicago Mercantile Exchange Historical Data Service in real-world equity, fixed interest, commodity, futures, derivatives, and foreign exchange markets. Their Global Open Grid market counterparts will do deals with the various Grid Clearing Houses for the provision of anonymised real-time price and quantity information. These data will be collected, second by second, from the Service Level Agreements that the GCHs are continually registering over the various classes of grid fabric service being traded.

Simulating a Digital Economy: Recent Research Findings

This paper has argued that an economic metaphor will outperform the mathematical and biological metaphors currently used for resource discovery, resource allocation and job scheduling in computational grids. Walrasian general equilibrium was rejected as a paradigm for efficiently operating a commercial Global Open Grid or World Virtual Machine. Instead, the Hayekian catallaxy paradigm will be utilised for facilitating the decentralised discovery and allocation of grid fabric resources, together with efficient job scheduling, in an unstructured peer-to-peer Web overlay network that scales indefinitely, having no single point of failure nor any potential bottlenecks liable to choke off the flow of critical control information.

This *desideratum* presently is being researched by computer scientists and heterodox economists in the Department of Computing at Imperial College London. The GOG-MAGOG Project has seven separate phases, of which the only the first has been completed to date: To *design* a grid middleware that will enable the autonomous emergence of a globe-spanning grid and an international market for spot and forward trading in grid fabric services, thereby also stimulating a more competitive eCommerce in other goods and services to develop on a level playing field. The second phase is currently underway: To *simulate* growth of the GOG network (and operation of its MAGOG nodes) in a mathematical model, thus confirming the small worlds principle and checking that Hayek’s catallaxy guarantees ‘order for free’ – as is known to occur in similar complex adaptive systems.

Within the second phase, although the GOG network’s growth has not yet been simulated, its structure and MAGOG’s stability are investigated in two recent papers. Harder (2007) defines GOG as a connected graph of vertices (nodes) and edges (links). Seller nodes offer the services of central processing units (CPU compute cycles) for minimum ask prices and buyer nodes communicate maximum bid prices, with all nodes using double message flooding via neighbouring links and each message having a preset ‘hopcount’ or time-to-live (TTL). Nodes are ‘unsatisfied’ while searching, then switch into a ‘satisfied’ state once a deal is consummated. The simulations were run in the static context only “... without a state change

of the nodes once they have reached a satisfied state. This showed that the nodes are able to set prices to ensure that the maximum number of nodes reaches the satisfied state.” Also, “in such a world ... areas of the network can lose out as they fail to find matches before the TTL times out.”

Harder & Martínez (2008) improved and extended this network study using the *igraph* package of Csàrdi & Nepusz (2006). The authors took a ‘2 x 2 x 2 x 2 x 2’ approach in testing two connected graphs (random vs scale-free), two P2P middleware designs (naive vs MAGOG), two initial price setting rules (fixed vs random), and two subsequent price adjustment rules (mean price vs delta price) in two contexts (static vs dynamic). Naive nodes immediately pass on unmatched messages; MAGOG buffers them in case matching messages arrive before their 10-hop TTLs expire. (In the dynamic context, after ‘a time’ nodes change from being satisfied back to the unsatisfied state, while those that have stayed unsatisfied ‘too long’ alter their prices.) Networks ranging from 4,096 to 1,048,576 nodes were simulated, with several system efficiency metrics being computed: ratio of satisfied nodes, deals per node, messages per buffer, average execution time, etc.

The authors conclude “We have shown ... that the system actually tends to an ‘equilibrium’, where the ... utilization of resources, volume of contracts and messages in buffers get stable, generally independently of the size of the network. The system takes longer to reach equilibrium for the MAGOG version, although ... parameters achieve more stable values once the equilibrium has been reached.” Unfortunately, the authors did not put any ceiling on bid prices (size of computing budgets), nor any floor under ask prices (unit costs of CPU compute cycles). As a result, the evolution of deal prices trended ever-upward for the random connected graph network and ever-downward for the scale-free alternative. When this oversight has been corrected, we expect the time-series plots of deal prices to instead display a horizontal trend or ‘flatline’ – after some initial transients due to Hayek’s bouts of learning through ‘false trading’.

Conclusion

With the aim of realising a World Virtual Machine that uses MAGOG middleware, the Imperial College Internet Centre has submitted a funding proposal to the Engineering and Physical Sciences Research Council (EPSRC). The remaining five phases are: To code MAGOG, using mainly the Java language and Web Services; To test this initial code on a local computer cluster (simulating the GOG network) and make all necessary modifications; To test MAGOG Alpha on the PlanetLab testbed, currently comprising 825 nodes at 406 sites (PlanetLab 2007); To test for autonomous growth of its own P2P network; To release MAGOG Beta onto the Internet, evaluate responses from serious users, then incorporate all necessary changes; and To achieve planet-wide distribution of the MAGOG middleware servent.

This final step will involve free Internet downloads from multiple websites on all continents, thus encouraging the world-spanning Global Open Grid (GOG) to emerge, then enter its initial phase of rapid growth in take-up. This is just *one possible future* for the Internet, Web and Grid. A mature planet-wide digital economy⁷ is envisaged, one in which not only most grid fabric services are commercially traded (with some freely offered and others

⁷ The EPSRC Digital Economy Programme has awarded its first-ranked Connecting Communities for the Digital Economy research grant to Professor John Darlington, Director of the Imperial College Internet Centre, as reported at <http://gow.epsrc.ac.uk/ViewPanelROL.aspx?PanelId=4618&RankingListId=6079>.

cooperatively shared) on a global P2P network, but also *any* commodity whose characteristics and bid or ask price can be inscribed on an XML message. This paper has demonstrated that, despite mainstream economics often sidelining much that does not fit within the Neoclassical School's general equilibrium orthodoxy, a heterodox economic paradigm *is* being taken seriously in mainstream computer science.

Appendix A: Axioms, Assumptions and Corollaries of General Equilibrium

- complete markets: at time zero all economic agents make trades in both spot and futures markets for commodities that maximise their lifetime utility and profits, where these commodities differ over time, space and all possible future events (changes in weather and climate, wars, insurrections, earthquakes, tsunamis, etc.);
- perfect rationality: all agents maximise their utility *qua* consumers and their profits *qua* producers (n.b. this is an irreducibly non-computable problem);
- perfect foresight: we live in an ergodic world that has calculable, insurable risk, but no irreducible uncertainty;
- perfect competition: perfect information (all agents know all prices), homogeneous products (perfect substitutes), free entry and exit in all industries (no barriers), atomistic agents (all are price-takers), and equal access (to all production technologies);
- perfect mobility of the factors of production over space, e.g. no controls over immigration or capital flows;
- perfect divisibility: one can purchase an infinitely small quantity of any commodity;
- a central auctioneer conducts the *tâtonnement* (or groping) towards the 'true prices';
- no 'false trading': no deals may be consummated at 'false prices' (alternatively, all products are 'contingent commodities' in that 'false trades' can be *unscrambled* once the auctioneer announces the economy's GE vector of 'true prices');
- only relative – not absolute – prices, are announced by the auctioneer (although all prices can be expressed relative to some chosen good or service, and this *numéraire* commodity may be gold, silver, pigs, peanuts, labour, jam, ...);
- agents have unlimited computational capacity;
- costless enforcement of contracts and property rights;
- no money: there is no need for a medium of exchange or a store of value (even if the *numéraire* is gold or silver, that does not make it money);
- no difference between private and social rates of return or discount rates;
- no loan defaulters: all debts are paid on time and in full;
- no "one-way arrow of time": logical – not historical – time passes;
- no hysteresis or path dependency, e.g. bankrupt agents are costlessly resurrected;
- no transactions, information, search, assimilation, computation, or learning costs;
- no monotonically increasing returns to scale or scope;
- no externalities: there are no benefits (costs) beyond those privately enjoyed (borne) by each individual agent;
- non-satiation: no matter how much of any commodity an agent is consuming, he/she always wants more of it;
- all consumer preference maps are homothetic, fixed at birth and do not change with age, life experience or advertising; and
- all commodities are 'strong gross substitutes': all agents' preference maps and production sets are convex (extreme example: if its price relative to food went down far enough, agents would choose to eat poison instead).

Appendix B: CATNETS: An EU Project for Implementing Catallaxy

The European Union has funded the CATNETS project in Framework Programme 6. See Eymann, *et al.* (2003) for a discussion.

The objective is to determine the applicability of a decentralized economic self-organisation mechanism for resource allocation in application layer networks (ALNs). One of the main problems of future grid network technology is the efficient provisioning of services to clients by a scalable and dynamic resource allocation (matching) mechanism. The CATNETS project investigates a ‘free market’ economic self-organisation approach, the ‘Catallaxy’, as the basis for realizing self-organising resource allocation in ALNs. CATNETS targets are (1) to implement and evaluate Catallactic resource allocation as a ‘proof-of-concept’ prototype in a real ALN and (2) to test the properties and limitations of the Catallaxy approach using network simulation. Both approaches will share a common design of Catallaxy mechanisms applied to computer networks, including the appropriate messaging protocols needed to achieve dynamic negotiation and self-organisation, metadata and the utility functions of the network agents.

The term ALN integrates different Internet overlay network approaches, like grid and P2P systems. The allocation of resources in these networks (e.g. matching of demand and supply, deployment of service instances and service discovery) can principally be conducted in a centralised (e.g. using resource brokers) or in a decentralized fashion (e.g. using self-organising mechanisms). Centralised approaches reach their limits with increasing network size and growing numbers of elements; self-organising approaches thus gain attention ... The performance measurements of the proof-of-concept implementation will be compared against the simulation results, with the goal of being able to make a substantiated statement on the applicability of economic self-organisation as a major component of ALN networks. A positive evaluation of the Catallactic approach would have a high potential impact, with new possibilities for resource brokering in future ALNs, and maybe for self-organisation in computing in general.

Appendix C: Comparison of MAGOG with CATNETS

The MAGOG Project was begun in ignorance of the prior existence of the CATNETS Project, which serendipitously *also* is based on the catallaxy paradigm. There are nine salient differences between these two independently conceived grid middlewares:

CATNETS is designed for use in scientific, engineering, enterprise, etc. grids, whereas the aim of MAGOG is as general as its name implies: ‘Middleware for Activating the Global Open Grid’. GOG does not yet exist; it will be activated as an unstructured P2P network as grid fabric service user/providers download and install the MAGOG servent.

CATNETS utilises at least eight different types of complex mobile intelligent software ‘economic agents’, whereas MAGOG utilises five types of simple XML messages called ‘GridBees’. All intelligence resides in the MAGOG servents installed in host devices, each comprising a node of the Global Open Grid (GOG).

The CATNETS economic agents have a complex ‘set of generic negotiation mechanisms’ built in. In MAGOG the middleware simply ‘splits the difference’ between a minimum ask-price that is less than or equal to a maximum bid-price for the same grid fabric service. It is open to any MAGOG user to ‘plug in’ code that encapsulates a different negotiating strategy.

In CATNETS, user applications sometimes ‘will make request for resources to the base platform, which will in turn, forward it to the Catalactic middleware’ and this ‘will require the implementation of a connector’. In other scenarios the application ‘will make requests directly to the Catalactic middleware’. In MAGOG, by contrast, GridBee messages released by a user and a provider may meet in *any* middleware Pub, realise they match, then later close the deal.

CATNETS supports two generic markets, one for resources and the other for services. In MAGOG there is just one: the market for grid fabric services. Also, the deal-making capability of MAGOG's GridBees is so general that *any* good or service can be traded using this middleware.

CATNETS may require Grid Containers and Globus Resource Allocation Managers from GlobusToolkit4, whereas MAGOG requires no resource allocation managers (because the emergent Catallaxy takes care of all that) and uses no components from GT4.

In CATNETS ‘Agent's roles ... as buyer, seller, broker, auctioneer, depends on the specific market being implemented.’ Thus different versions must be crafted for many specialised applications, such as Cat-COVITE for the building and construction industry. By contrast, the same MAGOG middleware is utilised by any and all applications.

CATNETS requires ‘Implementation of instrumentation, measurement, analysis and visualization of agent's behaviour in a fully distributed architecture’. In MAGOG there are no intelligent agents, just XML messages.

CATNETS ‘... allows the implementation of diverse market models on top of a single middleware infrastructure’, where this diversity refers to various computerised auctions and commodity trading systems. MAGOG does not need to implement any of these models of the real world economy, because it operates in the real world economy.

References

CacheLogic (2006). <http://www.cachelogic.com/home/pages/isp/p2ptoday.php>

Csàrdi, G & Nepusz, T (2006). “The igraph software package for complex network research”. *InterJournal Complex Systems*. 2006: 1695.

Torsten Eymann, Michael Reinicke, Felix Freitag, Leandro Navarro, Oscar Ardaiz, and Pau Artigas (2005). “A Hayekian self-organization approach to service allocation in computing systems”. *Advanced Engineering Informatics*. 19(3): 223-233, July 2005.

Harder, Uli (2007). “A Peer-to-Peer Market for Grid Computing”. *Proceedings of the 6th Workshop on Process Algebra and Stochastically Timed Activities*: 105-109. Imperial College London. 26-27 July 2007. <http://pastaworkshop.org/2007/proceedings>

Harder, Uli & Martínez, Fernando (2008). “Simulation of a Peer-to-Peer Market for Grid Computing”. *Proceedings of the 15th International Conference on Analytical and Stochastic Modelling Techniques and Applications*. Nicosia, Cyprus. 4-6 June 2008. *Lecture Notes in Computer Science*. 5055: 234-248. Springer-Verlag.

- Hayek, Friedrich A (1935a). Socialist calculation I: the nature of the problem.
- ___ (1935b). Socialist calculation II: the state of the debate.
- ___ (1940). Socialist calculation III: the competitive solution.
- ___ (1937). Economics and knowledge.
- ___ (1945). The use of knowledge in society.
- ___ (1948). The meaning of competition.
- In Bruce Caldwell (General Editor), "The Collected Works of F. A. Hayek", 12 volumes. Chicago: University of Chicago Press, 1990-2007
- ___ (1973). *Law, Legislation and Liberty: A New Statement of the Liberal Principles of Justice and Political Economy, vol. I, Rules and Order*. London, Routledge & Kegan Paul; Chicago, University of Chicago Press, 1973
- Holland, John (1986). *Induction: Processes of Inference, Learning, and Discovery*. [with Holyoak, K J, Nisbett, R E and Thagard, P R] MIT Press, 1986
- IBM (2005). "An architectural blueprint for autonomic computing". 3rd edition, June 2005 <http://www-03.ibm.com/autonomic/pdfs/AC%20Blueprint%20White%20Paper%20V7.pdf>
- Internet Centre (2007). Internet Economics" <http://www.internetcentre.imperial.ac.uk/projects>
- Lange, Oskar (1936). "On the economic theory of socialism: Part I". *Review of Economic Studies*. 4(1), 53-71, 1936
- National Science Digital Library (2007). <http://expertvoices.nsdlib.org/cornell-info204/2007/04/27/wikipedia-as-a-small-world-network>
- PlanetLab (2007). <http://www.planet-lab.org>
- Polanyi, Michael (1958). *Personal Knowledge: Towards a Post Critical Philosophy*. London: Routledge, 1958
- Richardson, Colin (2007). "Growing the Global Open Grid: Design Brief and Middleware Architecture". Technical Report. The Internet Centre, Imperial College London.
- Ritchie, G & Levine, J (2003). "A fast, effective local search for scheduling independent jobs in heterogeneous computing environments". In J Porteous (Ed.), *Proceedings of the 22nd Workshop of the UK Planning and Scheduling Special Interest Group (PLANSIG 2003)*: 178–183, 2003
- Tobin, J. (1969). "A General Equilibrium Approach to Monetary Theory." *Journal of Money, Credit and Banking*, 1969, 1(1), 15-29.
- Wikipedia (2007a). <http://en.wikipedia.org/wiki/XML>
- ___ (2007b). http://en.wikipedia.org/wiki/File_sharing
- ___ (2007c). http://en.wikipedia.org/wiki/BOINC#Projects_using_BOINC_Framework
- ___ (2007d). http://en.wikipedia.org/wiki/Web_services