Market-Based Systems

University of St Andrews
Department of Computer Science
Distinguished Lecture Series, March 2008

Lecture 2

Dave Cliff

Department of Computer Science
University of Bristol

dc@cs.bris.ac.uk

The Contents

NB This is a 3-hour partial tutorial overview of Market-Based Systems …in three 60min chunks

- partial as in incomplete: we can't cover everything in three hours
- partial as in biased: this is my version of the story…

- Lecture 1: Rationale and Background

  Here we'll find out why computer scientists should care about market-based systems, review some notable applications, and also cover some of the background economics. They call economics "the dismal science" for a reason, so that background economics stuff won't delay us too long...

- Lecture 2: Artificial Trading Agents for Fun and Profit

  This lecture tells the story of some of the best-known algorithms used for autonomous "trader-robots", and how they were found to consistently beat human traders.

- Lecture 3: What's hot, what's not, and where next: Tales from the City
  ----
  Looks at work on automatic optimization and design of trader-agents, and online market mechanisms, with particular reference to the current hot topics in the automated trading technology in the financial markets.

Recap: Gode & Sunder results

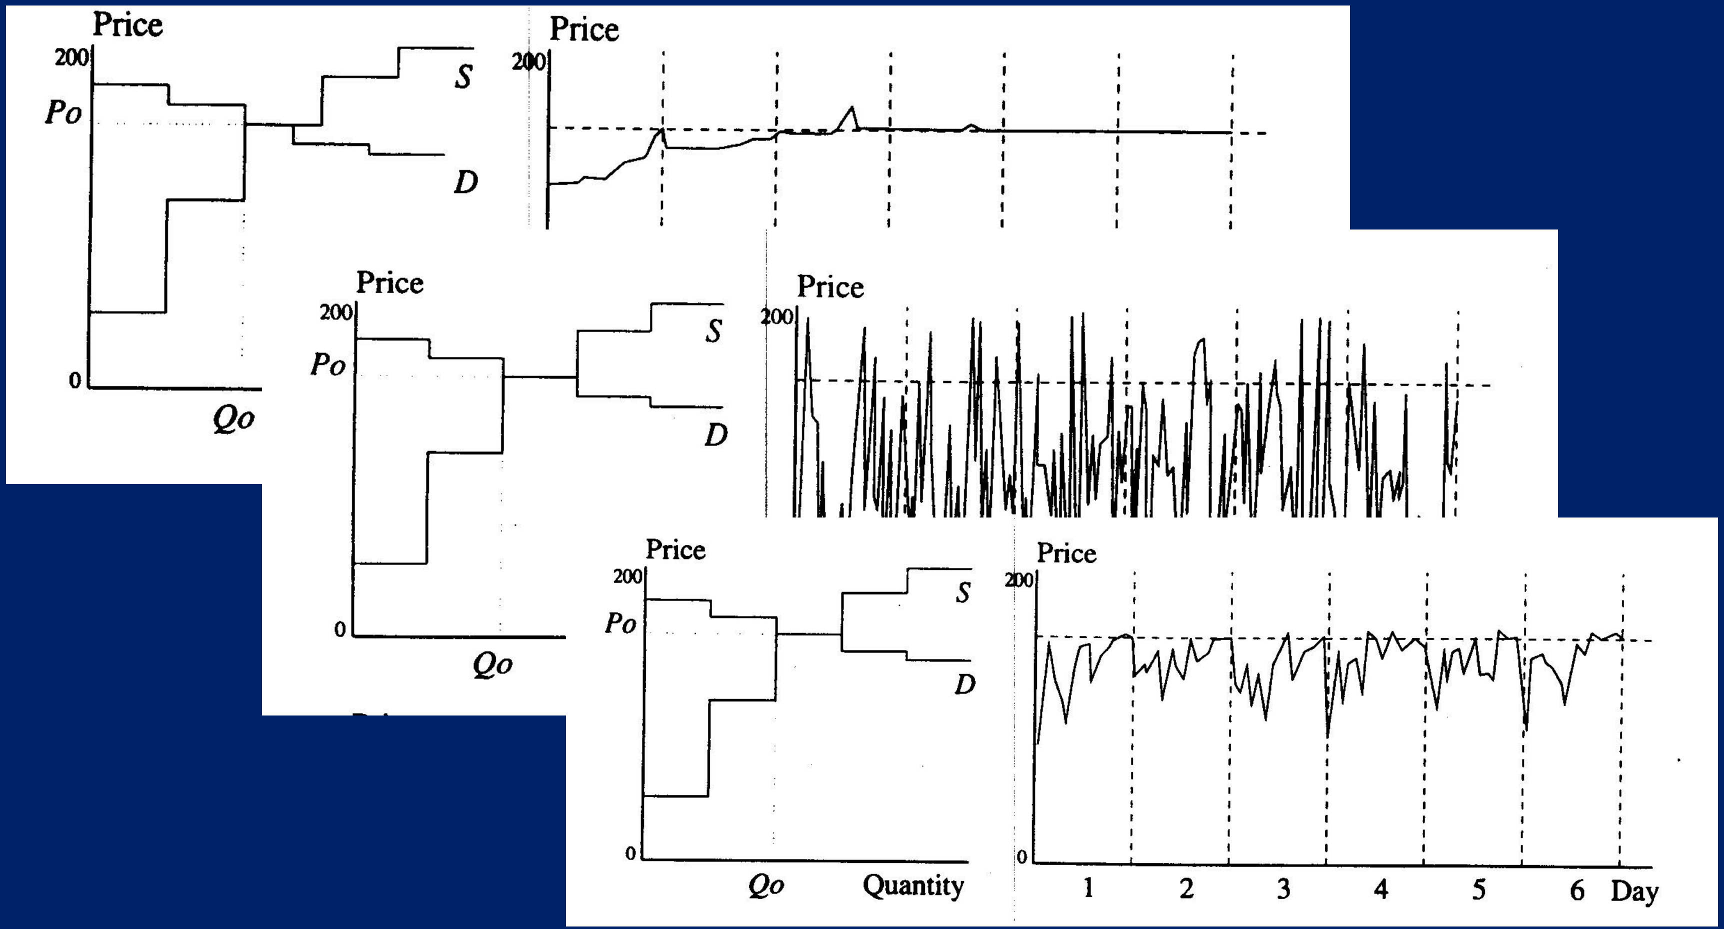

More than zero, actually

- D Cliff (1997) Minimal-intelligence agents for bargaining behaviors in market-based environments. HP Labs Technical Report HPLTR97-91. http://www.hpl.hp.com/techreports/97/HPL-97-91.html

- Critique of Gode&Sunder'93
  - Analyzed PDFs of ZI traders and proved that Gode&Sunder's results were artefactual
  - Predicted conditions when ZI-C traders would fail to equilibrate
  - Implemented ZI-trader system to empirically demonstrate failures predicted by analysis

- Developed ZI-Plus (ZIP) traders
  - Widrow-Hoff momentum learning
  - Demonstrated to succeed in markets where ZI-C's had failed
  - Human-like market dynamics in CDA and "retail" auctions
  - "open source" (published the C code) in 1997

Zero intelligence: not magic for the sellers

- Simple qualitative thinking about probabilities in ZI-C systems shows how they work…
- ZI-C seller generates offer prices at random, in the range from the minimum seller limit price S_{min} up to the maximum price allowed in the system, P_{max}.
- and so that gives us a probability density fn for the offer prices…

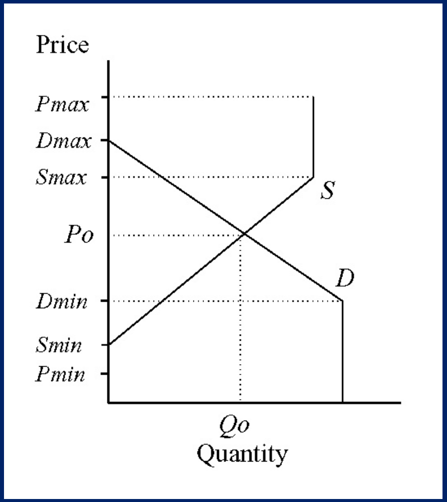

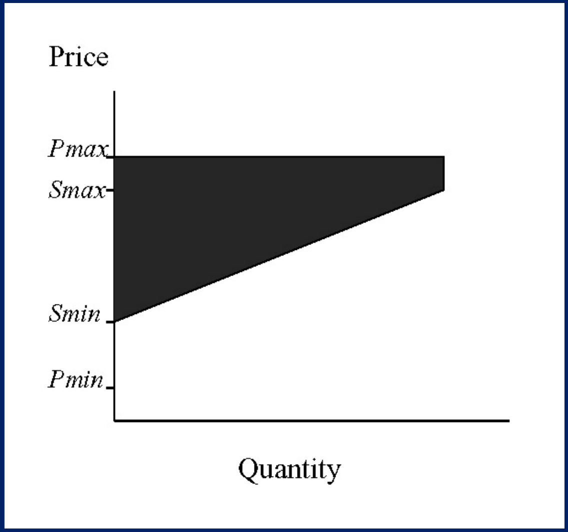

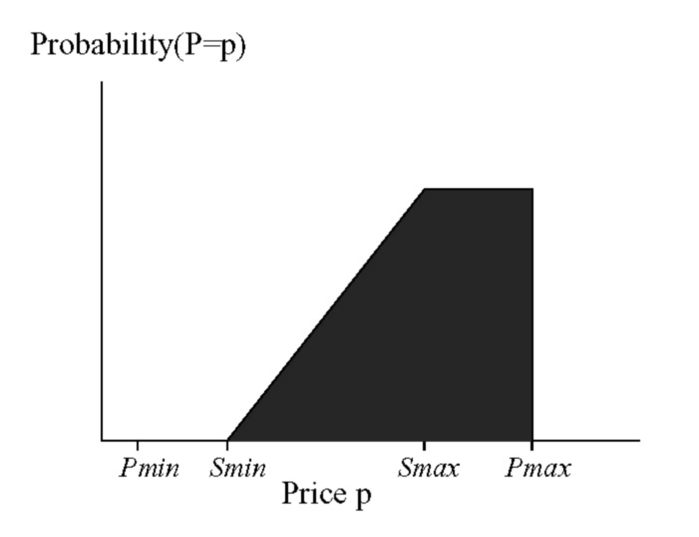

Zero intelligence: not magic for the buyers

- Same for buyers:

- ZI-C buyer generates offer prices at random, in the range from the maximum seller limit price D_{max} down to the minimum price allowed in the system, P_{min}.

- and so that gives us a probability density fn for the bid prices…

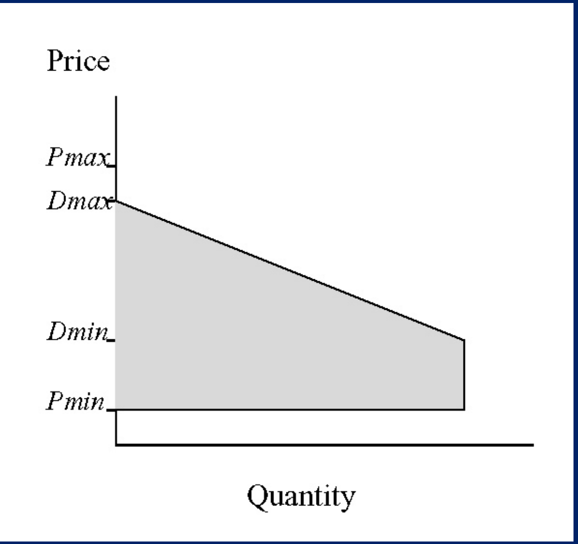

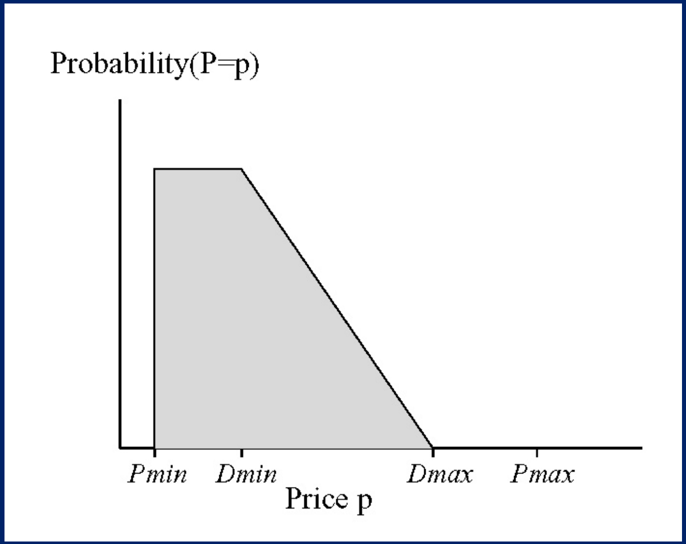

Zero intelligence: not magic at all

- Transactions only occur when a bid price and an offer price cross, so the pdf for transaction prices is going to be defined by the intersection of the bid-price pdf and the offer-price pdf
- Expected transaction price E(p) is in this case close to the equilibrium price, because the pdf is symmetric about its peak, and the peak is set by the equilibrium price.

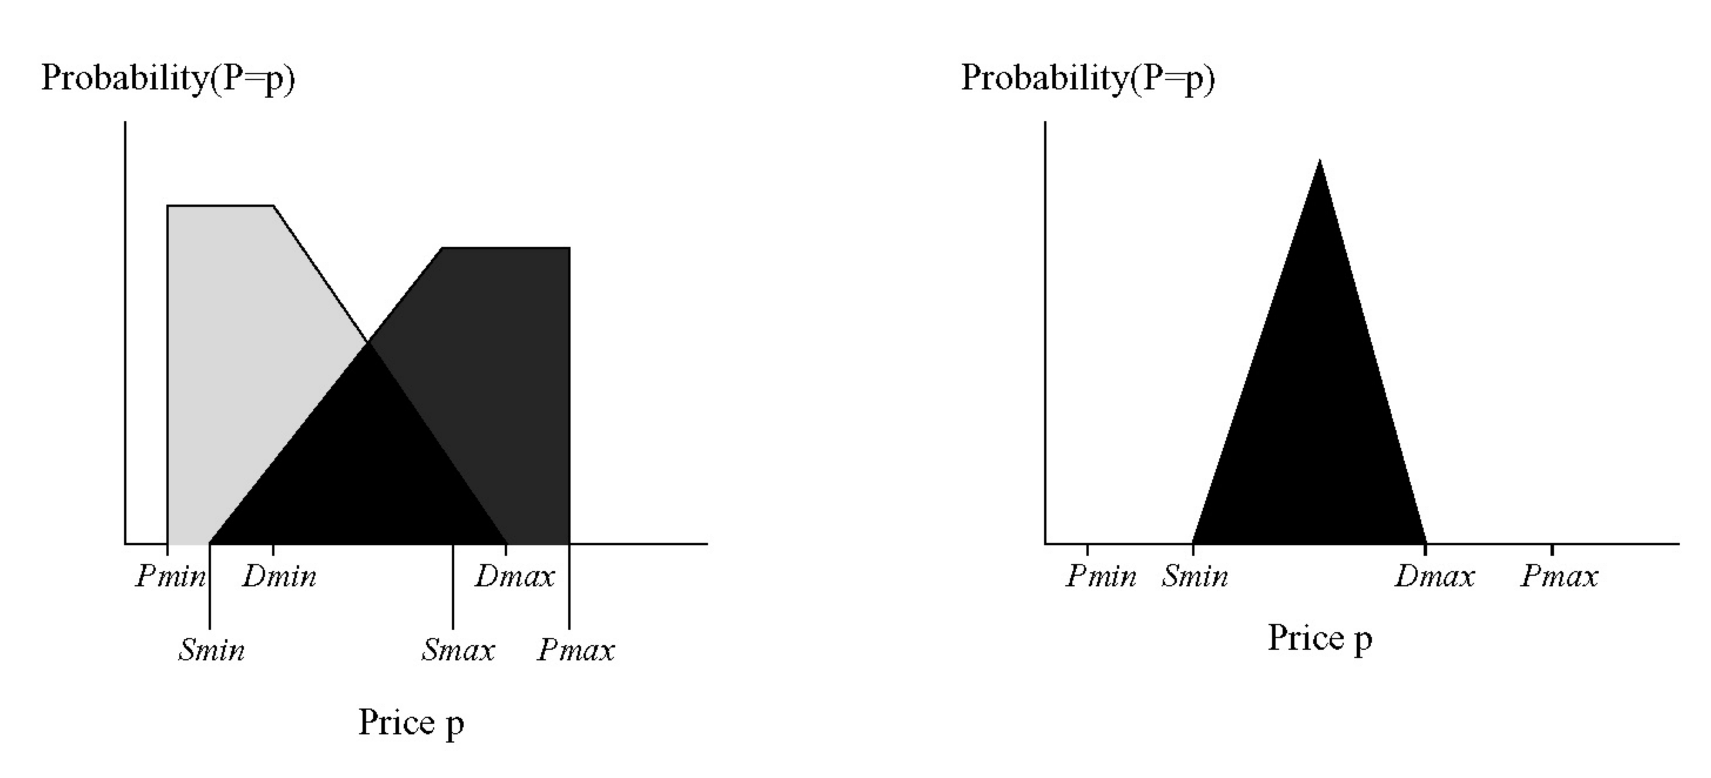

Zero intelligence: not much good

- But in situations where the transaction-price pdf is not symmetric about the peak determined by the equilibrium price, the expected transaction price E(p) will differ significantly (and predictably) from the equilibrium price P_0
- These analytically predictable failures to equilibrate have been demonstrated to occur in replications of Gode&Sunder's experiments, using supply and demand schedules such as this (similar to one of Smith's shown earlier), where $E(p) \neq P_0$ so we need a better trader than ZI-C

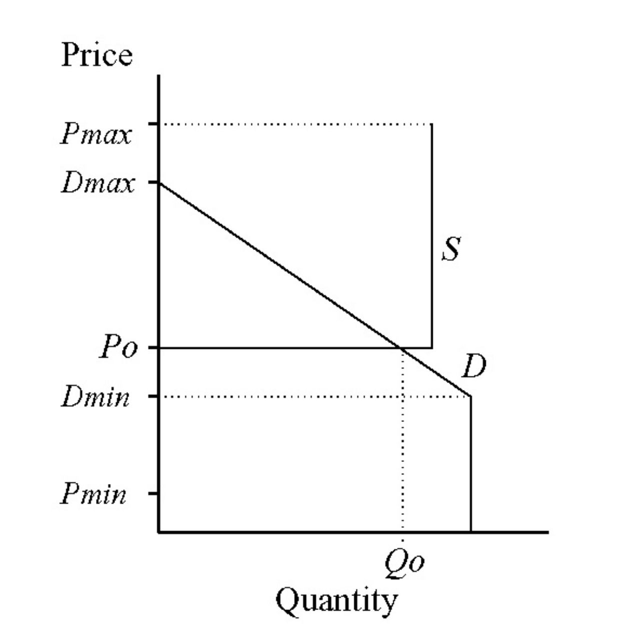

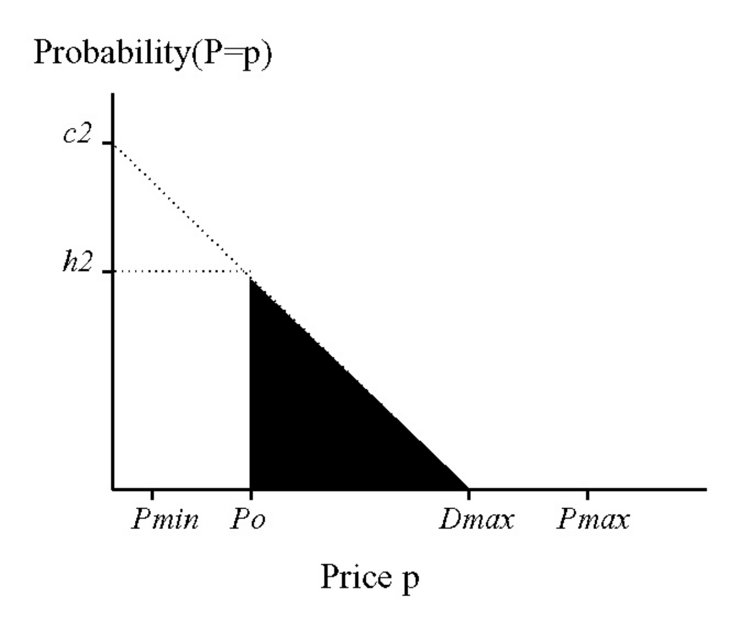

ZIP: qualitative margin heuristics

- Core of the ZIP algorithm is a minimal set of qualitative heuristics for adjusting trader *i*'s margin μ depending on *i*'s current quote-price *p*; on whether *i* is still active; and on the last quote *q*:

For SELLERS:

- if (*q.accepted* == true) then
  - forall sellers s_i
    - if $s_i.p \leq q.p$ then raise $s_i.\mu$
  - if (*q.type* == "bid") then
    - forall active sellers s_i
      - if $s_i.p \geq q.p$
      - then lower $s_i.\mu$
- else
  - if (*q.type* == offer) then
    - forall active sellers s_i
      - if $s_i.p_i \geq q.p$
      - then lower $s_i.\mu$

For BUYERS:

- if (*q.accepted* == true) then
  - forall buyers b_i
    - if $b_i.p \geq q.p$ then raise $b_i.\mu$
  - if (*q.type* =="offer") then
    - forall active buyers b_i
      - if $b_i.p \leq q.p$
      - then lower $b_i.\mu$
- else
  - if (*q.type* == bid) then
    - forall active sellers b_i
      - if $b_i.p \leq q.p$
      - then lower $b_i.\mu$

ZIP: quantitative margin adjustments

- ZIP algorithm is adaptive: adjusts margins up or down using simple machine learning rules
- Quote-price $p_i(t)$ set by limit price λ_i, and margin $\mu_i(t)$: $p_i(t) = \lambda_i \cdot (1+\mu_i(t))$
  - Seller A $\mu_i(t)$ in $[0,\infty]$ forall t; $\mu_i(t)$+= raises margin; $\mu_i(t)$-= lowers margin
  - Buyer $\mu_i(t)$ in $[-1,0]$ forall t; $\mu_i(t)$-= raises margin; $\mu_i(t)$+= lowers margin
- ZIP uses Widrow-Hoff learning rule to adjust actual output A wrt desired D using rate β:
  - $A(t+1) = A(t) + \Delta(t)$; where $\Delta(t) = \beta \cdot (D(t) - A(t))$
  - With momentum (damping) factor γ in $[0,1]$: $A(t+1) = \gamma \cdot A(t) + ((1-\gamma) \cdot \Delta_i(t))$; $0 \leq \gamma_i \leq 1$
- So for ZIP we have:
  - $\mu_i(t+1) = (p_i(t) + \Delta_i(t)) / \lambda_i - 1$
  - $\Delta_i(t) = \beta_i \cdot (\tau_i(t) - p_i(t))$; where target price $\tau_i(t) = (\mathbf{A}_i(t) + \mathbf{R}_i(t) \cdot q(t))$; $\mathbf{A}()$ & $\mathbf{R}()$ stochastic
- Giving:
  - $\mu_i(t+1) = (p_i(t) + \Gamma_i(t)) / \lambda_i - 1$; where $\Gamma_i(0)=0$ and $\Gamma_i(t+1) = \gamma_i \cdot \Gamma_i(t) + ((1-\gamma_i) \cdot \Delta_i(t))$

ZIP: better than humans, actually

(four years later, at IBM T.J. Watson Research Labs in 2001…)

- Tested trader-bot algorithms: ZI-C, ZIP, Kaplan's "Sniper", & IBM's own "MGD"
- Pitted human traders against trading agents in experimental economics lab
- Both ZIP and MGD beat humans
- HP's ZIP did at least as well as IBM's MGD traders (and maybe better?)
  - Average efficiencies: ZIPs=1.030; MGDs=1.023; Humans=0.876

"…the successful demonstration of machine superiority in the CDA and other common auctions could have a much more direct and powerful impact – one that might be measured in billions of dollars annually."

- R Das, J E Hanson, J O Kephart, & G Tesauro. Agent-human interactions in the continuous double auction. *Proceedings IJCAI-01*, Seattle, 2001

  http://www.research.ibm.com/infoecon/researchpapers.html

Todd Kaplan's "Lurking Sniper" Trader

- A surprisingly robust and effective trader algorithm entered into an early trading-agent contest, which outperformed all the competition including more complex trading algorithms that used explicit optimizing principles, statistically based predictions of future transaction prices, or learning algorithms.
- Simple too:
  - Sit quietly in the background doing nothing until the bid-offer spread drops to a sufficiently small value, or the offer is less than the smallest transaction price in the previous period, or there is not much time until the market closes. If any of these conditions is met, sniper jumps in and "steals the deal" so long as the deal makes the sniper a profit greater than its minimum threshold.
- Perhaps too simple?
  - Kaplan Snipers don't adapt to market activity
  - They are unable to infer the market's P_0; so they will snipe any deal, however far from equilibrium
  - They free-ride on the goodwill of other traders; and so are not much good when confronted with copies of themselves
- J. Rust, J. H. Miller, R. Palmer (1993). "Behavior of Trading Automata in a Computerized Double Auction Market" in D. Friedman & J. Rust (eds) *The Double Auction Market: Institutions, Theories, and Evidence*. pp.155—198. Addison Wesley.

Gjerstad-Dickhaut GD traders (and MGD)

- A relatively sophisticated algorithm that computes a belief function using data for recent market activity: calculates the belief function from history of *n* recent trades
- The belief function indicates the probability, for each possible bid or offer price, that a bid or offer would be accepted at that price.
- Uses cubic spline interpolation to compute values of the belief function for prices that do not occur in the history list
- Chooses a quote-price that maximizes the trader's expected gain: simple product of utility gain from trade at that price and probability of acceptance at that price
- Subsequently modified so that belief function is forced to show a zero probability of acceptance for bids lower than the previous lowest trade price and for offers higher than the previous highest trade price – this reduces volatility
- Modified GD known as MGD.
- S. Gjerstad and J Dickhaut (1998). Price formation in double auctions. *Games and Economic Behavior*, 22:1—29.

How to decide which is best?

- With multiple alternative trader algorithms, deciding which is best becomes an issue
- Attempting to form a full analytic (e.g. game-theoretical) understanding of the capabilities of algorithms such as ZI-C, KSniper, ZIP, or MGD is typically either:
  - impossible, or…
  - so difficult/laborious that it may as well be impossible, or…
  - requires so many simplifying assumptions that the end conclusions are of limited or zero relevance to the real system
- So empirical studies (simulation experiments) are the method of choice…
- Side-by-side comparisons can illustrate differences in response/performance between markets homogeneously populated by specific trader algorithms, varying the trader algorithm but keeping constant/repeatable the market supply & demand schedules and/or the dynamic changes in supply and demand
- For MBC applications, the assumption that all traders will run the same algorithm is plausible, e.g. resource allocation on a company's intranet grid or utility data center
- For many other applications (e.g. financial markets) homogeneity is implausible

Agent-vs-Agent Trading Contests: Santa Fe

- Seminal experiment in heterogeneous trader-agent market dynamics:
- J. Rust, J. H. Miller, R. Palmer (1993). "Behavior of Trading Automata in a Computerized Double Auction Market" in D. Friedman & J. Rust (eds) *The Double Auction Market: Institutions, Theories, and Evidence*. pp.155—198. Addison Wesley.
- A series of tournaments held at the Santa Fe Institute: participants from around the world submitted over 30 different algorithms to be pitted against each other in a series of computerized double-auction tournaments, with $10,000 in prizes paid out in proportion to profits earned by the trader agents.
- "We find that the top-ranked programs yield a fairly `realistic' working model of a [C]DA market in the sense that their collective behavior is consistent with the key `stylized facts' of human experiments. We also find that a very simple strategy is a highly effective and robust performer in these markets. This strategy was able to outperform more complex algorithms that use statistically based predictions of future transaction prices, explicit optimizing principles, or sophisticated `learning algorithms'." (p.157)
- …that very simple strategy was Kaplan's Sniper

Agent-vs-Agent Trading Contests: TAC

- Trading Agent Contest (TAC)
- Initiated at Wellman's Lab in Michigan, 2001; now run from the Swedish Institute for Computer Sci.
- Now two parallel contests:
  - TAC Classic - a "travel agent" scenario based on complex procurement on multiple simultaneous auctions
  - TAC SCM - a PC manufacturer supply-chain management scenario based on sourcing of components, manufacturing of PC's and sales to customers
- www.sics.se/tac

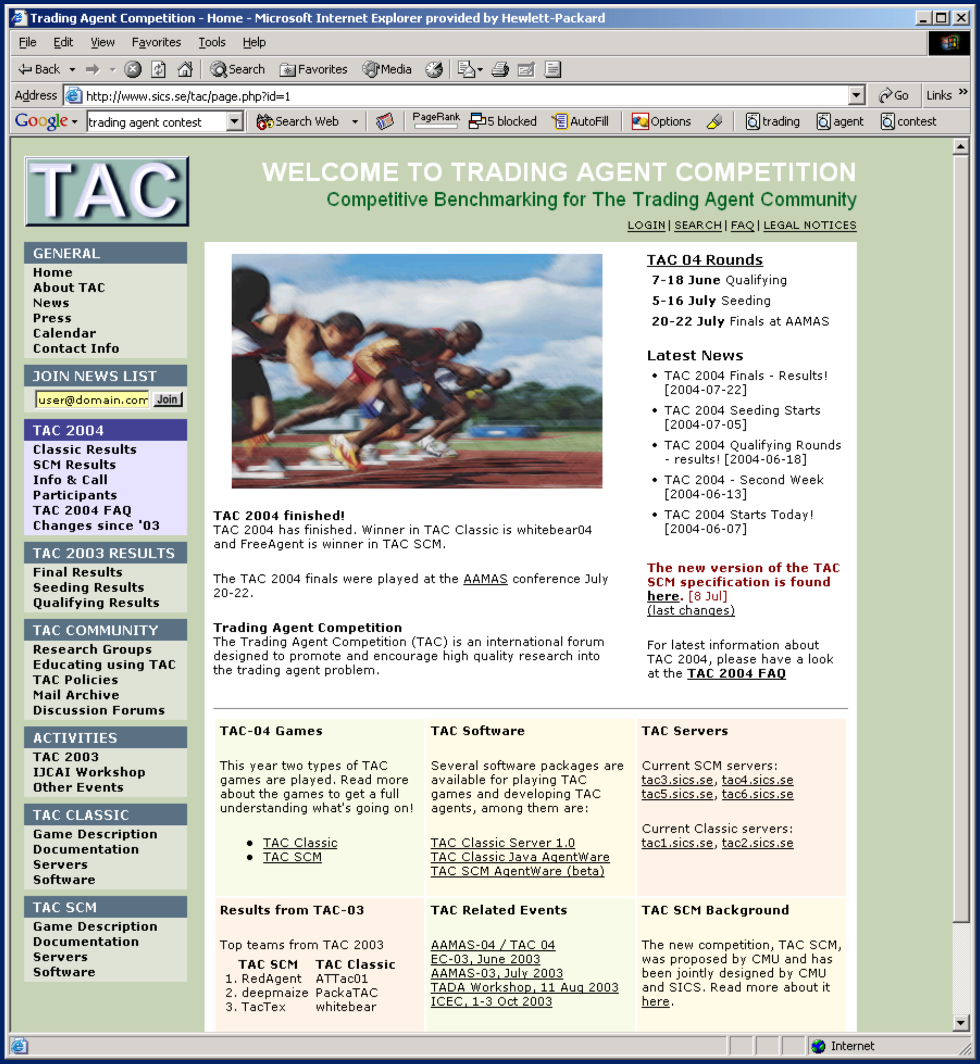

Agent-vs-Agent Trading Contests: IBM #1

G. Tesauro & R Das (2001). High-Performance Bidding Agents for the Continuous Double Auction. *IJCAI 2001 Economic Agents Models and Mechanisms Workshop* and/or *Proc Third ACM Conference on Electronic Commerce* pp.206-209.

http://www.research.ibm.com/infoecon/paps/dblauc.pdf

Tested populations of trading-agents using pairs of algorithms chosen from ZI (-C), Kaplan, ZIP, GD, and MGD, in real-time market experiments where traders had to trade multiple units with different limit prices…

…suggested modifications to ZIP and to GD

…MGD is the Modified GD; ZIP modifications so minor that they still call it ZIP

Performed homogeneous population tests for validation/reference

Performed balanced-group tests, in which:

"…buyers and sellers are evenly split between two types of [trader] agent [algorithm], and every agent of one type has a counterpart of the other type with identical limit prices. …we believe [this] test to be the fairest way to test two different algorithms against each other."

Agent-vs-Agent Trading Contests: IBM #1

Results of homogenous population tests...

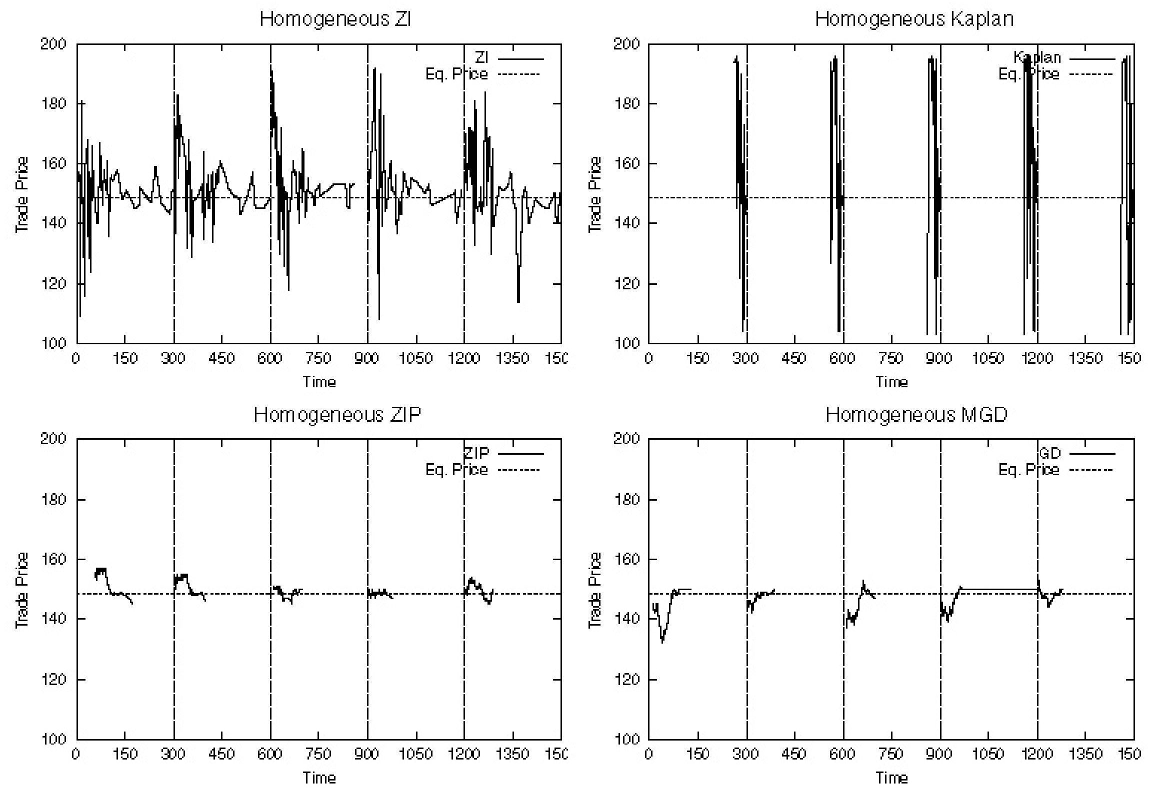

Figure 1: Trade price vs. time for different strategies (homogeneous population): (a) ZI, (b) Kaplan, (c) ZIP, and (d) MGD. The vertical dashed-lines indicate the start of a new trading period, while the horizontal dashed-lines represent the market equilibrium price.

Agent-vs-Agent Trading Contests: IBM #1

Results of balanced group tests...

Groups	Wins	Surplus Difference	Efficiency	Trade Ratio	Price (std dev)
Kaplan vs ZI	0-100	-1080.0	0.964	1.055	150.6 (±17.3)
ZIP vs. ZI	100-0	+1037.0	0.955	1.003	151.7 (±8.3)
ZIP vs. Kaplan	99-1	+347.9	0.982	1.128	148.8 (±4.6)
GD vs. ZI	99-1	+575.1	0.978	1.049	150.7 (±10.0)
GD vs. Kaplan	93-7	+323.0	0.978	1.141	149.3 (±10.1)
GD vs. ZIP	36-63	-54.4	0.996	1.062	149.2 (±6.3)
MGD vs. ZI	99-1	+787.4	0.960	0.972	150.9 (±10.1)
MGD vs. Kaplan	98-2	+210.3	0.993	1.086	148.2 (±4.8)
MGD vs. ZIP	71-29	+44.9	0.997	1.054	148.6 (±2.9)

Table 3: Balanced group test results. Columns show: the number of wins for each group in 100 trials, the average difference between first group and second group surplus, population average efficiency, trade ratio, and final period trade price ± (standard deviation). The magnitude of surplus differences can be compared to the total population surplus of 2612.0 theoretically available in each experiment.

Agent-vs-Agent Trading Contests: IBM #2

W.E. Walsh, R. Das, G. Tesauro, & J O Kephart (2002). Analyzing Complex Strategic Interactions in Multi-Agent Systems" *Proc. Game Theoretic and Decision Theoretic Agents Workshop* , AAAI-02, Edmonton, Canada, 2002.
http://www.research.ibm.com/infoecon/paps/MultiAgentGame.pdf

- IBM explored the "population dynamics" of CDA markets with various mixtures of ZIP, GD, and Kaplan traders, and characterised those dynamics for all mixture ratios when running replicator dynamics experiments:
  - where each trader may occasionally and asynchronously examines the strategy of a randomly-chosen competitor
  - If that other strategy appears to be generating more payoff than the trader's current one, then the trader switches to that strategy
- Triangular "simplex" figures summarise the dynamics for these experiments
  - Each triangle encloses a space within which each point represents a particular mix of ZIP, GD, and Kaplan: the vertices represent all traders playing the same algorithm.
  - Curved flow vectors indicate the transition paths the system takes as it approaches the various equilibrium points. Gray-scale shading represents average speed of transition along the flow vectors on the space: darker is faster.

Agent-vs-Agent Trading Contests: IBM #2

W.E. Walsh, R. Das, G. Tesauro, & J O Kephart (2002). Analyzing Complex Strategic Interactions in Multi-Agent Systems" http://www.research.ibm.com/infoecon/paps/MultiAgentGame.pdf

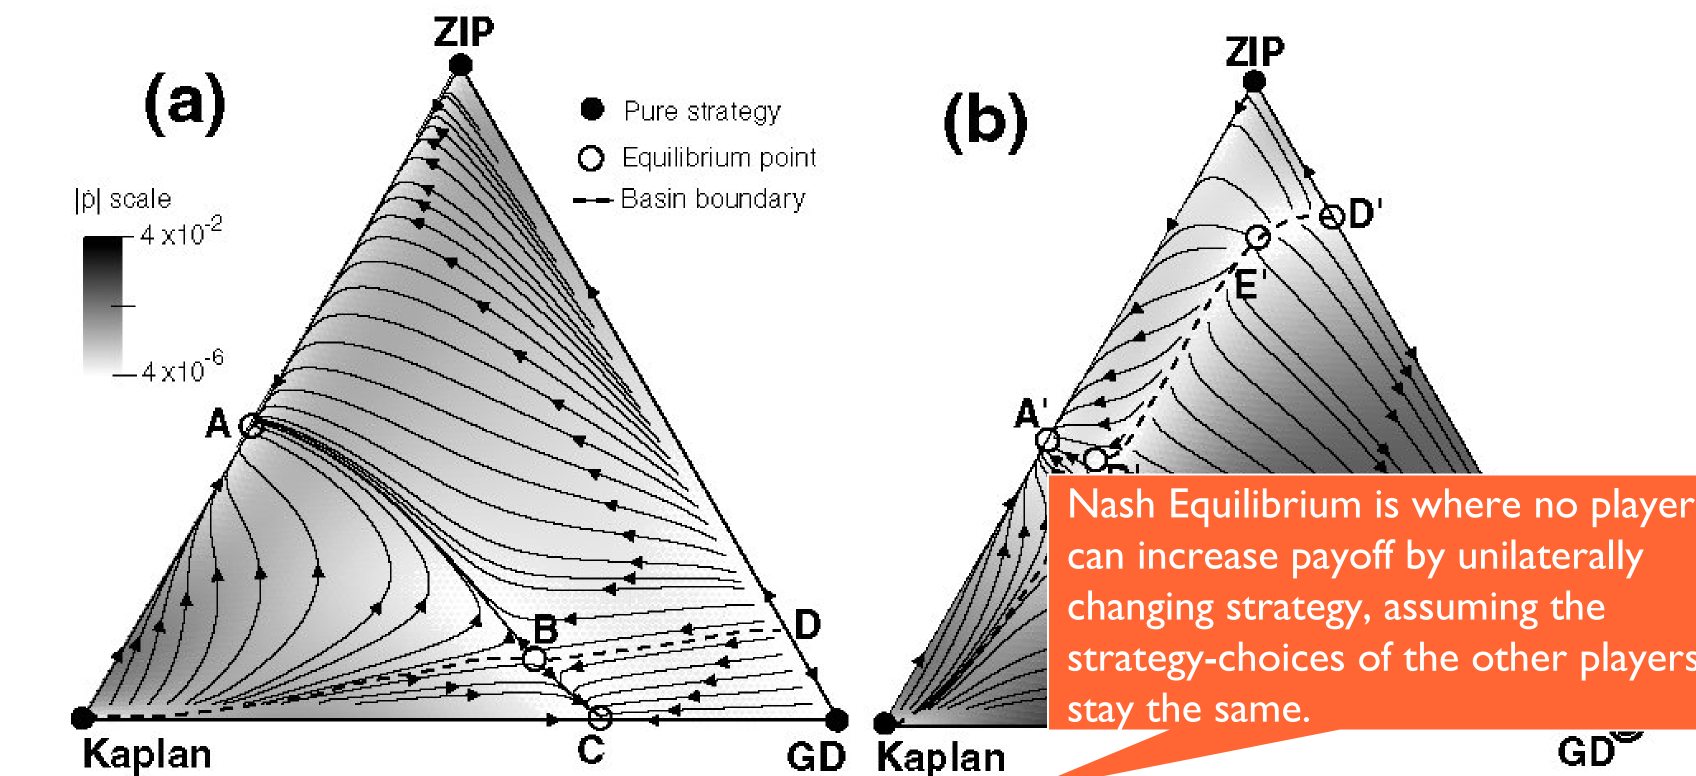

Nash Equilibrium is where no player can increase payoff by unilaterally changing strategy, assuming the strategy-choices of the other players stay the same.

Figure 2: (a) Replicator dynamics for the CDA Game. Open circles are Nash equilibria with labels corresponding to those in Table 2. Other notations are similar to those in Figure 1. (b) Replicator dynamics for the CDA Game with perturbed payoffs, in which 5% of the ZIP and Kaplan agent payoffs was shifted to GD agents.

Agent-vs-Human Trading Contests: IBM #3

- Experiments at IBM T.J. Watson Research Labs in 2001
- Pitted human traders against trading agents in experimental economics lab
- Astonishingly, a first: no-one had thought of doing this before
- Tested trading-agent algorithms including Modified versions of ZIP and GD
- Both ZIP and MGD further modified to work with an order book
  - Book shows *n* current bids & offers, ordered best to worst.
  - e.g for *n*=3: BidBook(1.08,1.07,1.02) OfferBook(1.11,1.14,1.21)
- MGD is stll the Modified GD; ZIP modification is so trivial that the name is still ZIP.
- ZIP appears to do better than MGD (but not a lot of data)
- Finding #1: consistent off-equilibrium trading

- R Das, J E Hanson, J O Kephart, & G Tesauro. Agent-human interactions in the continuous double auction. *Proceedings IJCAI-01*, Seattle, 2001. http://www.research.ibm.com/infoecon/paps/AgentHuman.pdf

Agent-vs-Human Contests: IBM #3 Finding 1

- Finding #1: consistent off-equilibrium trading (horizontal dashed line is P_0)

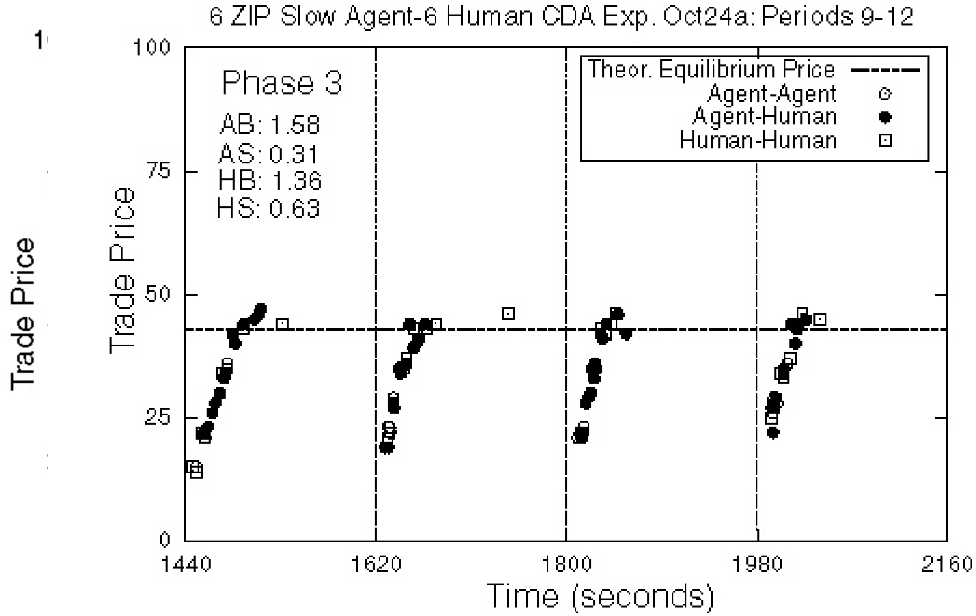

Figure 3: Trade price vs time for experiment Oct24a with 6 ZIP slow agents and 6 humans. Trading activity in periods 9-12 (out of 16 total) is shown. Other details are similar to those in Figure 2.

Figure 2: Trade pri
the start of a new tra
each phase, p^* is sh
humans with open s
average efficiency o

es indicate
prices. In
tween two
efer to the

GD Results
ZIP Results

Agent-vs-Human Contests: IBM #3 Finding 2

- Finding #2: ZIP and MGD both beat the humans

Experiment				Agent			Human		
ID	# Periods	# Trades	Interaction	Strategy	Surplus	Efficiency	# Traders	Surplus	Efficiency
Oct17	15	412	0.38	GD Fast	11058	1.016	5	6991	0.927
Oct18	15	504	0.29	ZIP Fast	11069	1.028	6	7023	0.652
Oct23	16	320	0.33	GD Fast	10495	0.999	3	4582	0.965
Oct24a	16	455	0.48	ZIP Slow	10696	1.032	6	9490	0.916
Oct24b	9	261	0.42	GD Fast	6808	1.026	6	6353	0.958
Oct25	16	433	0.49	GD Fast	12159	1.052	6	9708	0.840

Table 1: Summary of the six agent-human CDA experiments. For each experiment, the table presents: the number of trading periods, the total number of successful trades, the fraction of trades between agents and humans, the bidding strategy employed by all six agents, the number of human traders, and the aggregate agent and human performances in terms of total surplus accumulated over the entire experiment and the average efficiency.

- In all experiments, humans are beaten by traders; MGD scores best, but also scores worst
- MGD avg(S)=10130 avg(E)=1.023;   ZIP avg(S)=10883 avg(E)=1.030 ... so ZIP wins overall?

"…the successful demonstration of machine superiority in the CDA and other common auctions could have a much more direct and powerful impact – one that might be measured in billions of dollars annually."

- …at point of execution in the financial markets, humans might no longer make economic sense

In the news…

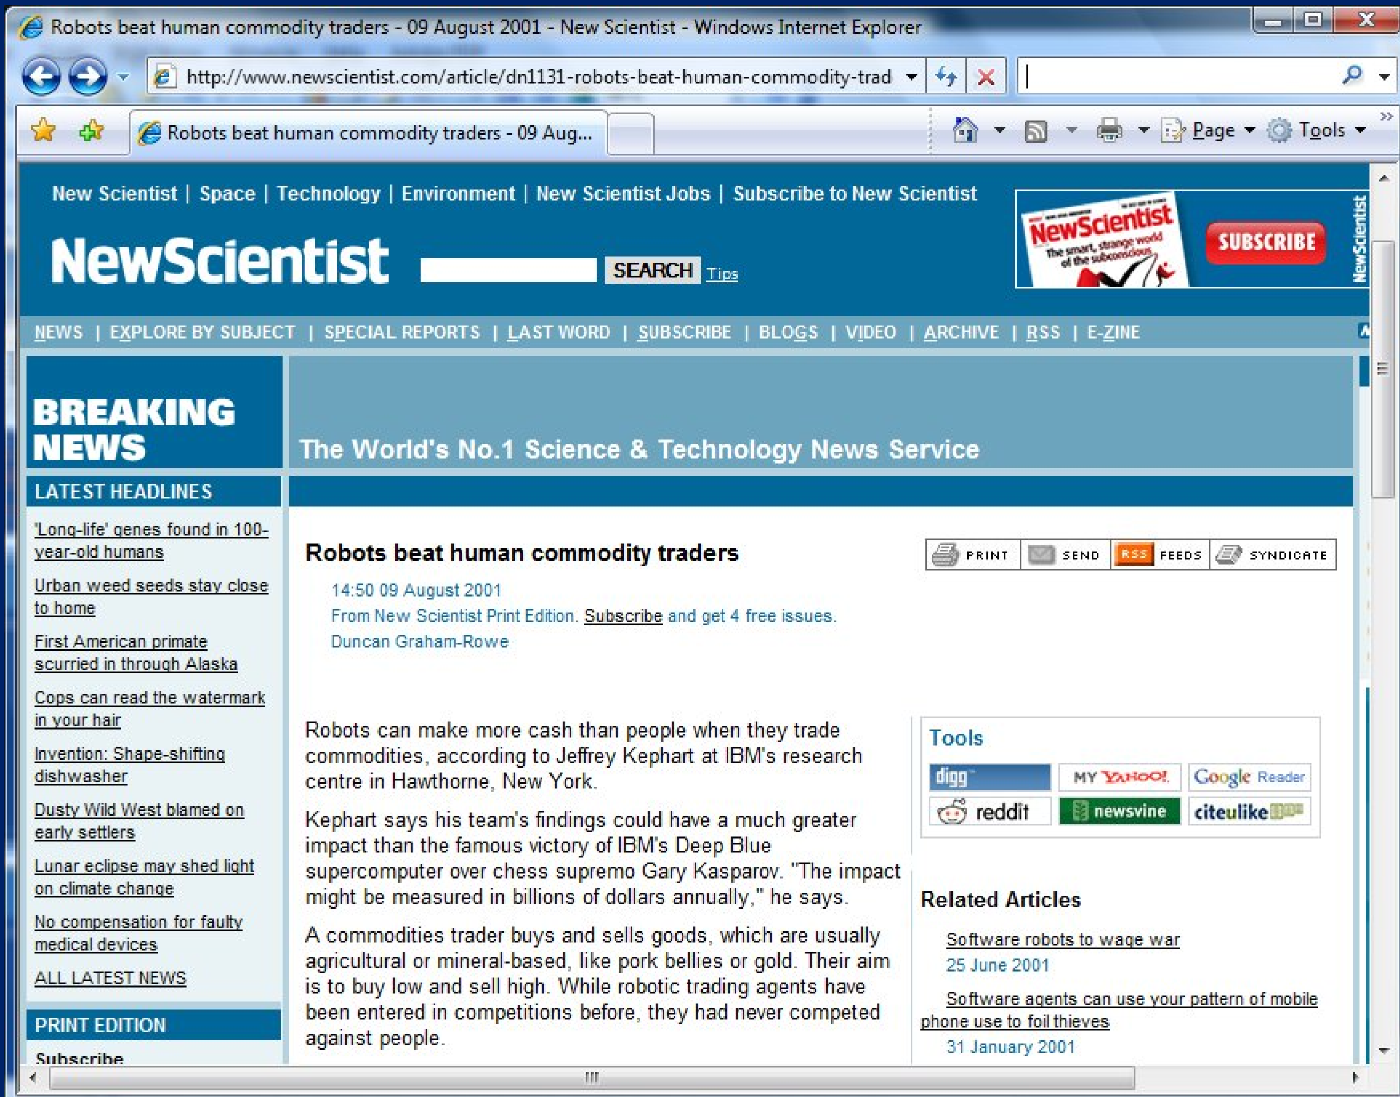

Robots beat human commodity traders - 09 August 2001 - New Scientist - Windows Internet Explorer

http://www.newscientist.com/article/dn1131-robots-beat-human-commodity-trad

New Scientist | Space | Technology | Environment | New Scientist Jobs | Subscribe to New Scientist

NewScientist SEARCH Tips

NEWS | EXPLORE BY SUBJECT | SPECIAL REPORTS | LAST WORD | SUBSCRIBE | BLOGS | VIDEO | ARCHIVE | RSS | E-ZINE

BREAKING NEWS

The World's No.1 Science & Technology News Service

LATEST HEADLINES

- 'Long-life' genes found in 100-year-old humans
- Urban weed seeds stay close to home
- First American primate scurried in through Alaska
- Cops can read the watermark in your hair
- Invention: Shape-shifting dishwasher
- Dusty Wild West blamed on early settlers
- Lunar eclipse may shed light on climate change
- No compensation for faulty medical devices
- ALL LATEST NEWS

PRINT EDITION

Subscribe

Robots beat human commodity traders

14:50 09 August 2001

From New Scientist Print Edition. Subscribe and get 4 free issues.

Duncan Graham-Rowe

Robots can make more cash than people when they trade commodities, according to Jeffrey Kephart at IBM's research centre in Hawthorne, New York.

Kephart says his team's findings could have a much greater impact than the famous victory of IBM's Deep Blue supercomputer over chess supremo Gary Kasparov. "The impact might be measured in billions of dollars annually," he says.

A commodities trader buys and sells goods, which are usually agricultural or mineral-based, like pork bellies or gold. Their aim is to buy low and sell high. While robotic trading agents have been entered in competitions before, they had never competed against people.

Tools

Related Articles

- Software robots to wage war
  25 June 2001
- Software agents can use your pattern of mobile phone use to foil thieves
  31 January 2001

And on and on and on it goes

- 2002: After the human-robot trials, IBM publish "GDX" algorithm
  - GDX dominates ZIP
  - IBM claim it "…may offer best performance of any published trading algorithm"
  - 2007: Iain Toft's UEA PhD thesis seems to show ZIP dominates GDX??
- 2005: "ZIP60"
  - Replace *Ca* & *Cr* constants with ***U****[C_low,C_hi]* ranges… 10 params… "ZIP10"
  - Have one vector of 10 params for Buyers, another for Sellers… "ZIP20"
  - One for each of the three cases in the "decision tree"… "ZIP60"
  - Original ZIP now re-named "ZIP08"
- 2006: P. Vytelingum's Southampton PhD "Adaptive Aggression" (AA) algorithm
  - AA dominates GDX
- 2008: Cliff "ZIP100" (ZIP60 in a party frock)
  - ZIP100 dominates AA (??)
  - Paper coming later in 2008.
- NB: "vanilla" ZI models, when done well, can still be very informative
  - E.g. J. Doyne Farmer's work on ZI modelling dynamics of LSE time-series

Using this in computing systems

- Traditional approach to resource allocation in IT:
  - Components/processes/agents submit their resource requirements to a central planner (eg, the OS)
  - Central planner computes optimal allocation

- Problems:
  - Not robust to central-planner failures
  - Not robust to inaccurate resource-requirement signals
    - Perhaps to due to faults/bugs, or even
    - Perhaps due to malign intention

- How to ensure resource-requirement signals are accurate?
- Answer: introduce payments
  - Payments can be thought of as a handicap to ensure honest signalling
  - In some cases, payments can be in a virtual currency

..And to more traditional control problems

- Xerox Parc built a multi-site air-conditioning system using market-based principles
- Scarce resource - conditioned air - must be allocated to consumers - aircon controls in individual offices
- Market-based system ensured balance of supply and demand across different sites- excess supply at one site was purchased cheaply by another site, ensuring matched supply and demand
- Relevant to data centres:

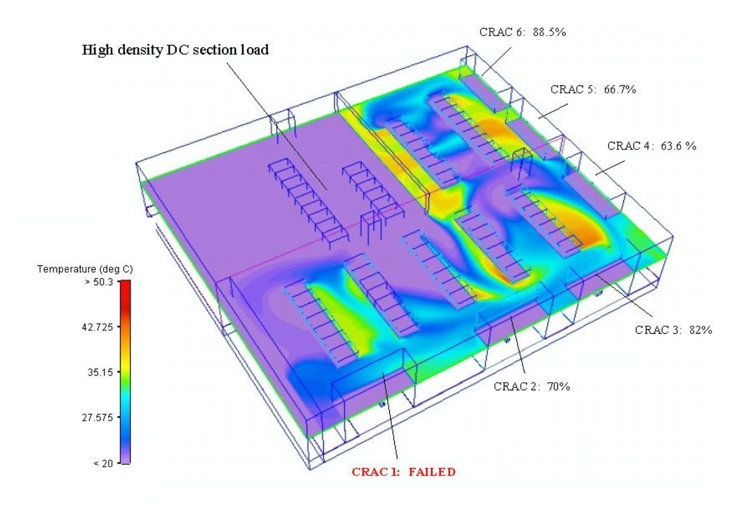

FIG. 5. Temperature Plane Plot after CRAC unit failure and load redistributed.

Market Based Control: early work

- The theme of using ideas from microeconomics for computational resource allocation and control was explored in depth by Miller & Drexler in three papers proposing the development of what they referred to as Agoric Systems: distributed computer systems, potentially involving multiple owners or vendors (i.e. federated), where scarce resources are automatically allocated by techniques inspired by the operation of free-market economies. All three papers appear in:
  - B. Huberman (ed): *The Ecology of Computation*. North-Holland, 1988;
- The next major collection on the topic of market-based control (MBC) gathered together papers describing market-based approaches for: network bandwidth allocation (Miller et al); RAM allocation (Harty & Cheriton); air conditioning operation (Clearwater et al) ; pollution regulation (Marron & Bartels) and job-shop scheduling (Baker); all in:
  - S. Clearwater (ed): *Market-Based Control.* World Scientific, 1995.
- Not one of these papers described a system that was both decentralized and automatic
  - Some were automatic but centralized
  - others were decentralized but required humans to make final decisions
- For further discussion see pp.11-13 of D Cliff & J Bruten. *Minimal-Intelligence Agents for Bargaining Behaviors in Market-Based Environments*. Technical Report HPL-97-91. http://www.hpl.hp.com/techreports/97/HPL-97-91.html

Market Based Control: recent work

- Recent work has been more promising, providing genuinely decentralized or distributed approaches:

- J Cheng & M Wellman (1998) The WALRAS algorithm: A convergent distributed implementation of general equilibrium outcomes. *Computational Economics* 12:1-24.
- M Wellman, W Walsh, P Wurman, & J MacKie-Mason (2001). Auction protocols for decentralized scheduling. *Games and Economic Behavior*, 2001. http://ai.eecs.umich.edu/people/wellman/pubs/geb01wwwmm.html
- M Wellman, J MacKie-Mason, D Reeves, & S Swaminathan) (2003). Exploring bidding strategies for market-based scheduling. *Fourth ACM Conference on Electronic Commerce*, 2003. http://ai.eecs.umich.edu/people/wellman/pubs/dexter02.html
- A Byde, M Salle, C Bartolini (2003) *Market-Based Resource Allocation for Utility Data Centers*. http://www.hpl.hp.com/techreports/2003/HPL-2003-188.pdf
- K Lai, B Huberman, & L Fine (2004) *Tycoon: A Distributed Market-Based Resource Allocation System.* http://arxiv.org/PS_cache/cs/pdf/0404/0404013.pdf

Michigan Group

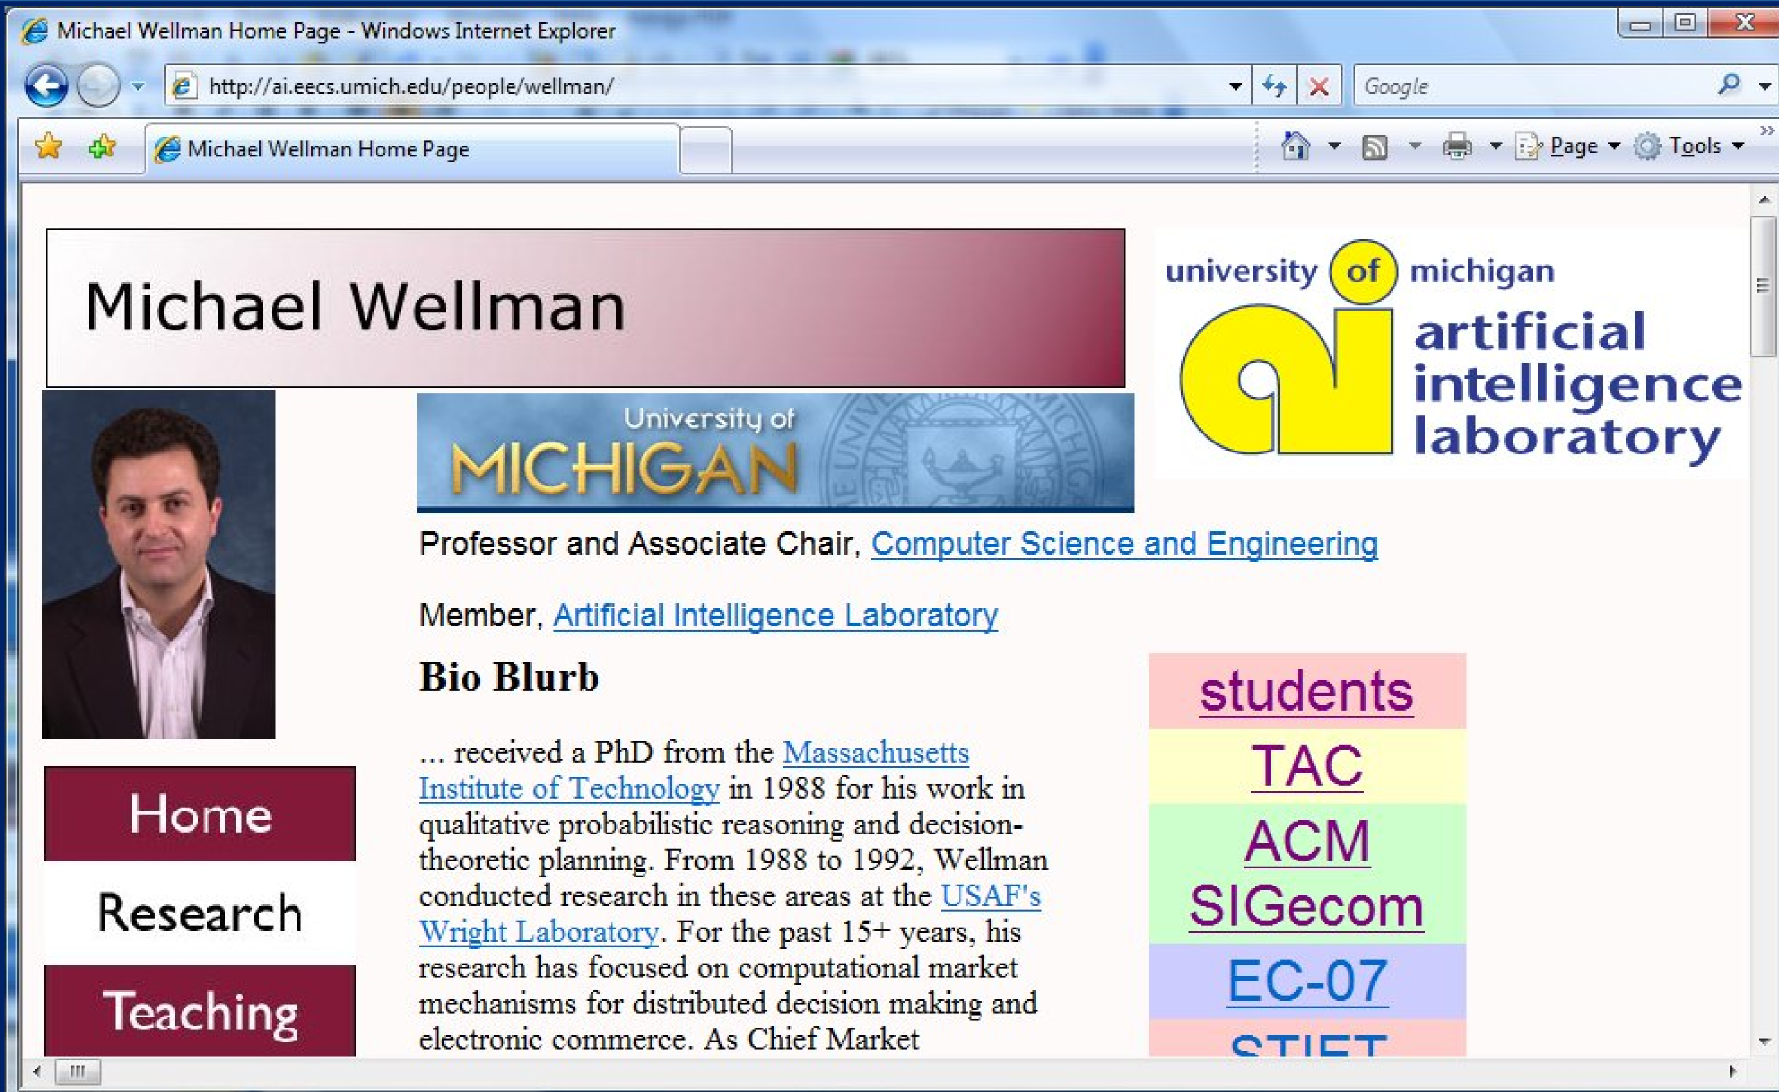

www.sics.se/tac/

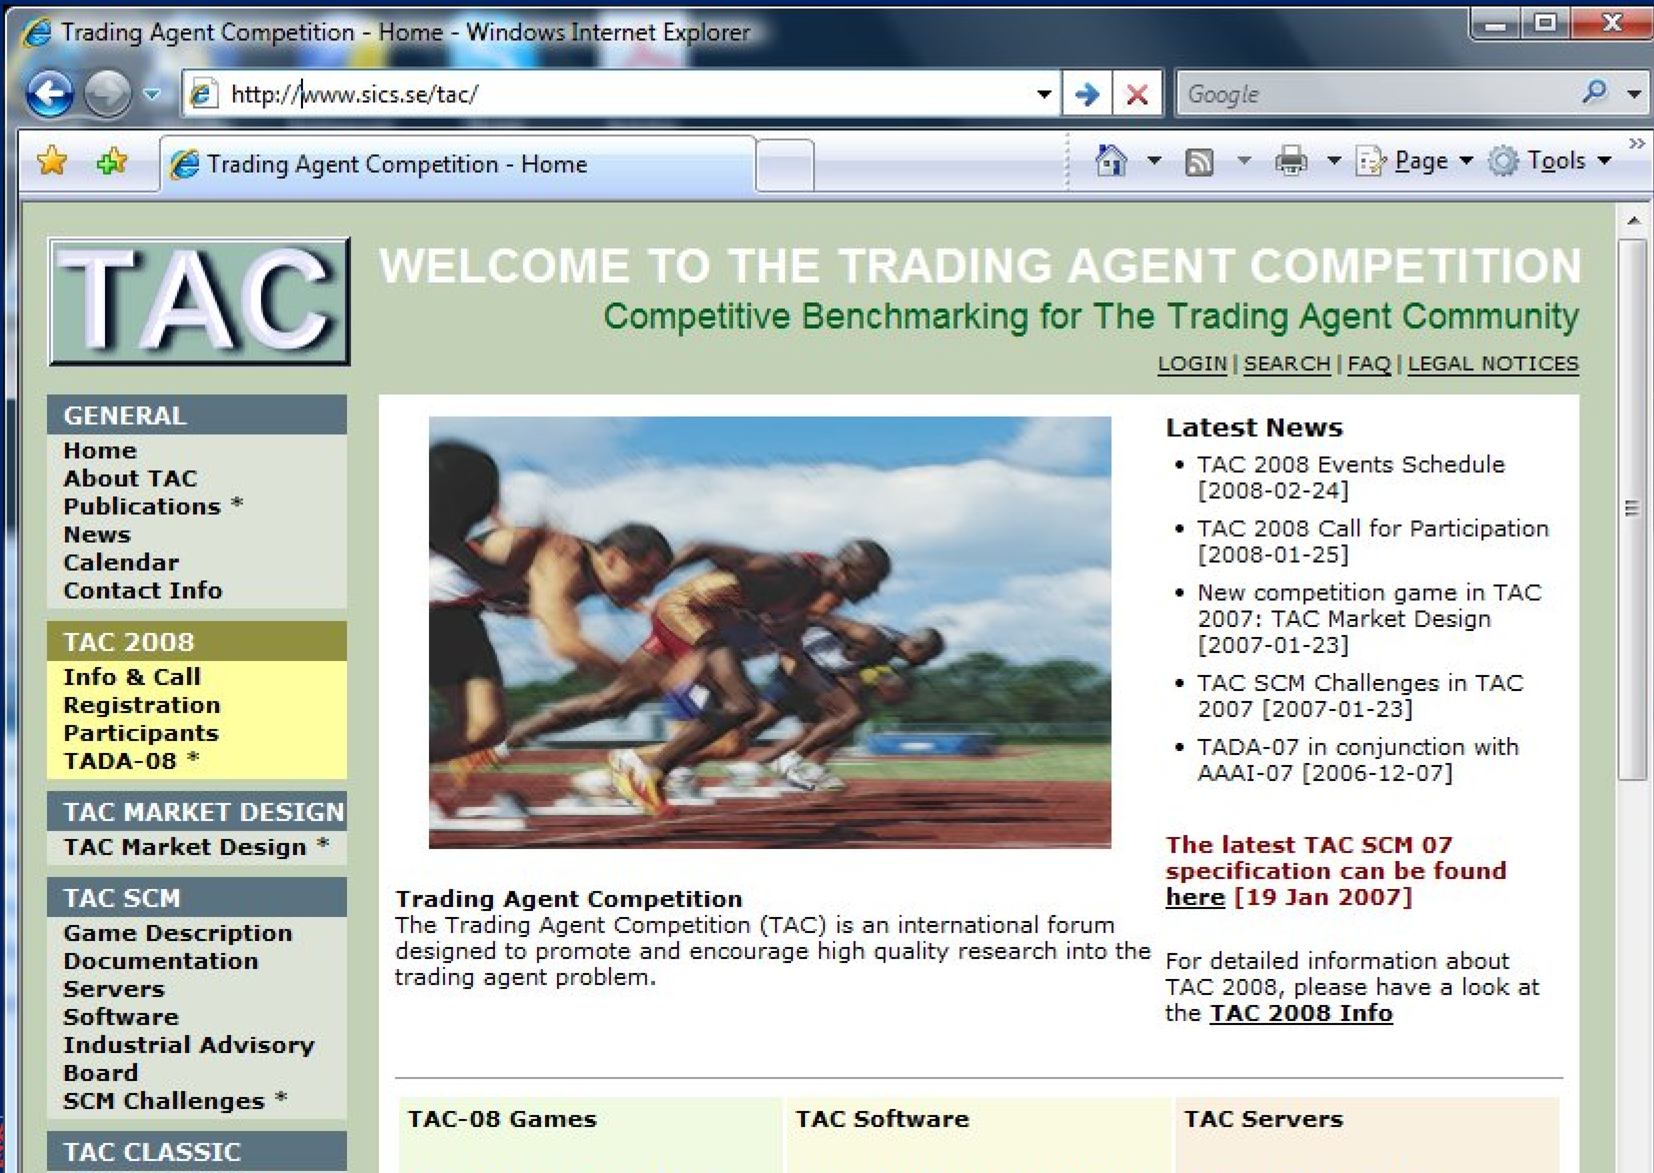

Southampton Group

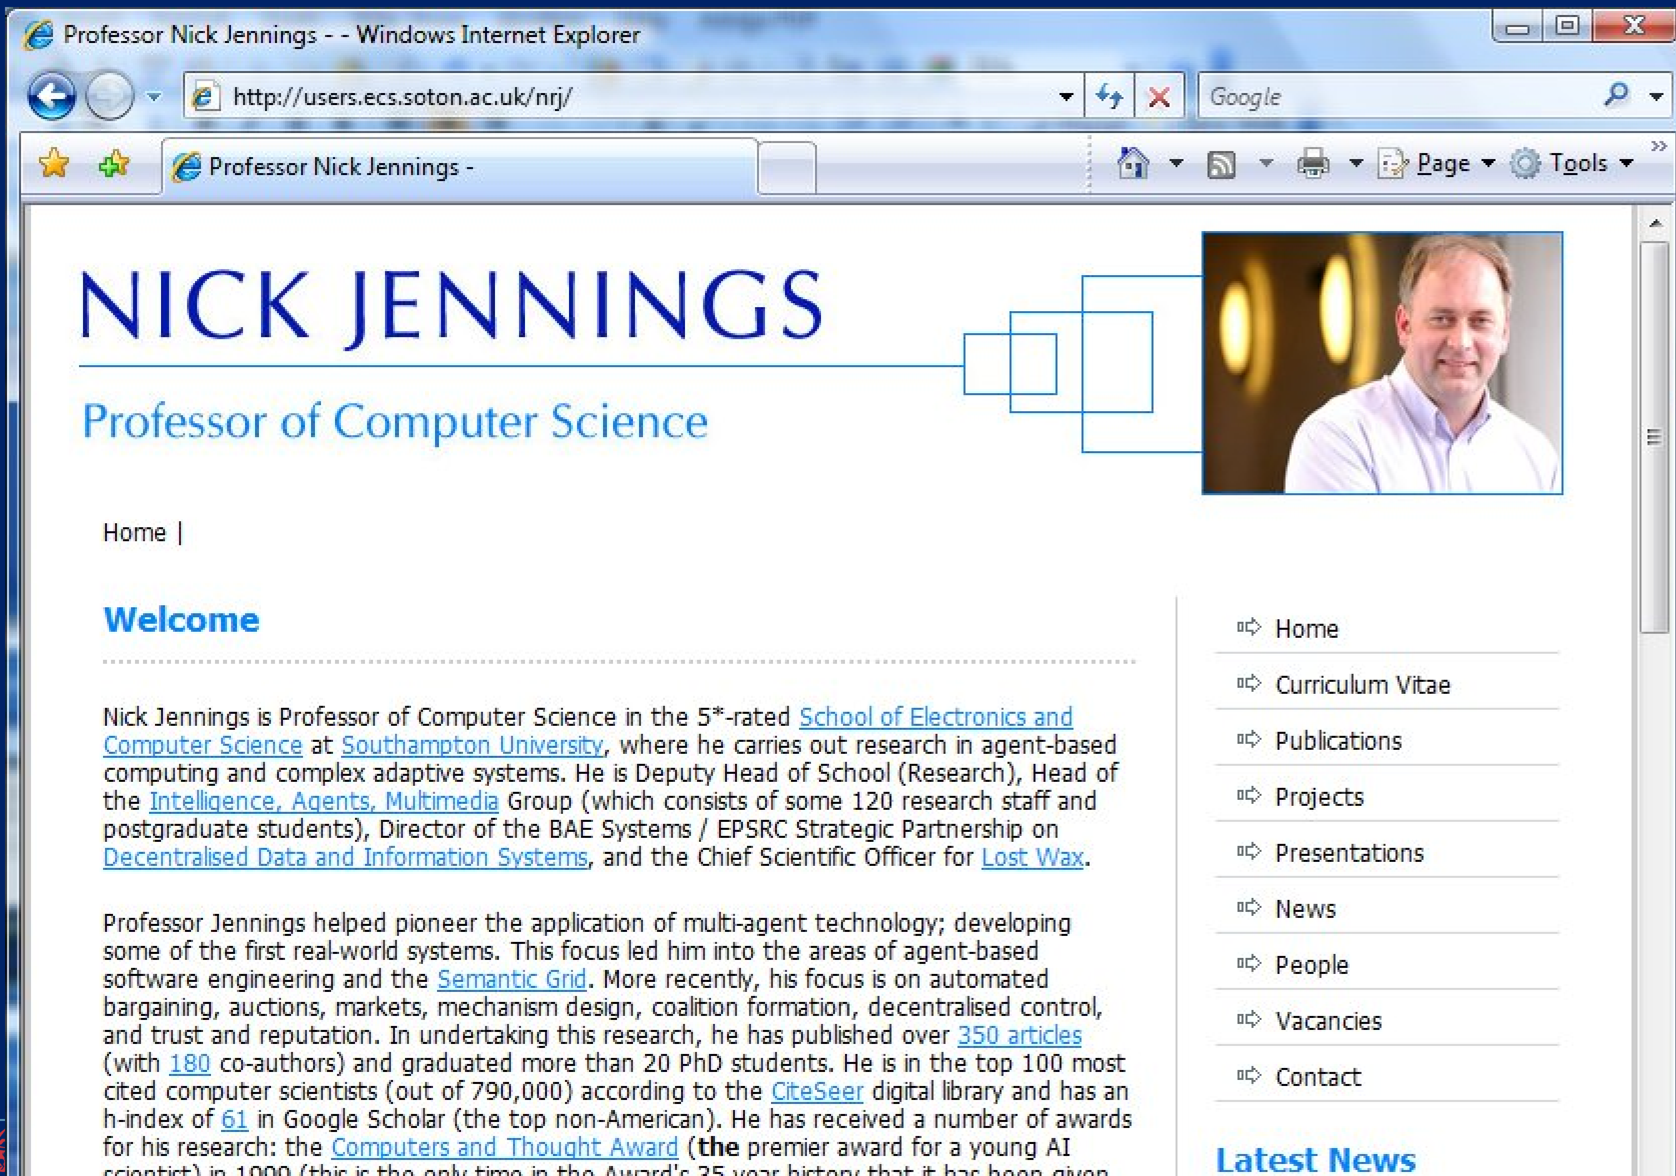

NICK JENNINGS

Professor of Computer Science

Home |

Welcome

Nick Jennings is Professor of Computer Science in the 5*-rated School of Electronics and Computer Science at Southampton University, where he carries out research in agent-based computing and complex adaptive systems. He is Deputy Head of School (Research), Head of the Intelligence, Agents, Multimedia Group (which consists of some 120 research staff and postgraduate students), Director of the BAE Systems / EPSRC Strategic Partnership on Decentralised Data and Information Systems, and the Chief Scientific Officer for Lost Wax.

Professor Jennings helped pioneer the application of multi-agent technology; developing some of the first real-world systems. This focus led him into the areas of agent-based software engineering and the Semantic Grid. More recently, his focus is on automated bargaining, auctions, markets, mechanism design, coalition formation, decentralised control, and trust and reputation. In undertaking this research, he has published over 350 articles (with 180 co-authors) and graduated more than 20 PhD students. He is in the top 100 most cited computer scientists (out of 790,000) according to the CiteSeer digital library and has an h-index of 61 in Google Scholar (the top non-American). He has received a number of awards for his research: the Computers and Thought Award (**the** premier award for a young AI scientist) in 1999 (this is the only time in the Award's 35 year history that it has been given

- Home
- Curriculum Vitae
- Publications
- Projects
- Presentations
- News
- People
- Vacancies
- Contact

Latest News

www.marketbasedcontrol.com

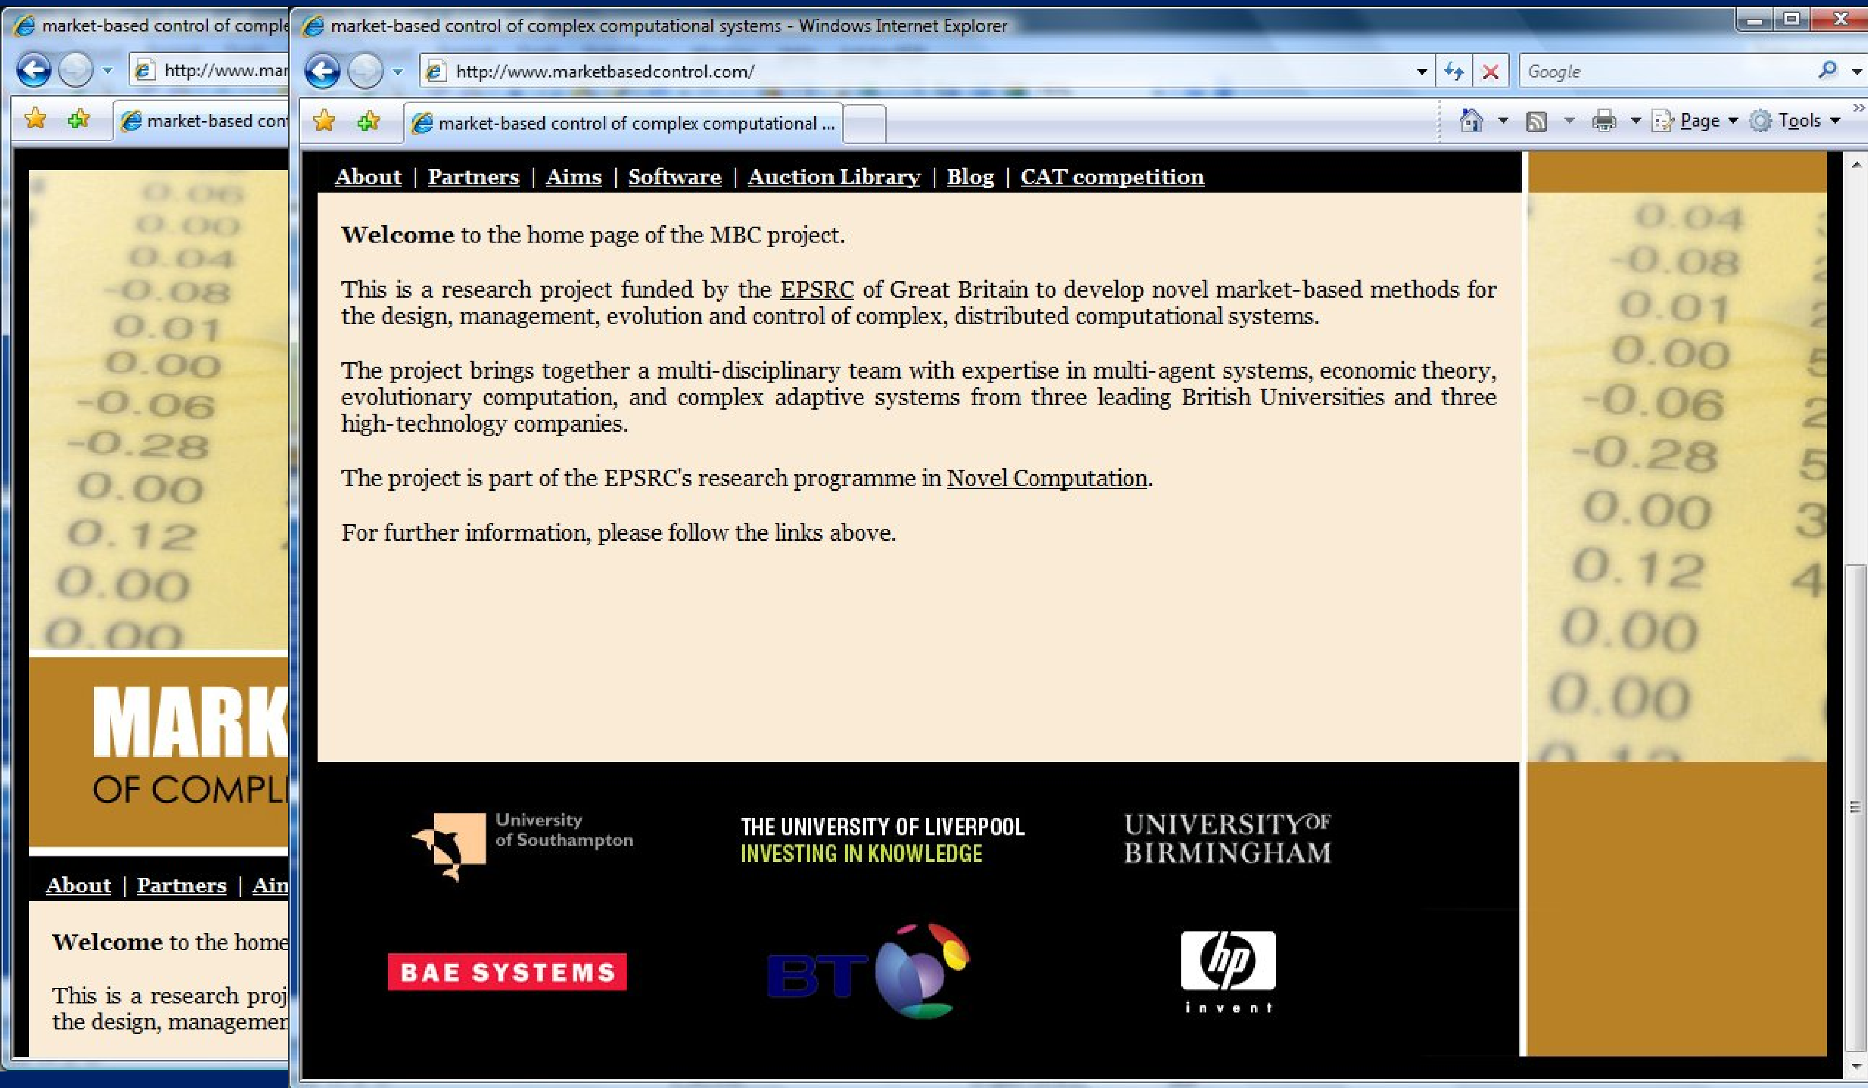

Bernardo Huberman: PARC to HPLPA

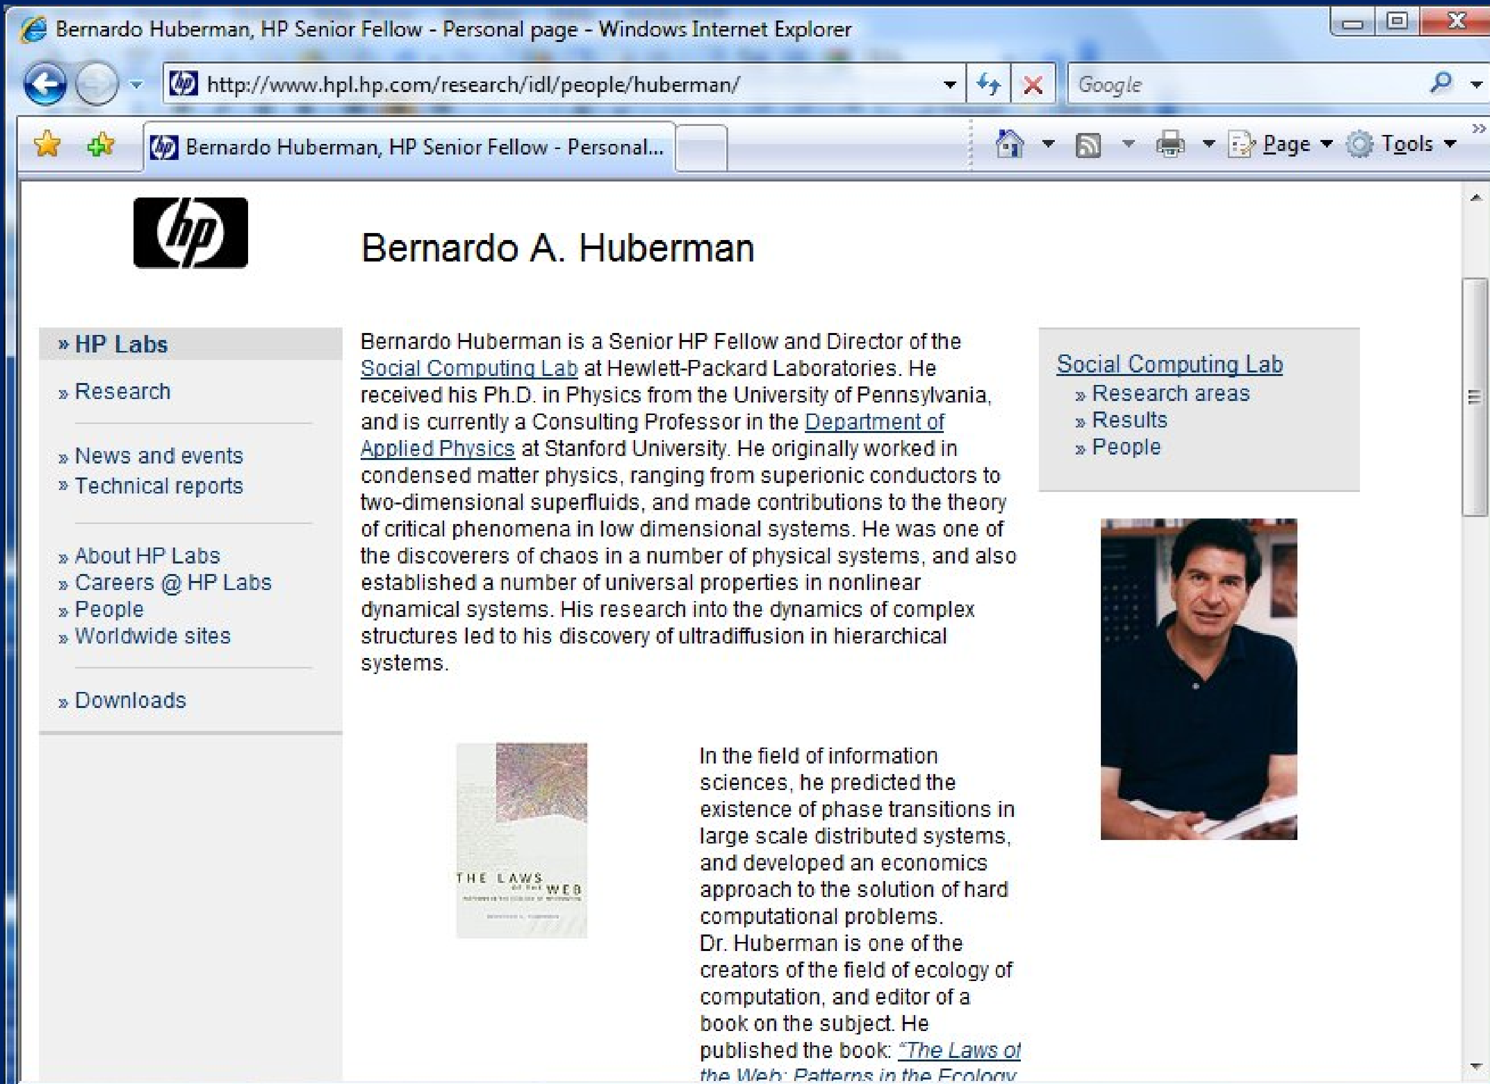

Bernardo Huberman, HP Senior Fellow - Personal page - Windows Internet Explorer

http://www.hpl.hp.com/research/idl/people/huberman/

Bernardo Huberman, HP Senior Fellow - Personal...

Page ▾ Tools ▾

Bernardo A. Huberman

» HP Labs
» Research
» News and events
» Technical reports
» About HP Labs
» Careers @ HP Labs
» People
» Worldwide sites
» Downloads

Bernardo Huberman is a Senior HP Fellow and Director of the Social Computing Lab at Hewlett-Packard Laboratories. He received his Ph.D. in Physics from the University of Pennsylvania, and is currently a Consulting Professor in the Department of Applied Physics at Stanford University. He originally worked in condensed matter physics, ranging from superionic conductors to two-dimensional superfluids, and made contributions to the theory of critical phenomena in low dimensional systems. He was one of the discoverers of chaos in a number of physical systems, and also established a number of universal properties in nonlinear dynamical systems. His research into the dynamics of complex structures led to his discovery of ultradiffusion in hierarchical systems.

THE LAWS OF THE WEB

In the field of information sciences, he predicted the existence of phase transitions in large scale distributed systems, and developed an economics approach to the solution of hard computational problems. Dr. Huberman is one of the creators of the field of ecology of computation, and editor of a book on the subject. He published the book: "The Laws of the Web: Patterns in the Ecology

Social Computing Lab
» Research areas
» Results
» People

HP *Tycoon:* tycoon.hpl.hp.com

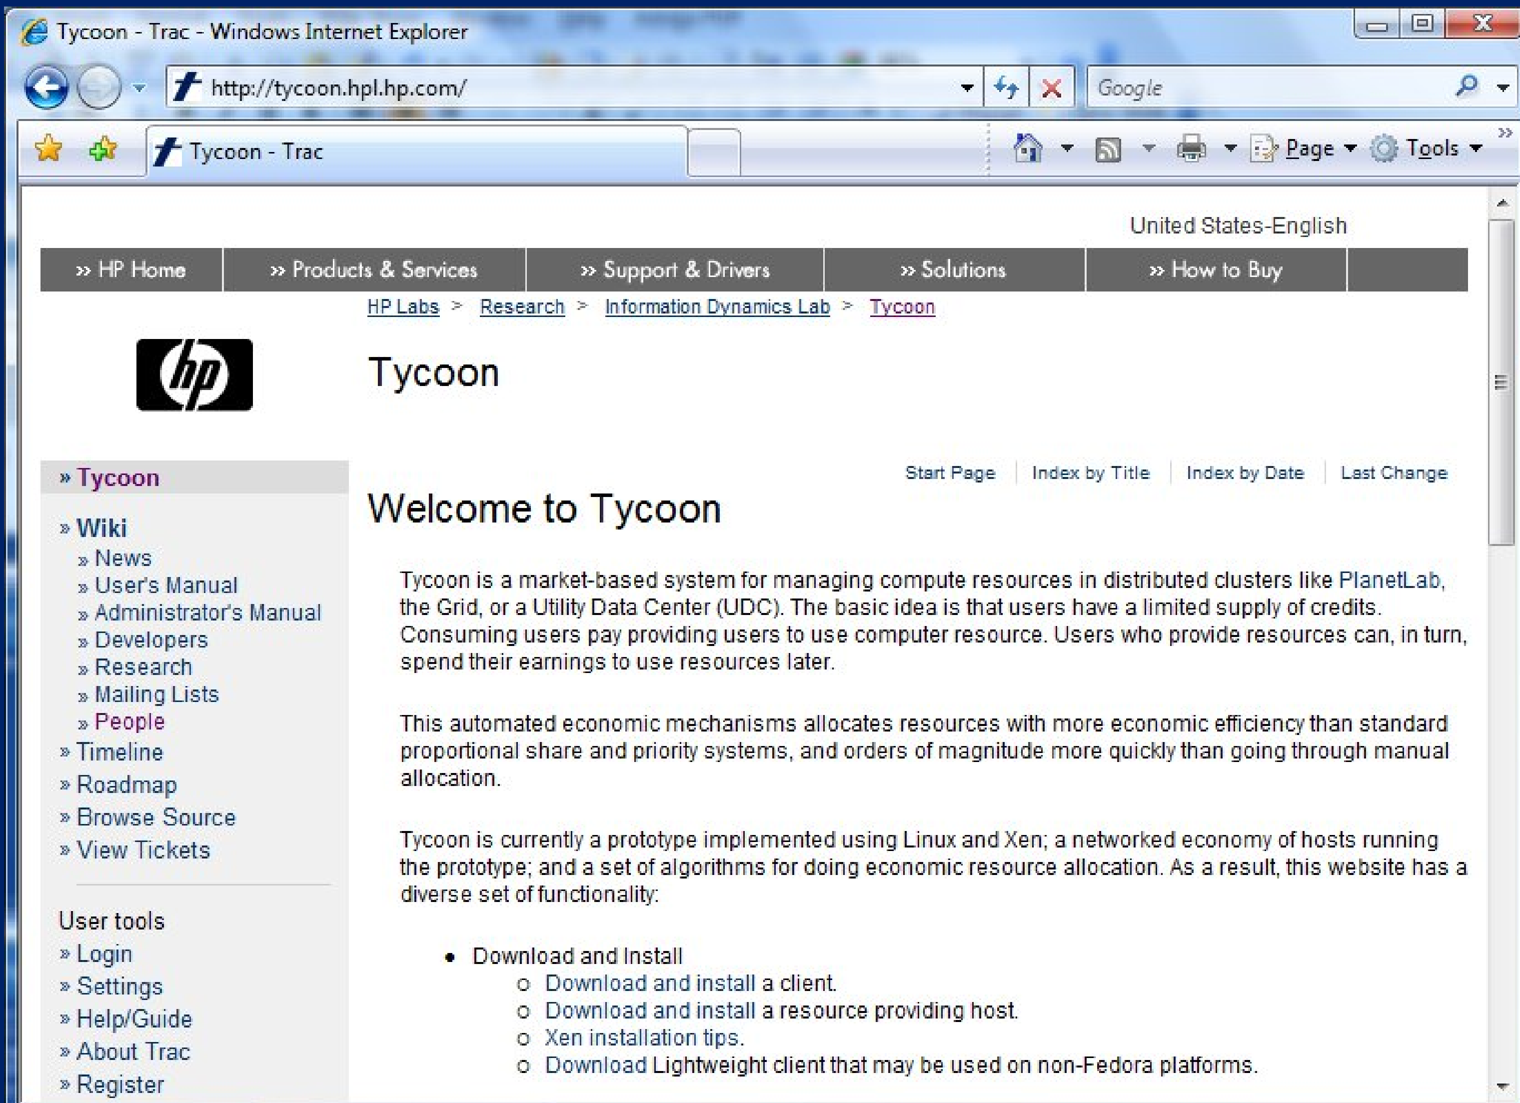

A selection of recent HP papers…

A Byde, M Salle, C Bartolini (2003) *Market-Based Resource Allocation for Utility Data Centers.* http://www.hpl.hp.com/techreports/2003/HPL-2003-188.pdf

K Lai, B Huberman, & L Fine (2004) *Tycoon: A Distributed Market-Based Resource Allocation System.* http://arxiv.org/PS_cache/cs/pdf/0404/0404013.pdf

S. Clearwater & B. A. Huberman. (2005) “Swing Options”. *Proc. 11th International Conference on Computing in Economics.*

http://www.hpl.hp.com/research/idl/papers/swings/swings.pdf

A. Byde (2006) A Comparison Between Mechanisms for Sequential Compute Resource Auctions. *Proc. AAMAS 06.*

http://www.hpl.hp.com/personal/Andrew_Byde/publications/2006.05%20-%20aamas.pdf

UK "E-science" for Grid : www.lesc.ic.ac.uk/markets

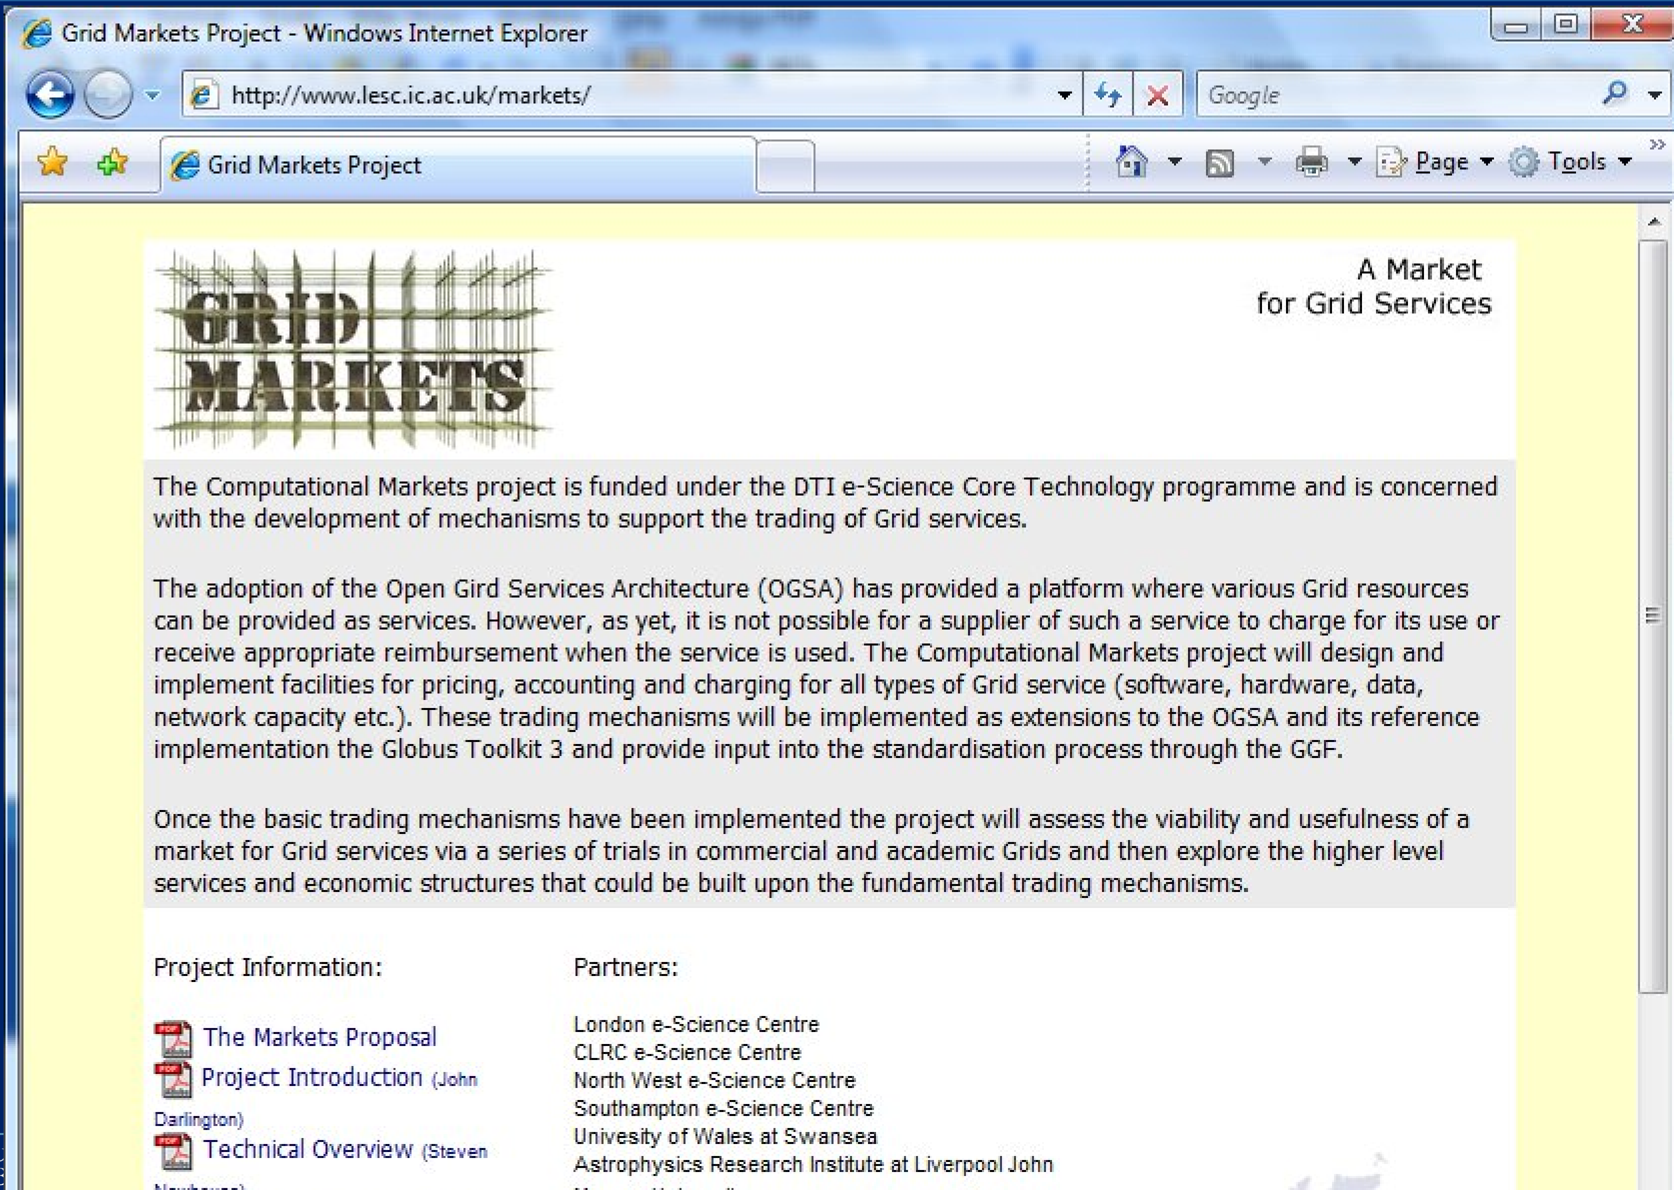

In the dealing rooms…

- It happens that the robot traders are of significant interest to real-world finance houses with multi-million-dollar flows on their trading floors…
- Most of the costs of running a trading floor are salaries & bonuses
- Robot traders don't need salaries or bonuses
- do the math

D.Hart-Davis /DHD Photo Gallery (http://gallery.hd.org)

Meanwhile, tales from the city…

- “Algorithmic Trading” became a hot topic in financial markets c. ~2002
- Initial focus on:
  - Automated execution for reduced “market impact” (VWAP)
  - Managing risk (ARM)
- But “feature creep” led to systems that, essentially, replace human traders
- Lower latency is usually better
  - Mathematical sophistication is useful, but “if you snooze, you lose.”
  - “Quant” skills needed, but algorithmic sophistication counts too
  - Simple algorithms, like ZIP, are very low latency; and ZIP was open-source

…more on this in Lecture 3.

Lecture 2: Summary

- Some intelligence/adaptation/learning is useful/necessary in trader agents
- ZIP & (M)GD the first two adaptive minimal algorithms for general CDA
- Lots of related works since then
- ZIP & MGD still the only algorithms shown to outperfrom humans
- Ability to deal with dynamically changing supply and demand is important
  - Something that much of the agents literature too often overlooks
- Simplicity = low-latency = attractive in the financial markets
- Open-sourcing ZIP was a very smart move, in retrospect