IBM Research

Autonomic Computing: The First Decade

Jeff Kephart (kephart@us.ibm.com)
IBM Thomas J Watson Research Center
Hawthorne, NY, USA

Outline

- Birth
- Formative Years
- What Have we Accomplished?
  - And what we have not?

In the beginning there was Chaos

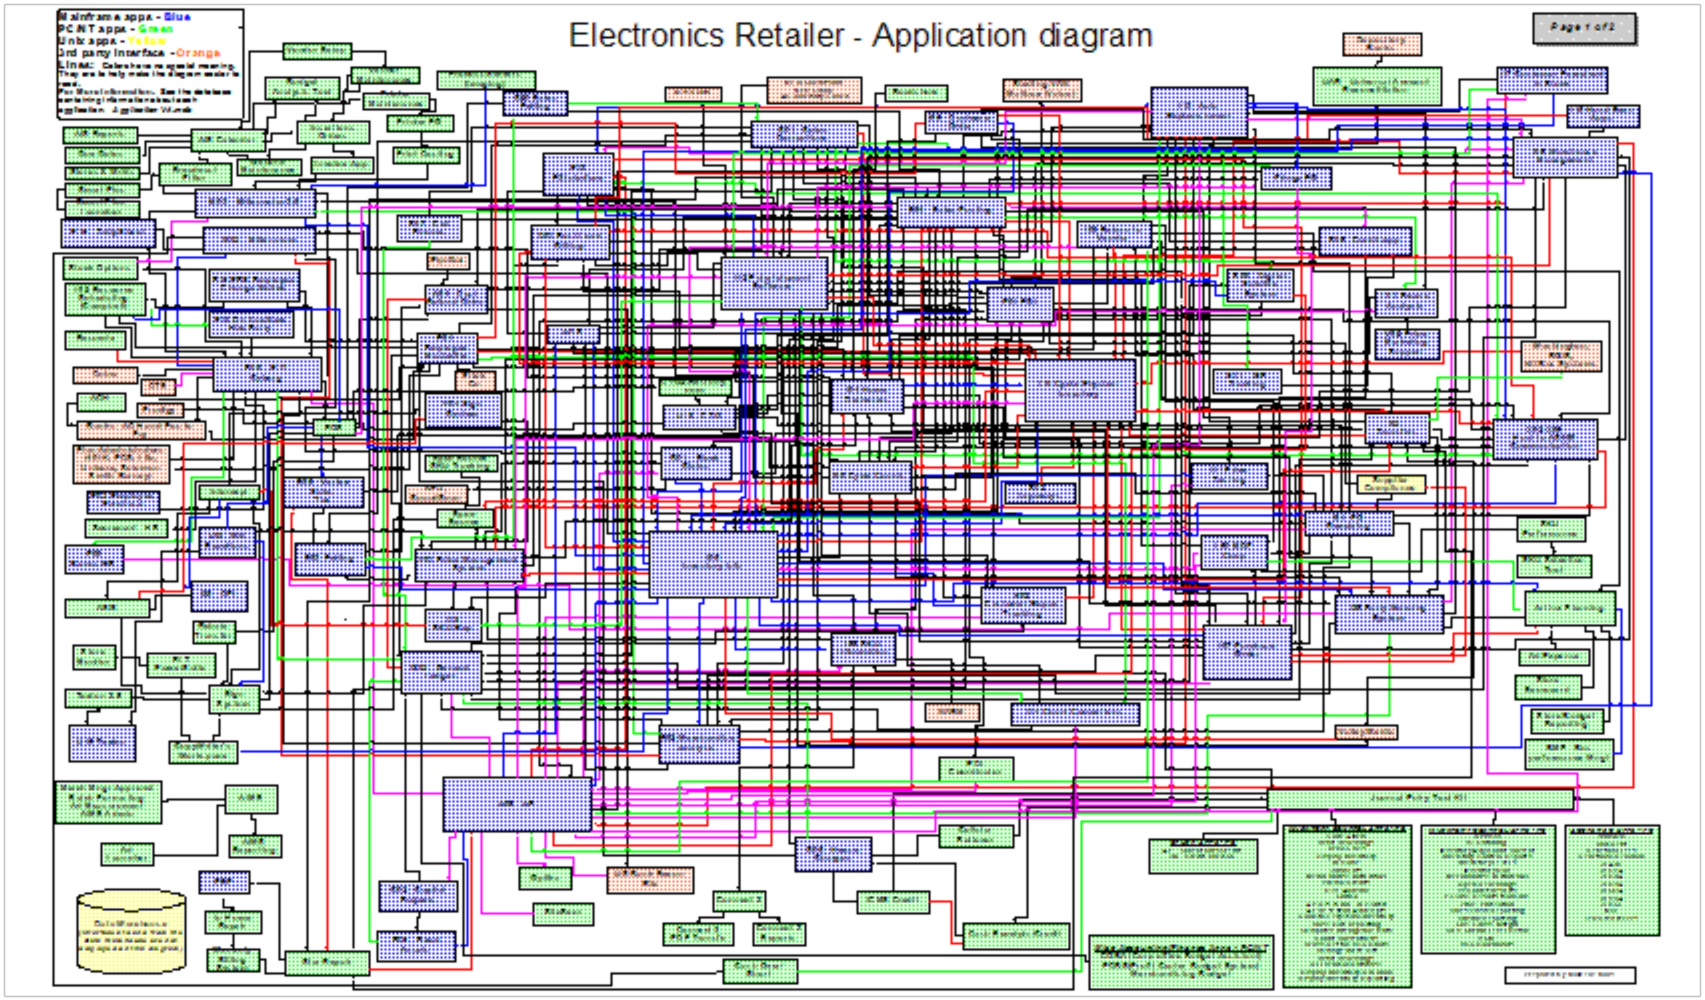

Where it All Began:
The Autonomic Computing Manifesto

- IBM Senior Research VP Paul Horn first set forth the idea of Autonomic Computing in keynote to National Academy of Engineers
- Harvard University, October 2001
- Autonomic Computing Manifesto released immediately thereafter

THE INFORMATION TECHNOLOGY INDUSTRY LOVES TO PROVE THE IMPOSSIBLE POSSIBLE.

We obliterate barriers and set records with astonishing regularity. But now we face a problem springing from the very core of our success — and too few of us are focused on solving it.

More than any other I/T problem, this one — if it remains unsolved — will actually prevent us from moving to the next era of computing. Interestingly enough, it has little to do with the usual barriers that preoccupy us.

It's not about keeping pace with Moore's Law, but rather dealing with the consequences of its decades-long reign. It's not directly related to how many bits we can squeeze into a square inch, or how thinly we can etch lines in silicon. In fact, a continued obsession with the smaller/faster/cheaper triumvirate is really a distraction.

It's not a barrier of "machine intelligence," either, that threatens our progress. It has less to do with building "thinking machines" that embody the popular conception of artificial intelligence (AI) than automating the day-to-day functioning of computing systems. It may sound odd coming from the creators of Deep Blue, but we don't really need a better chess-playing supercomputer — or sentient machines and androids programmed to love and laugh — to overcome the largest obstacle standing in our way.

The obstacle is complexity. Dealing with it is the single most important challenge facing the I/T industry.

It is our next Grand Challenge.

How will this possibly help?

By embedding the complexity in the system infrastructure itself—both hardware and software — then automating its management. For this approach we find inspiration in the massively complex systems of the human body.

Think for a moment about one such system at work in our bodies, one so seamlessly embedded we barely notice it: *the autonomic nervous system*.

It tells your heart how fast to beat, checks your blood's sugar and oxygen levels, and controls your pupils so the right amount of light reaches your eyes as you read these words. It monitors your temperature and adjusts your blood flow and skin functions to keep it at 98.6°F. It controls the digestion of your food and your reaction to stress—it can even make your hair stand on end if you're sufficiently frightened. It carries out these functions across a wide range of external conditions, always maintaining a steady internal state called homeostasis while readying your body for the task at hand. SEE FIGURE 3

But most significantly, it does all this without any conscious recognition or effort on your part. This allows you to think about what you want to do, and not how you'll do it: you can make a mad dash for the train without having to calculate how much faster to breathe and pump your heart, or if you'll need that little dose of adrenaline to make it through the doors before they close.

It's as if the autonomic nervous system says to you, *Don't think about it — no need to. I've got it all covered.*

THAT'S PRECISELY HOW WE NEED TO BUILD COMPUTING SYSTEMS—AN APPROACH WE PROPOSE AS

autonomic computing.

FIGURE 3

The Autonomic Nervous System at Work

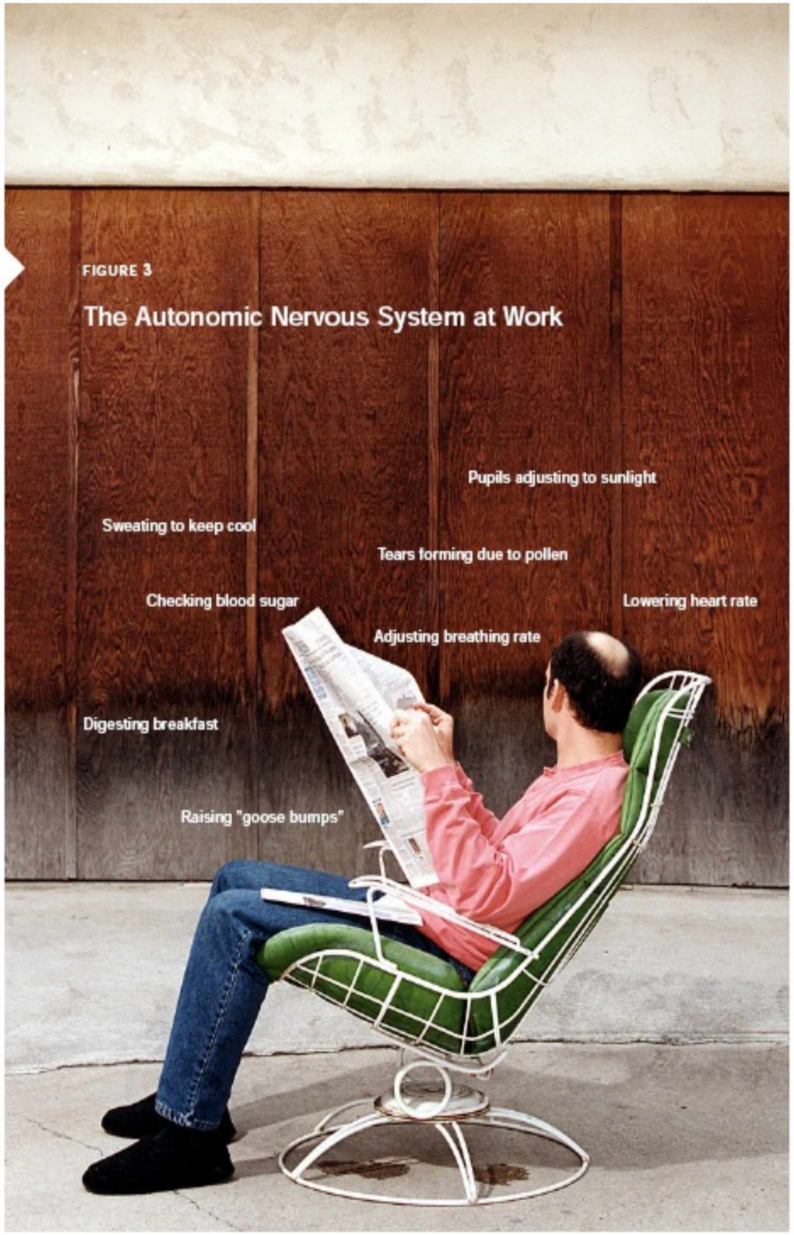

Eight Key Elements of an Autonomic Computing System

1 *To be* AUTONOMIC, a computing system needs to "know itself"—and comprise components that also possess a system identity.

2 *An* AUTONOMIC COMPUTING SYSTEM must configure and reconfigure itself under varying and unpredictable conditions.

3 *An* AUTONOMIC COMPUTING SYSTEM never settles for the status quo — it always looks for ways to optimize its workings.

4 *An* AUTONOMIC COMPUTING SYSTEM must perform something akin to healing — it must be able to recover from routine and extraordinary events that might cause some of its parts to malfunction.

5 *A* VIRTUAL WORLD is no less dangerous than the physical one, so an autonomic computing system must be an expert in self-protection.

6 *An* AUTONOMIC COMPUTING SYSTEM knows its environment and the context surrounding its activity, and acts accordingly.

7 *An* AUTONOMIC COMPUTING SYSTEM cannot exist in a hermetic environment.

8 *Perhaps* MOST CRITICAL for the user, an autonomic computing system will anticipate the optimized resources needed while keeping its complexity hidden.

This was soon boiled down to four …

Outline

- Birth
- Formative Years
- What Have we Accomplished?
  - And what have we not?

IBM's Internal Realignment to Support AC

- Created new Autonomic Computing group within Systems Management division
  - Alan Ganek, VP of Autonomic Computing
  - Autonomic Computing architecture board
- Created a new Autonomic Computing department within Research Division in 2002
  - Approximately 20 individuals
  - Approximately 100 researchers working on AC across IBM
- Created a new Joint Program to guide and fund AC Research
  - Dave Kaminsky/Tom Corbi and Jeff Kephart

IBM wanted to help drive a new research agenda

THIS IS

BIGGER THAN ANY *single*

I/T COMPANY.

It's a vision that requires the involvement of the top minds in the technology community. That's why we're forming an advisory board of leading academic and industry thinkers to help define our autonomic research agenda. We're also consulting with a large group of customers and I/T partners as part of our eLiza project to define a strategy for bringing autonomic innovations to products.

We call on our academic colleagues to drive exploratory work in autonomic computing. We propose that the research community recognize it as an important field of academic endeavor. We also call on our partners at government labs to collaborate with us on crucial projects in this area. We plan to fund a regular stream of academic grants to support research in this area, and we call on others in the I/T industry to do the same.

Paul Horn

Paul Horn, *Senior Vice President*

IBM RESEARCH

Autonomic Computing Advisory Board

- We established an AC Advisory Board in 2002
  - **Sponsors**: IBM Research VPs Alfred Spector, Robert Morris, Tilak Agrawala
  - **Chair**: Jeff Kephart

- Mission
  - Help define appropriate research agendas and curricula
  - Contribute insights on what are the relevant problems
  - Stimulate interest in AC issues of relevance within and across their respective fields
  - Endorse and legitimize autonomic computing within industry and academia

- We recruited 8 top academics and 5 key industry experts
  - Professors of AI, Distributed Systems, Grid Computing

- We presented IBM's AC research and solicited
  - Feedback on our research
    - More on self-healing, self-protection, human interaction; deeper work on policy; clarity architecture; build system prototypes
  - Advice on how to enlist academia to work on the great AC challenges

AC Advisory Board Recommendations for Recruiting Academia

1. Publish a well-placed, high-quality manifesto
2. Show that AC is radical, revolutionary, world-changing
   a. Publicize IBM's own high-quality research in AC
   b. Target top academics
      - Define problem in their specific terms
      - If they write good papers, rest of field will follow
3. Demonstrate *industry-wide* interest in AC (not just IBM hype)
4. Organize, sponsor, and participate in workshops, conferences
   a. International conferences and workshops
   b. Special IBM AC workshops

AC Advisory Board Recommendations for Recruiting Academia

1. Publish a well-placed, high-quality manifesto

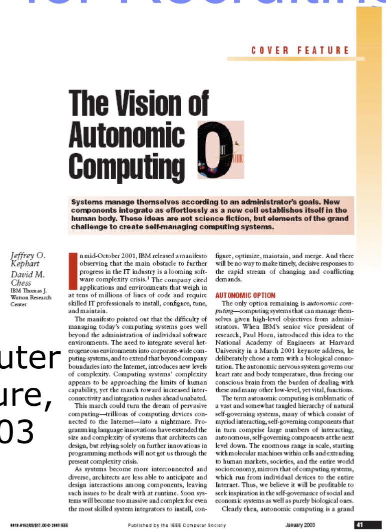

COVER FEATURE

The Vision of Autonomic Computing

Systems manage themselves according to an administrator's goals. New components integrate as effortlessly as a new cell establishes itself in the human body. These ideas are not science fiction, but elements of the grand challenge to create self-managing computing systems.

Jeffrey O. Kephart
David M. Chess
IBM Thomas J. Watson Research Center

In mid-October 2001, IBM released a manifesto observing that the main obstacle to further progress in the IT industry is a looming software complexity crisis.[1] The company cited applications and environments that weigh in at tens of millions of lines of code and require skilled IT professionals to install, configure, tune, and maintain.

The manifesto pointed out that the difficulty of managing today's computing systems goes well beyond the administration of individual software environments. The need to integrate several heterogeneous environments into corporate-wide computing systems, and to extend that beyond company boundaries into the Internet, introduces new levels of complexity. Computing systems' complexity appears to be approaching the limits of human capability, yet the march toward increased interconnectivity and integration rushes ahead unabated.

This march could turn the dream of pervasive computing—trillions of computing devices connected to the Internet—into a nightmare. Programming language innovations have extended the size and complexity of systems that architects can design, but relying solely on further innovations in programming methods will not get us through the present complexity crisis.

As systems become more interconnected and diverse, architects are less able to anticipate and design interactions among components, leaving such issues to be dealt with at runtime. Soon systems will become too massive and complex for even the most skilled system integrators to install, configure, optimize, maintain, and merge. And there will be no way to make timely, decisive responses to the rapid stream of changing and conflicting demands.

AUTONOMIC OPTION

The only option remaining is *autonomic computing*—computing systems that can manage themselves given high-level objectives from administrators. When IBM's senior vice president of research, Paul Horn, introduced this idea to the National Academy of Engineers at Harvard University in a March 2001 keynote address, he deliberately chose a term with a biological connotation. The autonomic nervous system governs our heart rate and body temperature, thus freeing our conscious brain from the burden of dealing with these and many other low-level, yet vital, functions.

The term autonomic computing is emblematic of a vast and somewhat tangled hierarchy of natural self-governing systems, many of which consist of myriad interacting, self-governing components that in turn comprise large numbers of interacting, autonomous, self-governing components at the next level down. The enormous range in scale, starting with molecular machines within cells and extending to human markets, societies, and the entire world socioeconomy, mirrors that of computing systems, which run from individual devices to the entire Internet. Thus, we believe it will be profitable to seek inspiration in the self-governance of social and economic systems as well as purely biological ones.

Clearly then, autonomic computing is a grand

Published by the IEEE Computer Society January 2003 41

IEEE Computer Cover Feature, January 2003

2. Show that AC is radical, revolutionary, world-changing
   a. Publicize IBM's own high-quality research in AC
   b. Target top academics
      - Define problem in their specific terms
      - If they write good papers, rest of field will follow

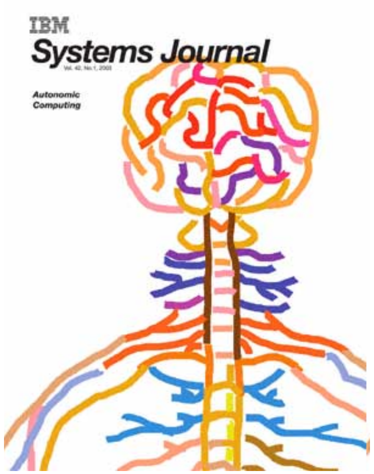

IBM Systems Journal Vol. 42, No. 1, 2003 Autonomic Computing

Targeting top professors; organizing conferences: Timeline

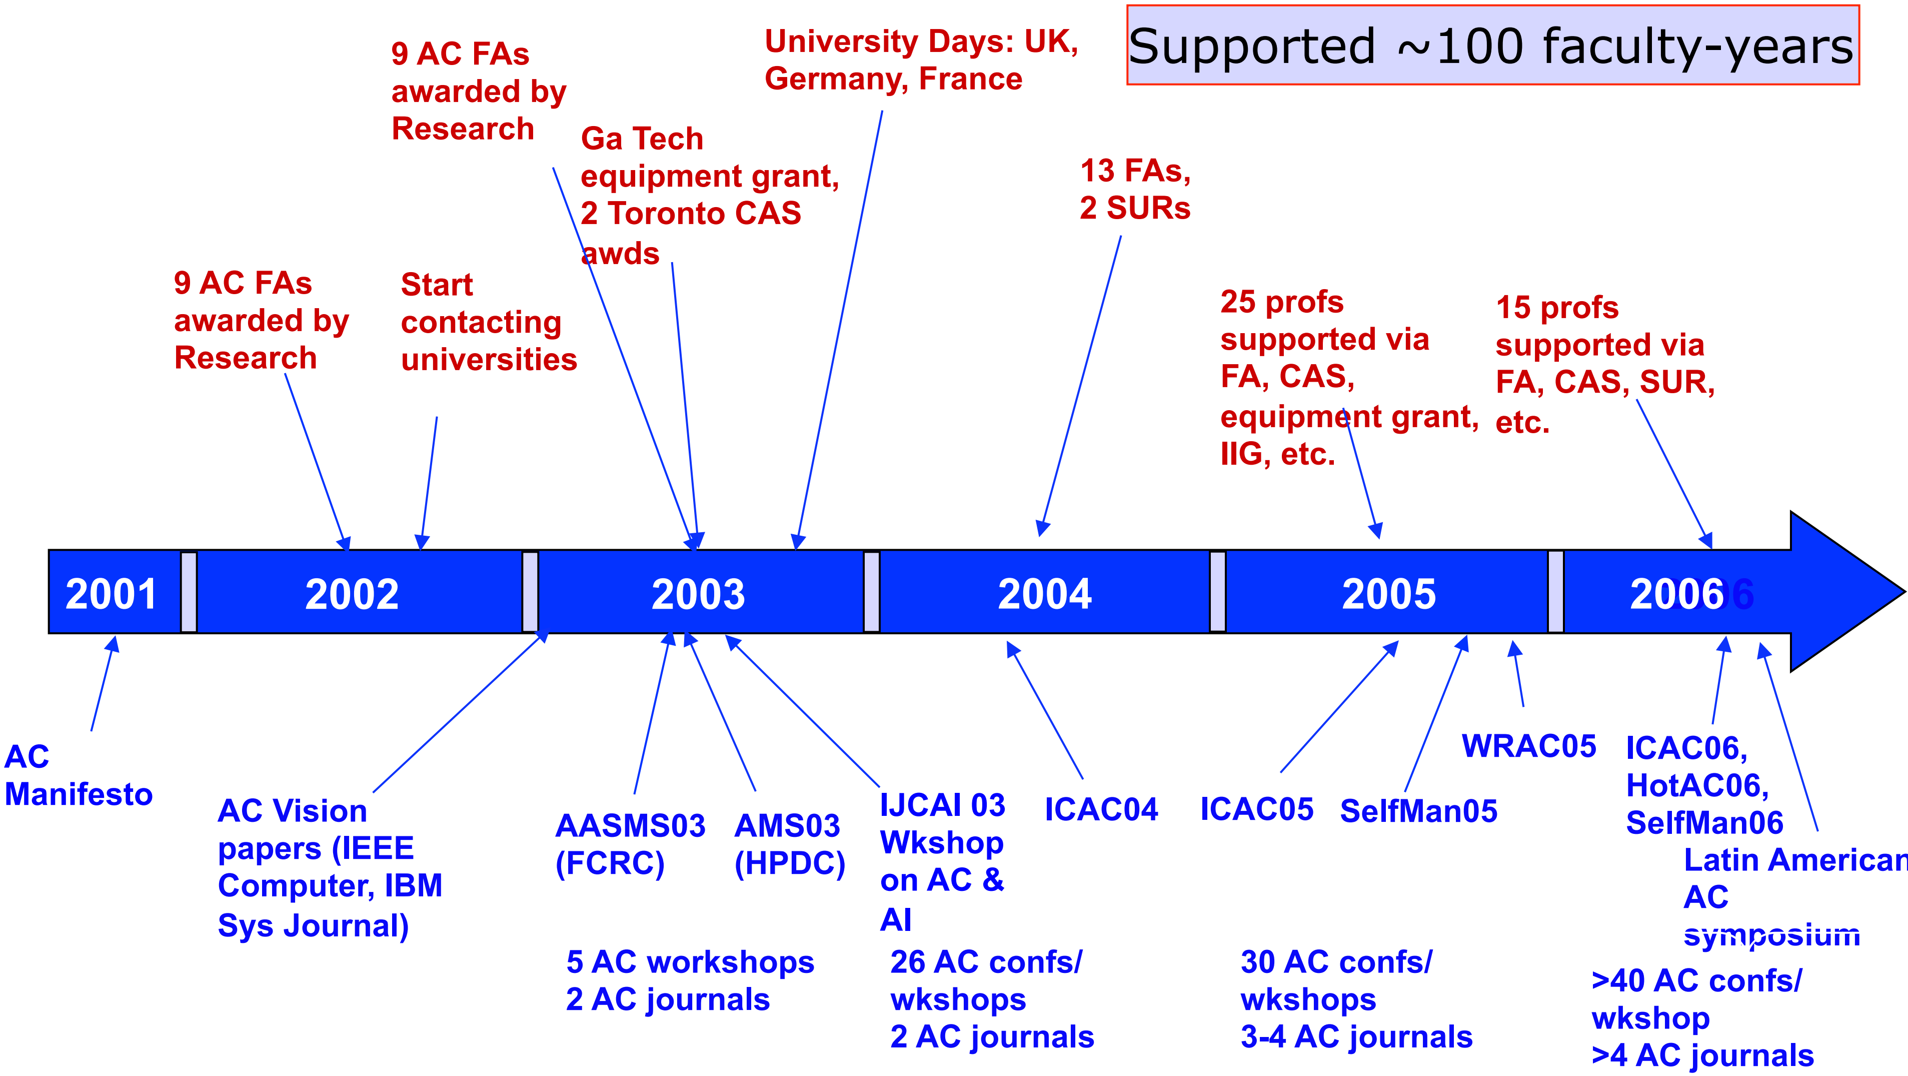

Organize and sponsor workshops, conferences

AASMS 03

Algorithms & Architectures for Self-Managing Systems

Federated Computing Research Conference

June ’03, San Diego, CA

Chase (Duke), Goldszmidt&Keeton (HP), Kephart&Tetzlaff (IBM)

AMS 03

Active Middleware Workshop on Autonomic Computing

High Performance Distributed Computing

June ’03 Seattle, WA

Hariri (Arizona), Parashar (Rutgers)

ICAC ‘04

International Conference on Autonomic Computing

May 17-18, 2004, New York

Establish an AC research community to work together to realize the vision of large-scale self-managing systems

ICAC ‘05

June 13-16, 2005, Seattle

10 demos
3 tutorials

Develop and nurture the AC research community

ICAC ‘06

June 12-16, 2006, Dublin

12 demos; 4 tutorials; 5 workshops

AC is catching on!

Excerpt from Report to Paul Horn on "Autonomic Computing as an Academic Discipline", late 2006

I'm flabbergasted!

- Initially spurred by our efforts
  - Faculty awards, equipment grants
  - Workshops, conferences (some IBM Academy)
  - University visits
  - Several classes taught by IBMers at Duke, UNC, St Andrews, Brazil
- But increasingly on its own
  - AC classes being taught around the world
    - >30 universities have AC content in their curricula
    - "Self-Managing Systems", Shivnath Babu, Duke University
    - "Autonomic Computing", Omer F. Rana, Cardiff U., UK, ½ day seminar.
    - "Parallel and Distributed Computing", Manish Parashar, Rutgers.
  - Government support: EPSRC in UK funds "Semantic Grid and Autonomic Computing Programme"
  - Over a dozen AC workshops, conferences initiated by non-IBMers
  - Publications
    - IEEE Task Force on Autonomous and Autonomic Systems newsletter
    - Special Issue of IEEE Internet Computing Jan 2007 on AC
    - ACM Transactions on Complex Adaptive Systems
  - Web site: www.autonomiccomputing.org

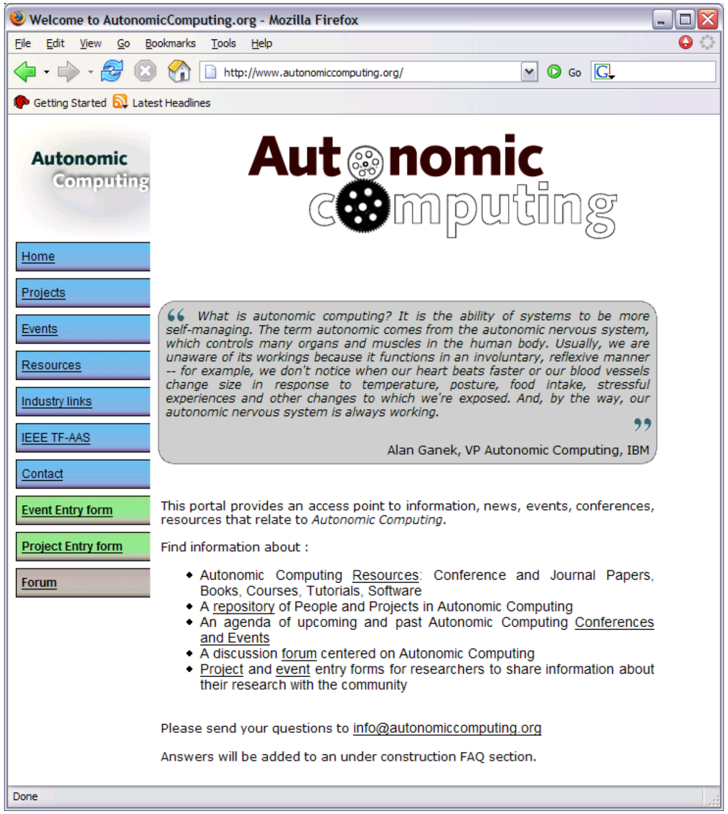

Outline

- Birth
- Formative Years
- What Have we Accomplished?
  - And what have we not?

Have we kept the momentum (five years later)?

- Over 8000 papers on autonomic computing
  - Approximately 160 ICAC papers (2% of literature)
- Over 200 patents issued on autonomic computing
  - >100 more under evaluation
- Nearly 200 conferences or workshops solicit papers on autonomic computing
- Government funding
  - FP6: Situated autonomic communications
    - ANA, BioNETS, CASCADAS, HAGGLE, ACCA
  - FP7: Self-awareness in autonomic systems

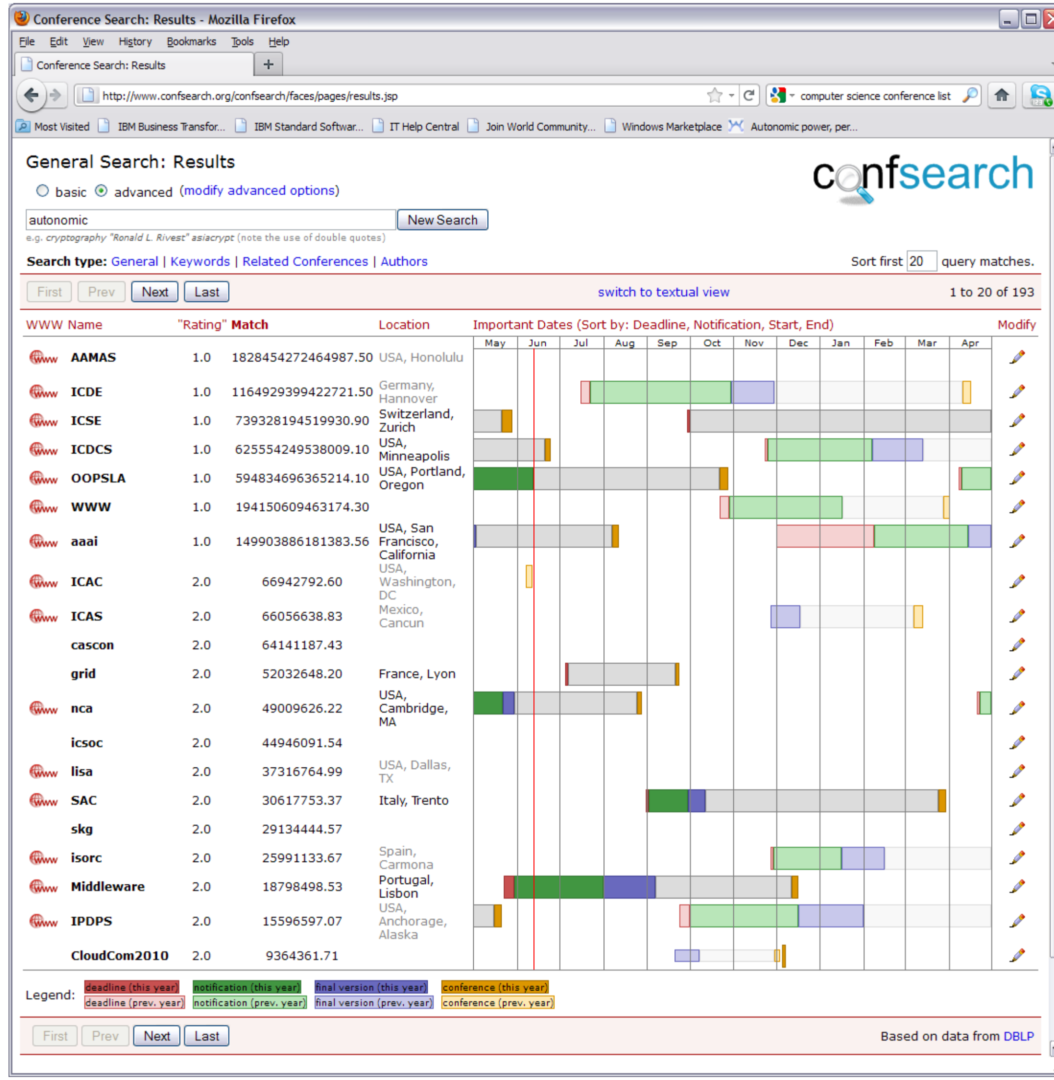

Let’s take a closer look at how AC is doing as a field

- Run Harzing’s Publish or Perish with queries “Autonomic Computing” and “International Conference on Autonomic Computing”
  - Uses Google Scholar; finds top 1000 papers in terms of citation counts
- Put structured data in spreadsheet
- Cleanse the data
- Identify interesting trends

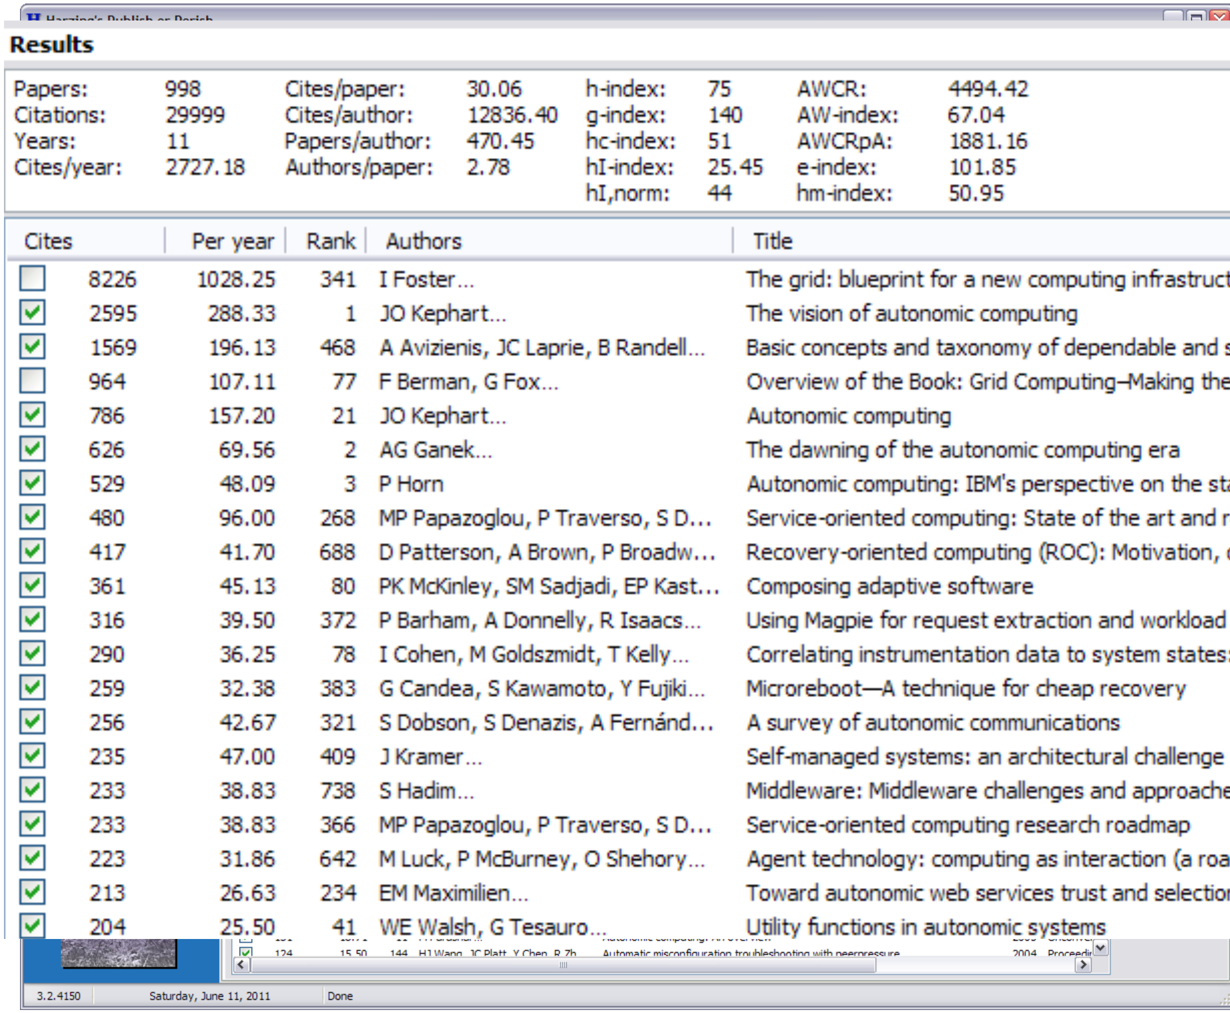

Results

Papers:	998	Cites/paper:	30.06	h-index:	75	AWCR:	4494.42
Citations:	29999	Cites/author:	12836.40	g-index:	140	AW-index:	67.04
Years:	11	Papers/author:	470.45	hc-index:	51	AWCRpA:	1881.16
Cites/year:	2727.18	Authors/paper:	2.78	hI-index:	25.45	e-index:	101.85
				hI,norm:	44	hm-index:	50.95

Cites	Per year	Rank	Authors	Title
8226	1028.25	341	I Foster...	The grid: blueprint for a new computing infrastructu
2595	288.33	1	JO Kephart...	The vision of autonomic computing
1569	196.13	468	A Avizienis, JC Laprie, B Randell...	Basic concepts and taxonomy of dependable and se
964	107.11	77	F Berman, G Fox...	Overview of the Book: Grid Computing–Making the G
786	157.20	21	JO Kephart...	Autonomic computing
626	69.56	2	AG Ganek...	The dawning of the autonomic computing era
529	48.09	3	P Horn	Autonomic computing: IBM's perspective on the stat
480	96.00	268	MP Papazoglou, P Traverso, S D...	Service-oriented computing: State of the art and res
417	41.70	688	D Patterson, A Brown, P Broadw...	Recovery-oriented computing (ROC): Motivation, de
361	45.13	80	PK McKinley, SM Sadjadi, EP Kast...	Composing adaptive software
316	39.50	372	P Barham, A Donnelly, R Isaacs...	Using Magpie for request extraction and workload m
290	36.25	78	I Cohen, M Goldszmidt, T Kelly...	Correlating instrumentation data to system states: A
259	32.38	383	G Candea, S Kawamoto, Y Fujiki...	Microreboot—A technique for cheap recovery
256	42.67	321	S Dobson, S Denazis, A Fernánd...	A survey of autonomic communications
235	47.00	409	J Kramer...	Self-managed systems: an architectural challenge
233	38.83	738	S Hadim...	Middleware: Middleware challenges and approaches
233	38.83	366	MP Papazoglou, P Traverso, S D...	Service-oriented computing research roadmap
223	31.86	642	M Luck, P McBurney, O Shehory...	Agent technology: computing as interaction (a road
213	26.63	234	EM Maximilien...	Toward autonomic web services trust and selection
204	25.50	41	WE Walsh, G Tesauro...	Utility functions in autonomic systems

3.2.4150 Saturday, June 11, 2011 Done

Autonomic Computing Papers (2002)

417	D Patterson, A Brown, P E	Recovery-oriented computing (ROC): Motivation, definition, techr	2002	UC Berkeley Tech Report
168	JP Bigus, DA Schlosnagle	ABLE: A toolkit for building multiagent autonomic systems	2002	IBM Systems …
123	N Zhong, J Liu…	In search of the wisdom web	2002	COMPUTER-LOS ALAMITOS-
84	R Sterritt	Towards autonomic computing: effective event management	2002	… Workshop, 2002. Proceedings. 27
63	A LaMarca, W Brunette, D	Plantcare: An investigation in practical ubiquitous systems	2002	UbiComp 2002: …
55	M Satyanarayanan	A catalyst for mobile and ubiquitous computing	2002	Pervasive Computing
52	D Paulson	Computer system, heal thyself	2002	Computer
44	GM Lohman…	SMART: Making DB2 (more) autonomic	2002	… of the 28th international conferenc
43	IBMA Computing	IBM's Perspective on the State of Information Technology	2002	White Paper, information available at
42	SS Lightstone, G Lohman	Toward autonomic computing with DB2 universal database	2002	ACM SIGMOD Record
39	S Elnaffar, P Martin…	Automatically classifying database workloads	2002	Proceedings of the eleventh …
27	E Mainsah	Autonomic computing: the next era of computing	2002	Electronics & Communication Engine
24	RK Sahoo, M Bae, R Vilalt	Providing persistent and consistent resources through event log a	2002	Workshop on Self- …
21	CH Crawford…	eModel: addressing the need for a flexible modeling framework ir	2002	Modeling, Analysis and Simulation of
20	DA Patterson	Recovery oriented computing: A new research agenda for a new	2002	Keynote address, HPCA
15	WW Gibbs	Autonomic computing	2002	Scientific American
14	E Schwartz	IBM Offers a Peek at Self-Healing PCS: Autonomic computing in	2002	Date Alleged: Nov
14	MN Huhns…	Robust software	2002	Internet Computing, IEEE
12	D Pescovitz	Helping computers help themselves	2002	Spectrum, IEEE
10	LD Paulson	IBM begins autonomic-computing project	2002	Computer
9	Y Tohma	Fault tolerance in autonomic computing environment	2002	
8	A Wolfe	News analysis: IBM sets its sights on autonomic computing	2002	IEEE Spectrum
6	DJ Clancy	NASA challenges in autonomic computing	2002	Almaden Institute
4	JY Chung…	„Beyond e-Marketplace & Next Generation e-Business: Grid, Aut	2002	4th International Conference on Elect
4	YS Tan, B Topol, V Vellan	Implementing service Grids with the service domain toolkit	2002	IBM Corporation
3	E Grishikashvili, N Badr, D	Autonomic computing: A service-oriented framework to support tl	2002	Proceeding of 3rd …
3	AZ Spector	Challenges and opportunities in autonomic computing	2002	Proceedings of the 16th international
2	J Kephart	Technology challenges of autonomic computing	2002	OOPSLA
2	R Sterritt	Towards autonomic computing: effective event management Sof	2002	Proceedings. 27th Annual NASA God

Autonomic Computing Papers (2002)

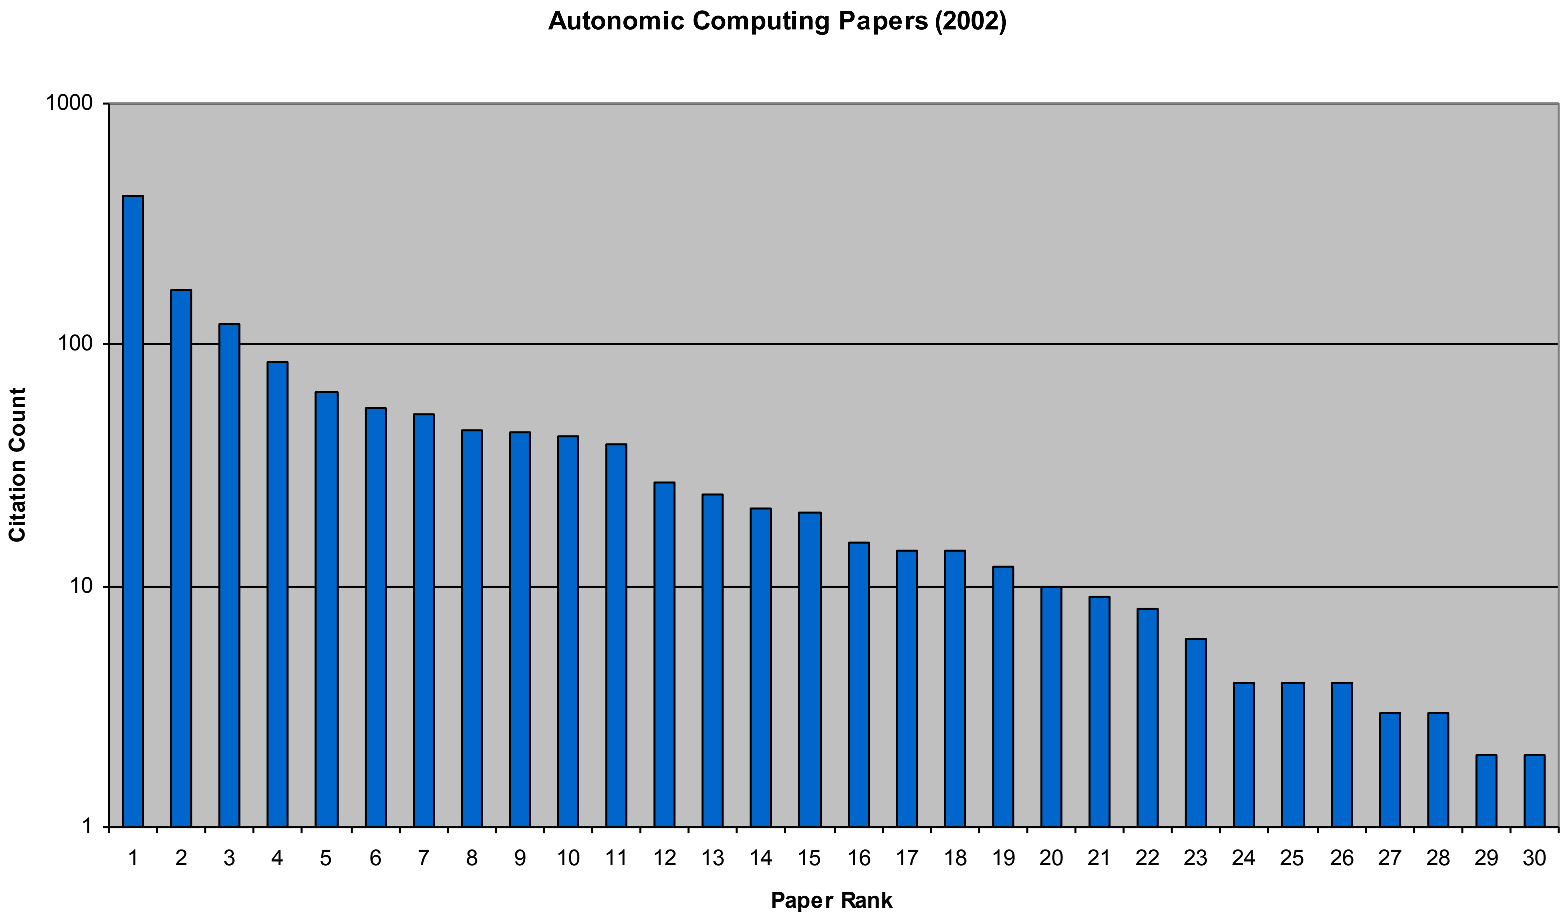

Autonomic Computing Papers (2003)

2595	JO Kephart…	The vision of autonomic computing	Computer
626	AG Ganek…	The dawning of the autonomic computing era	IBM Systems Journal
147	H Kreger	Fulfilling the Web services promise	Communications of the ACM
108	R Sterritt…	Autonomic Computing-a means of achieving dependability?	Engineering of Computer-Based …
106	J Appavoo, K Hui, CAN S	Enabling autonomic behavior in systems software with hot swapping	IBM systems …
104	G Kaiser, J Parekh, P Gr	Kinesthetics extreme: An external infrastructure for monitoring distributed legacy systems	… Autonomic Computing …
104	R Sterritt…	Towards an autonomic computing environment	Database and Expert Systems …
99	AB Brown…	Undo for operators: Building an undoable e-mail store	Proceedings of the annual conferenc
96	M Agarwal, V Bhat, H Liu	Automate: Enabling autonomic applications on the grid	… Computing …
96	H Cervantes…	Automating service dependency management in a service-oriented component model	Proceedings of CBSE
86	DM Chess, CC Palmer…	Security in an autonomic computing environment	IBM Systems Journal
84	DF Bantz, C Bisdikian, D	Autonomic personal computing	IBM Systems …
84	Y Diao, JL Hellerstein, S	Managing web server performance with autotune agents	IBM Systems Journal
82	R Want, T Pering…	Comparing autonomic and proactive computing	IBM Systems Journal
80	F Heylighen…	The meaning of self-organization in computing	Information Systems
78	X Dong, S Hariri, L Xue, H	Autonomia: an autonomic computing environment	… Proceedings of the …
75	A Leff, JT Rayfield…	Service-level agreements and commercial grids	IEEE Internet Computing
73	V Markl, GM Lohman…	LEO: An autonomic query optimizer for DB2	IBM Systems Journal
73	D Capera	The AMAS theory for complex problem solving based on self-organizing cooperative agents	
70	C Sapuntzakis…	Virtual appliances in the collective: A road to hassle-free computing	Proceedings of the 9th conference o
67	P Buhler, JM Vidal…	Adaptive workflow= web services+ agents	… of the International Conference or
61	A Dan, H Ludwig, G Paci	Web service differentiation with service level agreements	White Paper, IBM Corporation
57	RJT Morris…	The evolution of storage systems	IBM Systems Journal
55	EM Maximilien…	Agent-based architecture for autonomic web service selection	Workshop on Web Services and Age
53	J Jann, LM Browning…	Dynamic reconfiguration: Basic building blocks for autonomic computing on IBM pSeries s	IBM Systems Journal
53	M Milenkovic, SH Robins	Toward internet distributed computing	Computer
52	C Boutilier, R Das, JO Ke	Cooperative negotiation in autonomic systems using incremental utility elicitation	… on Uncertainty in …
47	R Sterritt	Pulse monitoring: extending the health-check for the autonomic GRID	… Informatics, 2003. INDIN 2003. Pr
44	M Agarwal…	Enabling autonomic compositions in grid environments	
44	JA Redstone, MM Swift…	Using computers to diagnose computer problems	… of the 9th Workshop on Hot …
42	F Berman, G Fox…	Grid computing	
40	JM Deegan…	High reliability memory subsystem using data error correcting code symbol sliced comma	US Patent App. 10/723,055
40	T De Wolf…	Towards Autonomic Computing: agent-based modelling, dynamical systems analysis, and	Industrial Informatics, 2003. INDIN …
38	S Elnaffar, W Powley, D	Today's DBMSs: How autonomic are they	Database and Expert …
36	S Hariri, L Xue, H Chen, M	Autonomia: an autonomic computing environment	IEEE International …
29	DM Russell, PP Maglio, R	Dealing with ghosts: Managing the user experience of autonomic computing	IBM Systems Journal
28	G Lanfranchi, PD Peruta,	Toward a new landscape of systems management in an autonomic computing environmer	IBM Systems …
27	S Lightstone, B Schiefer,	Autonomic computing for relational databases: the ten-year vision	…, 2003. INDIN 2003. …
27	H Tianfield	Multi-agent autonomic architecture and its application in e-medicine	Intelligent Agent Technology, 2003.
27	T Eymann, M Reinicke, C	Self-organizing resource allocation for autonomic networks	

Autonomic Computing Papers (2003)

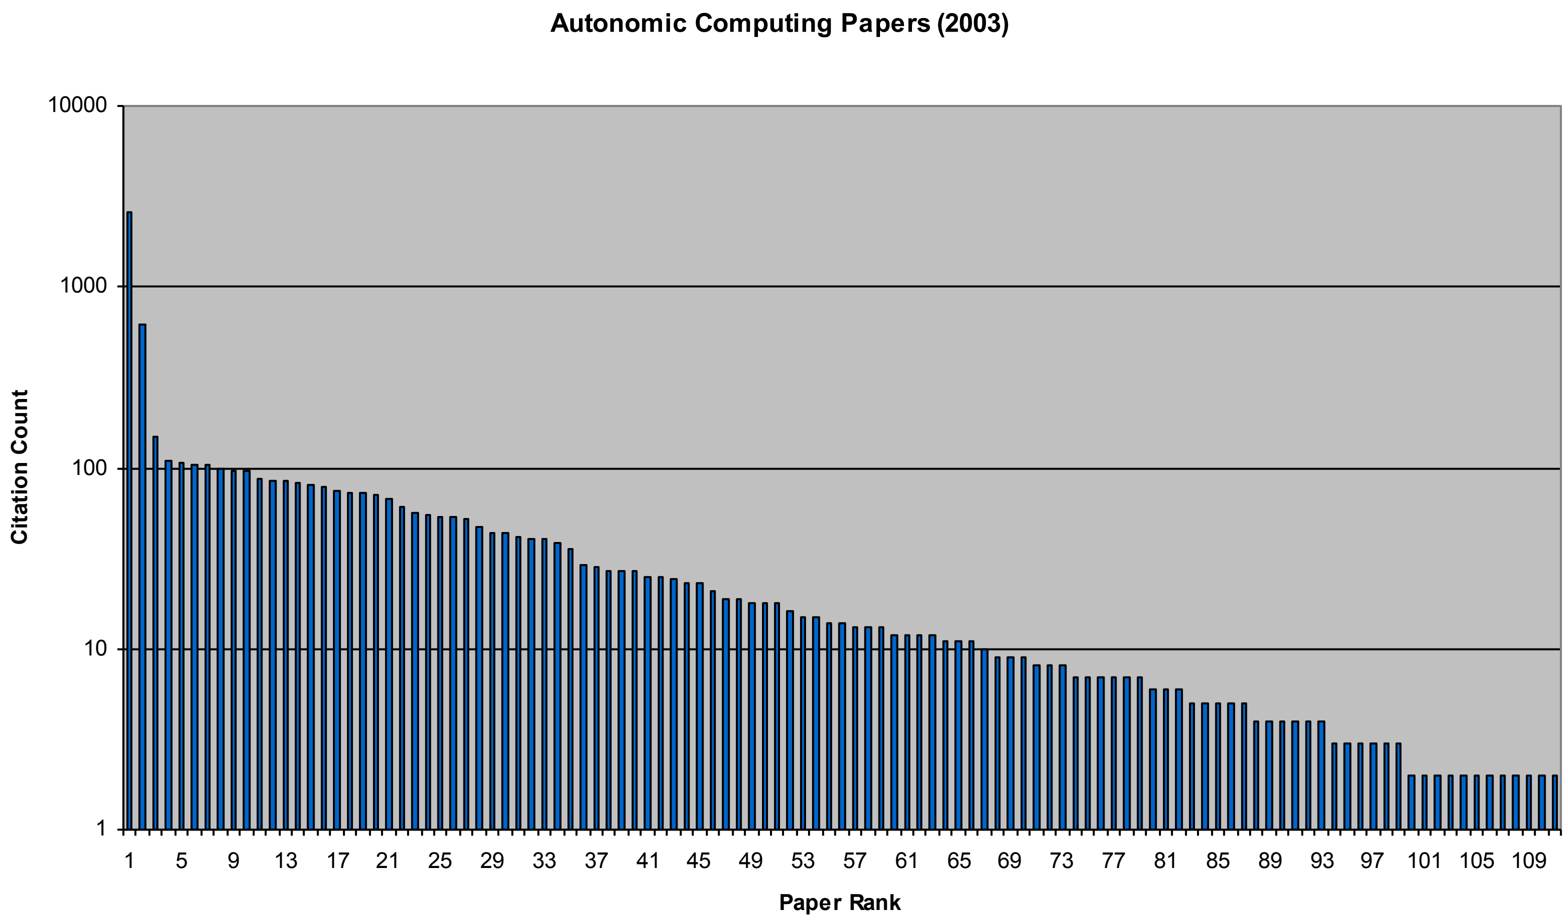

Autonomic Computing Papers (2004)

183 papers

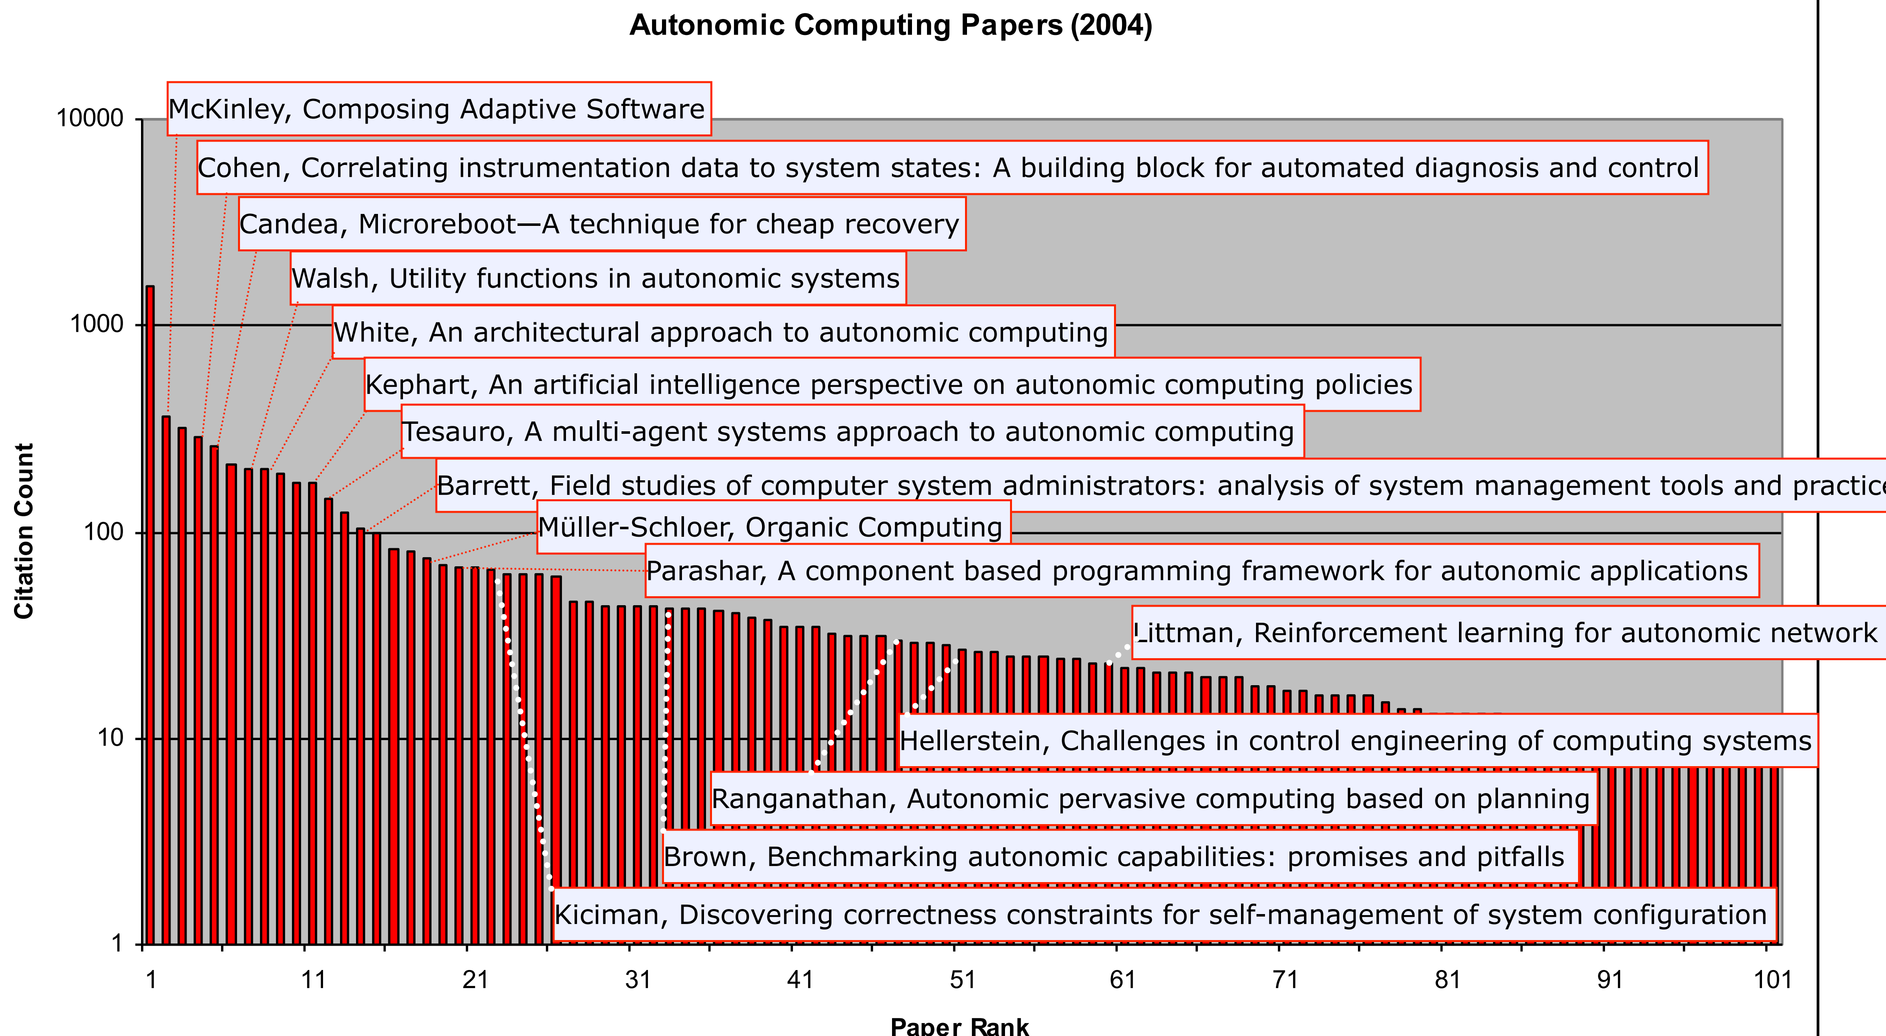

Autonomic Computing Papers (2004)
ICAC's impact

Total: 183 papers

ICAC04: 39 papers

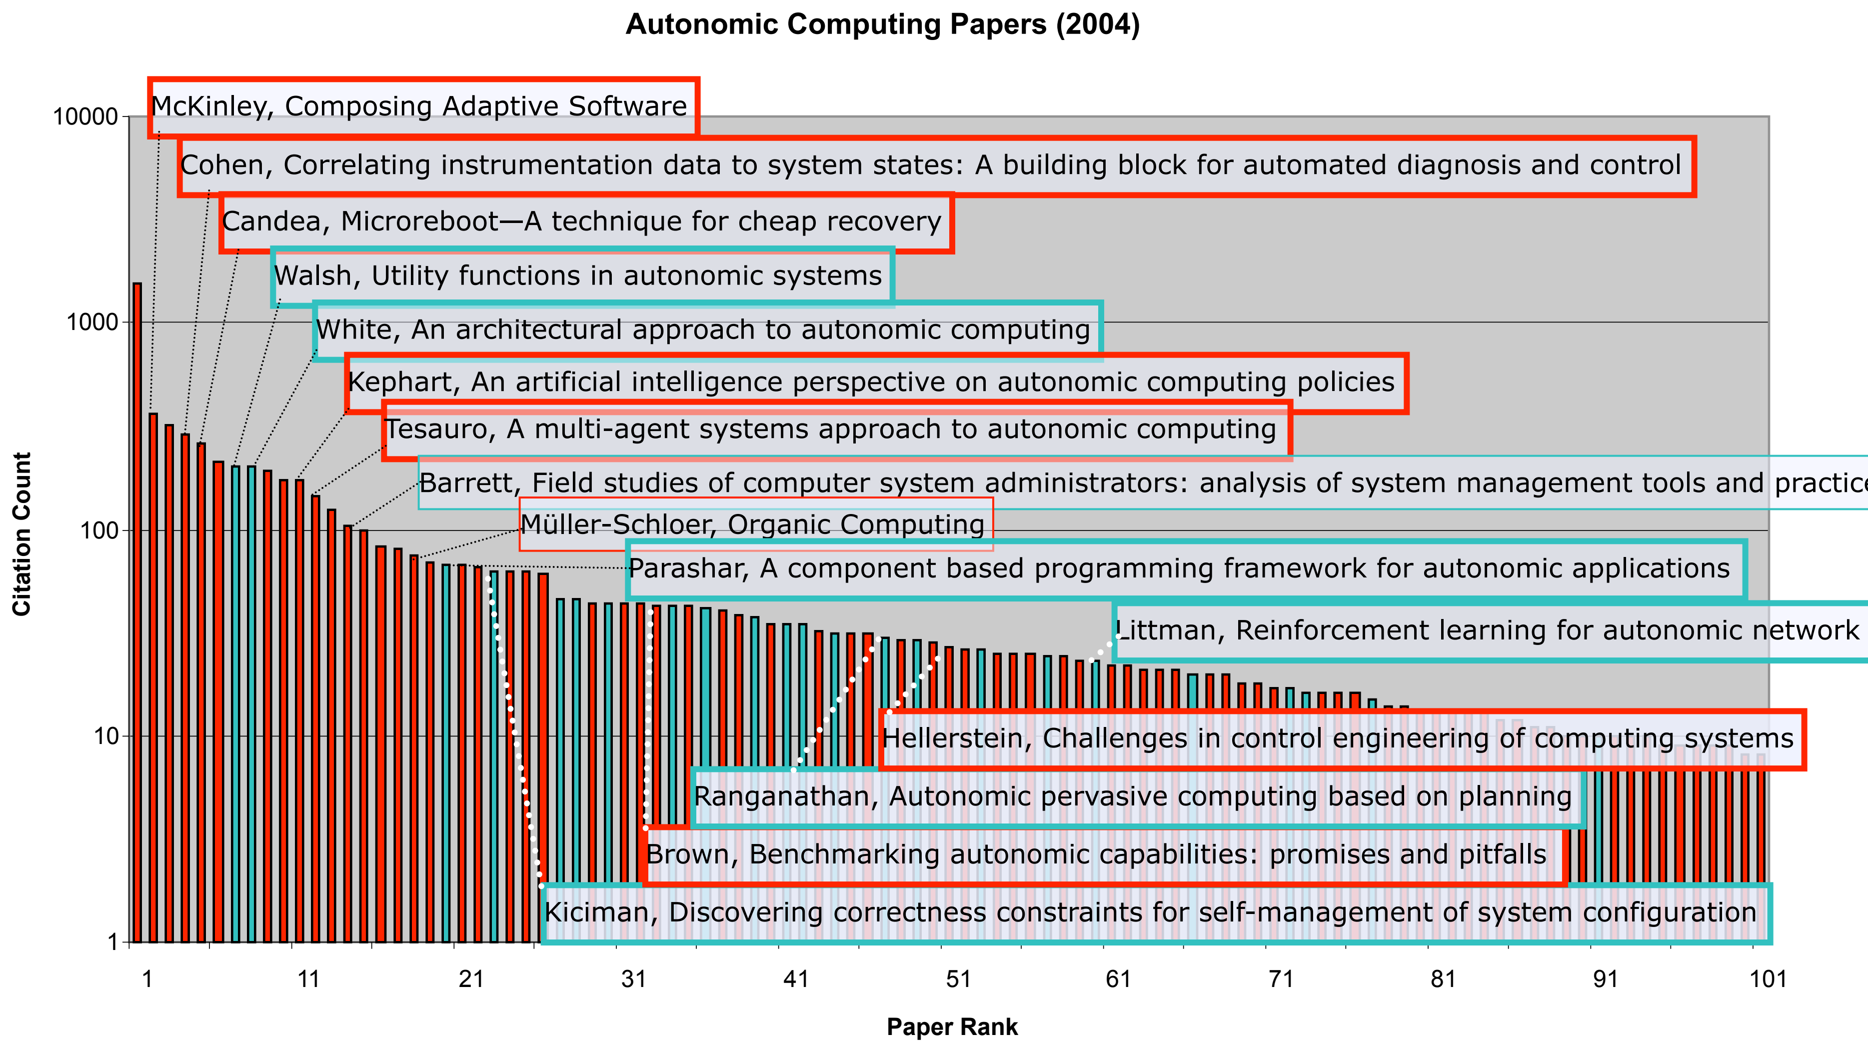

Wordle’s View of AC circa 2003

http://www.wordle.net/show/wrdl/3752036/Autonomic_Computing_paper_themes_2003

Wordle's View of AC circa 2004

http://www.wordle.net/show/wrdl/3752222/utonomic_Computing_paper_themes_2004

Analyzing AC trends over the first decade: a taxonomy

- Use original AC **vision** paper as basis for a taxonomy of papers

COVER FEATURE

The Vision of Autonomic Computing

Systems manage themselves according to an administrator's goals. New components integrate as effortlessly as a new cell establishes itself in the human body. These ideas are not science fiction, but elements of the grand challenge to create self-managing computing systems.

Jeffrey O. Kephart
David M. Chess
IBM Thomas J. Watson Research Center

In mid-October 2001, IBM released a manifesto observing that the main obstacle to further progress in the IT industry is a looming software complexity crisis.[1] The company cited applications and environments that weigh in at tens of millions of lines of code and require skilled IT professionals to install, configure, tune, and maintain.

The manifesto pointed out that the difficulty of managing today's computing systems goes well beyond the administration of individual software environments. The need to integrate several heterogeneous environments into corporate-wide computing systems, and to extend that beyond company boundaries into the Internet, introduces new levels of complexity. Computing systems' complexity appears to be approaching the limits of human capability, yet the march toward increased interconnectivity and integration rushes ahead unabated.

This march could turn the dream of pervasive computing—trillions of computing devices connected to the Internet—into a nightmare. Programming language innovations have extended the size and complexity of systems that architects can design, but relying solely on further innovations in programming methods will not get us through the present complexity crisis.

As systems become more interconnected and diverse, architects are less able to anticipate and design interactions among components, leaving such issues to be dealt with at runtime. Soon systems will become too massive and complex for even the most skilled system integrators to install, configure, optimize, maintain, and merge. And there will be no way to make timely, decisive responses to the rapid stream of changing and conflicting demands.

AUTONOMIC OPTION

The only option remaining is *autonomic computing*—computing systems that can manage themselves given high-level objectives from administrators. When IBM's senior vice president of research, Paul Horn, introduced this idea to the National Academy of Engineers at Harvard University in a March 2001 keynote address, he deliberately chose a term with a biological connotation. The autonomic nervous system governs our heart rate and body temperature, thus freeing our conscious brain from the burden of dealing with these and many other low-level, yet vital, functions.

The term autonomic computing is emblematic of a vast and somewhat tangled hierarchy of natural self-governing systems, many of which consist of myriad interacting, self-governing components that in turn comprise large numbers of interacting, autonomous, self-governing components at the next level down. The enormous range in scale, starting with molecular machines within cells and extending to human markets, societies, and the entire world socioeconomy, mirrors that of computing systems, which run from individual devices to the entire Internet. Thus, we believe it will be profitable to seek inspiration in the self-governance of social and economic systems as well as purely biological ones.

Clearly then, autonomic computing is a grand

0018-9162/03/$17.00 © 2003 IEEE Published by the IEEE Computer Society January 2003 41

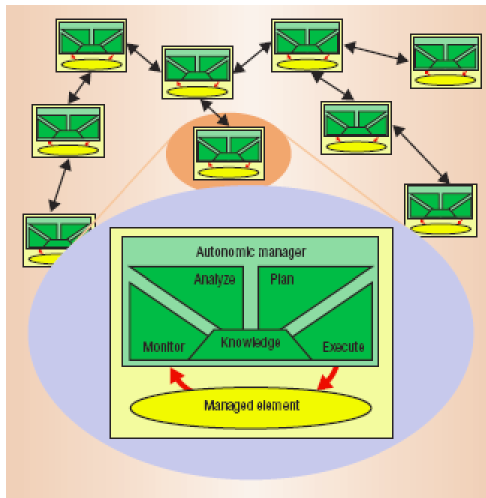

Figure 2. Structure of an autonomic element. Elements interact with other elements and with human programmers via their autonomic managers.

Architecture: Autonomic **elements** interact to produce system-level

self-configuration

self-healing

self-optimization, and

self-protection

- Engineering challenges
  - Element lifecycle; **software engineering**
  - Relationships: **services, standards, ontology,** negotiation
  - System: **policy**, **human interaction**, self-*
- Science challenges
  - **Machine learning, optimization & control**
  - Understanding and governing emergent **system behavior**

AC Paper Trends 2001-2010: Vision, Architecture, etc.

- Vision and architecture papers have settled to ~10%
  - **Appropriate**
- Human interaction study, never prevalent, became extinct in 2006
  - **BAD!**
- Study of system properties as a whole is rising steadily.
  - **GOOD!**

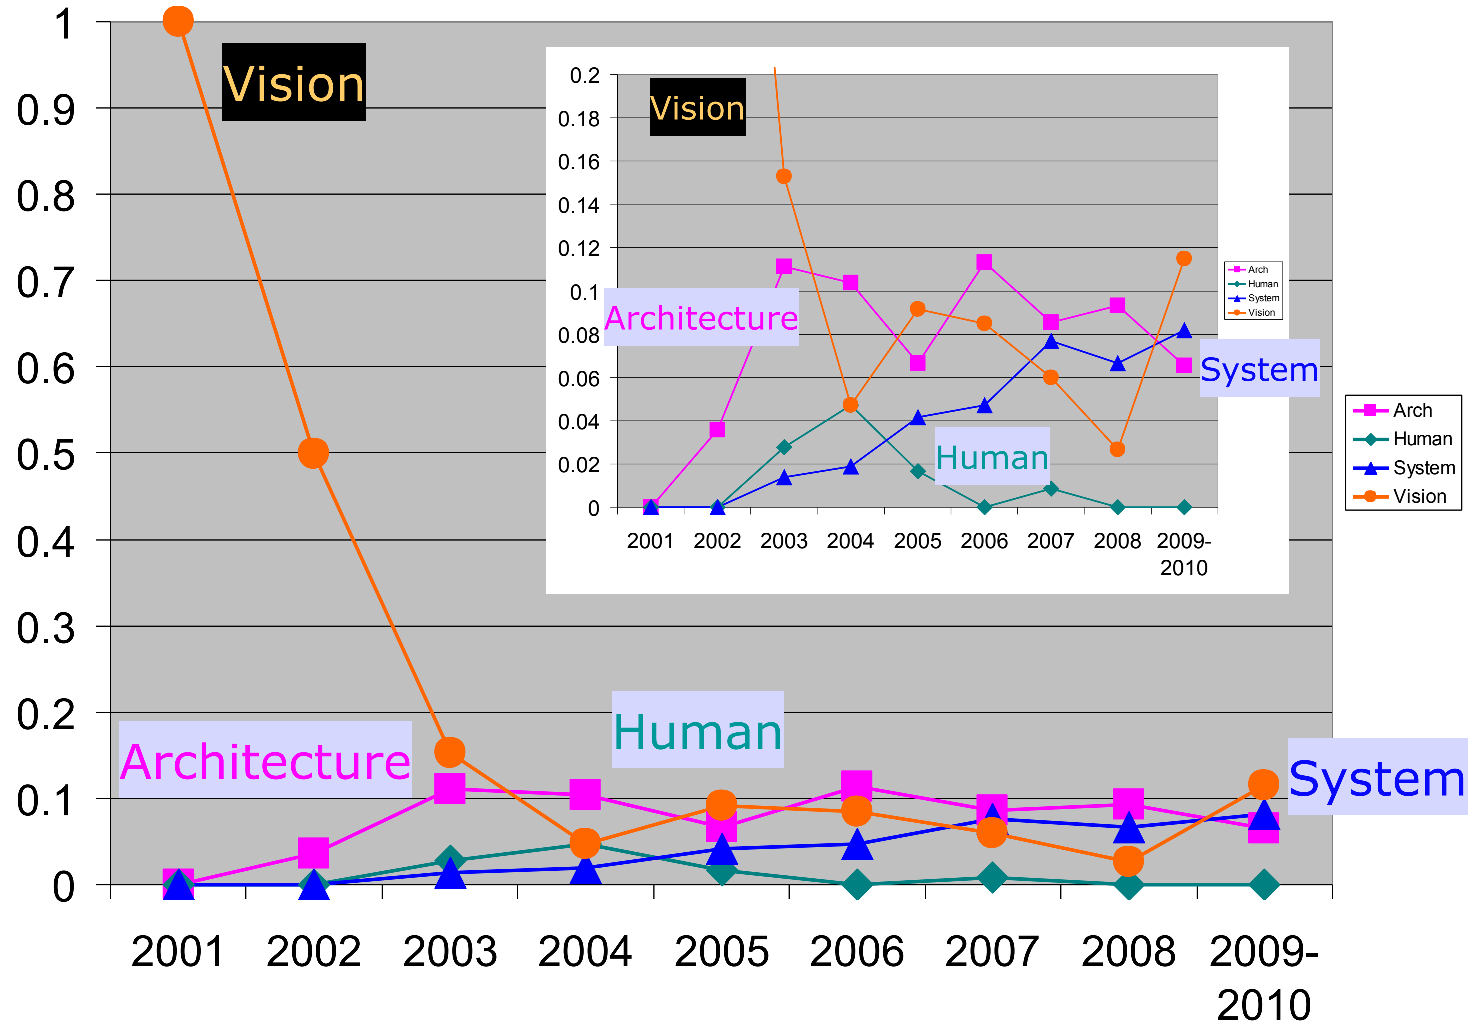

Vision

- Autonomic Computing
  - *Horn, Ganek, Kephart&Chess; Parashar&Hariri; Sterritt*
- Recovery-oriented computing
  - Don’t try to ensure 99.9999% up time for each component
  - Accept that faults are always going to happen; cope with them at system level
  - Micro-rebooting – minimize downtime by designing systems to be quickly rebootable at multiple levels
  - If it’s fast enough, occasional mistaken reboots are ok
  - *Patterson, Fox et al., UC Berkeley*
- Organic and bio-inspired computing
  - Use insights from biological systems to understand and exploit collective behavior
  - *KIT, BADS workshop; SASO; Richard Anthony*

No work on applying autonomic nervous system principles to autonomic computing !?!

A tale of two analogies

- Computer Viruses
  - Viruses replicate themselves by co-opting their host's resources
  - Analogies work on several levels
    - Macroscopic: epidemiology, evolutionary trends
    - Microscopic: immune system
  - Analogies help us
    - Understand the problem (science)
    - Ameliorate the problem (engineering)

- Autonomic Computing
  - Large-scale computing systems are becoming too complex for humans to manage. We need self-managing computing systems:
    - Self-configuring, Self-healing, Self-optimizing, Self-protecting
  - Autonomic nervous system automatically dilates pupils, increases respiratory rate, heart beat, etc.
  - Analogy to autonomic nervous system helps us describe the effect we want to achieve

Computer Viruses: Macroscopic Analogy

- Epidemiology
  - Individual = computer
  - *Social* network is important: can curtail spread relative to homogeneous mixing
- Evolutionary trends
  - Several great ages of computer viruses
    - File infectors
    - Boot infectors
    - Macro viruses
    - Worms
  - Heavily influenced by environment
  - Co-evolution with host (e.g. Microsoft Windows)
  - Overly virulent viruses are unsuccessful

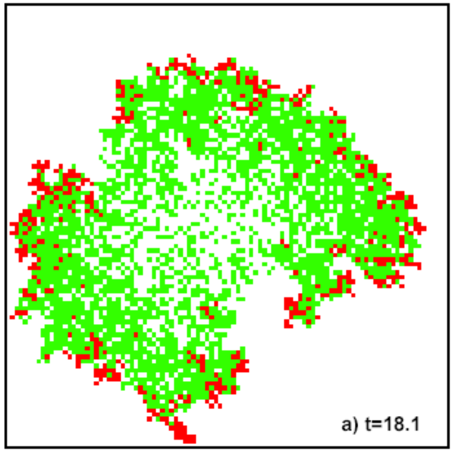

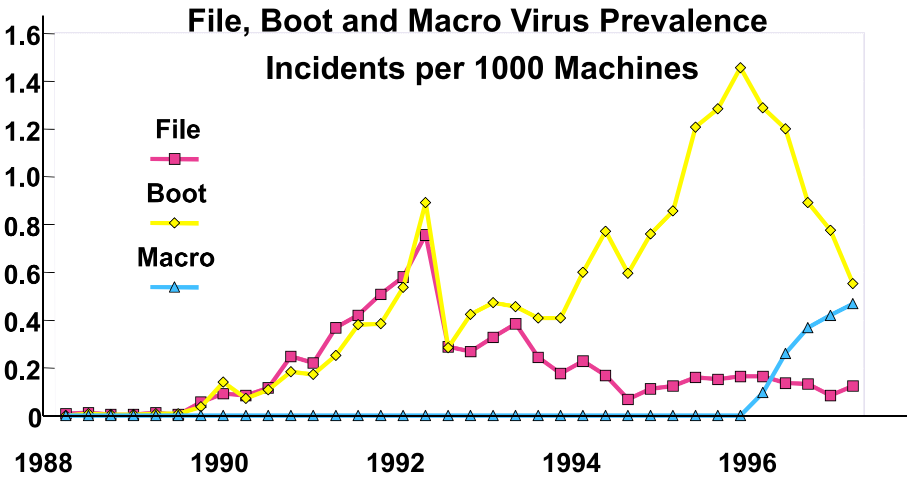

Computer Viruses: Immune System

- Recognize pathogen
  - **Unknown**: "Innate" immune system combines "Know thyself" with "Know thine enemy"
  - **Known**: Vertebrate immune system specifically detects tell-tale portions
- Eliminate it
  - **Biology**: Killer T cells destroy infected host cell to save host individual
  - **Computers**: Can often surgically remove virus from host cell
- Learn (if previously unknown)
  - **Biology**: Each individual does their own learning; vaccination helps
  - **Computers**: Learning can be shared

• Self/non-Self as proxy for Benign/Harmful
• Fight self-replication with self-replication

Immune System for Cyberspace

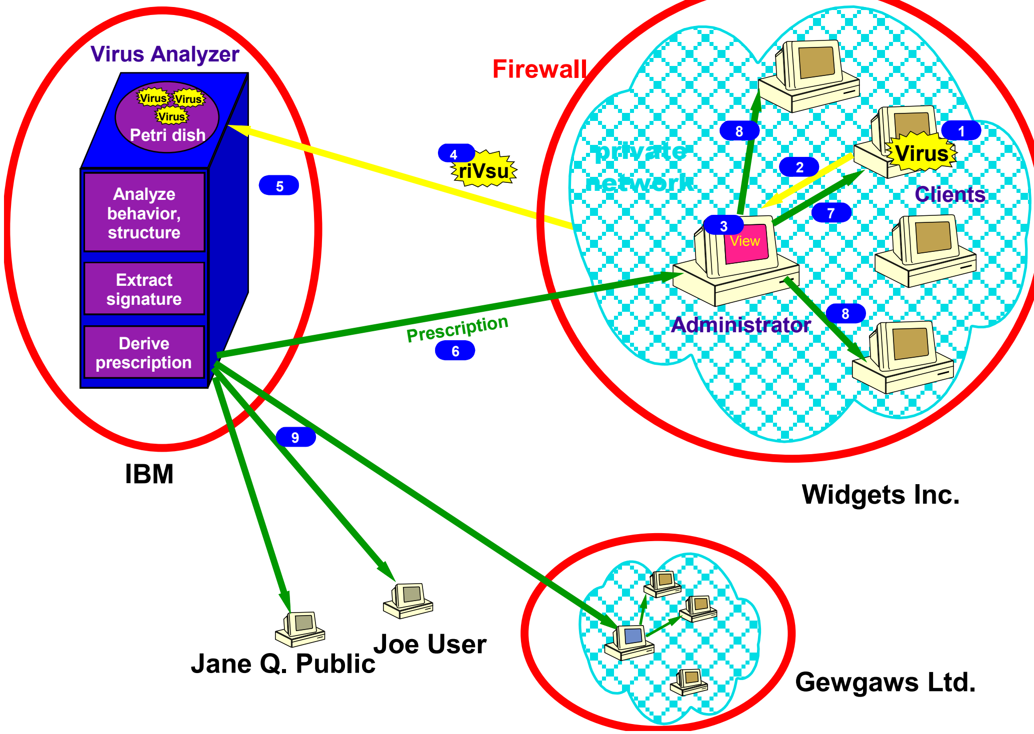

How can biological analogies be useful?

- **Marketing**: Describe the problem you're trying to solve; inspire others to solve it
- **Science**: Gain insight into the problem
  - Borrow mathematical techniques developed for related problems
  - Sometimes you end up contributing as much as you borrow (e.g. directed-graph epidemiology)
- **Engineering**: Derive techniques for solving the problem
  - Knowing that Nature has solved a related problem gives you hope
  - Even better, you may be able to adapt Nature's solutions to your problem
  - Even *wrong* theories about Nature's workings can be valuable!

Be open to what Nature has to teach you, but be judicious about what ideas you borrow.!

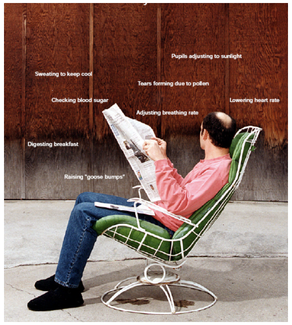

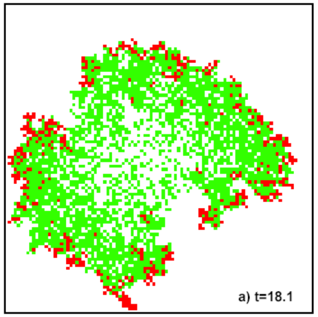

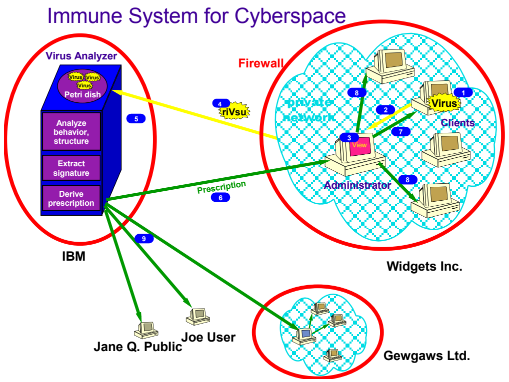

Kephart et al., *Fighting Computer Viruses,* Scientific American Nov 1997.

AC Paper Trends 2001-2010:
System architecture, policy, self-optimization

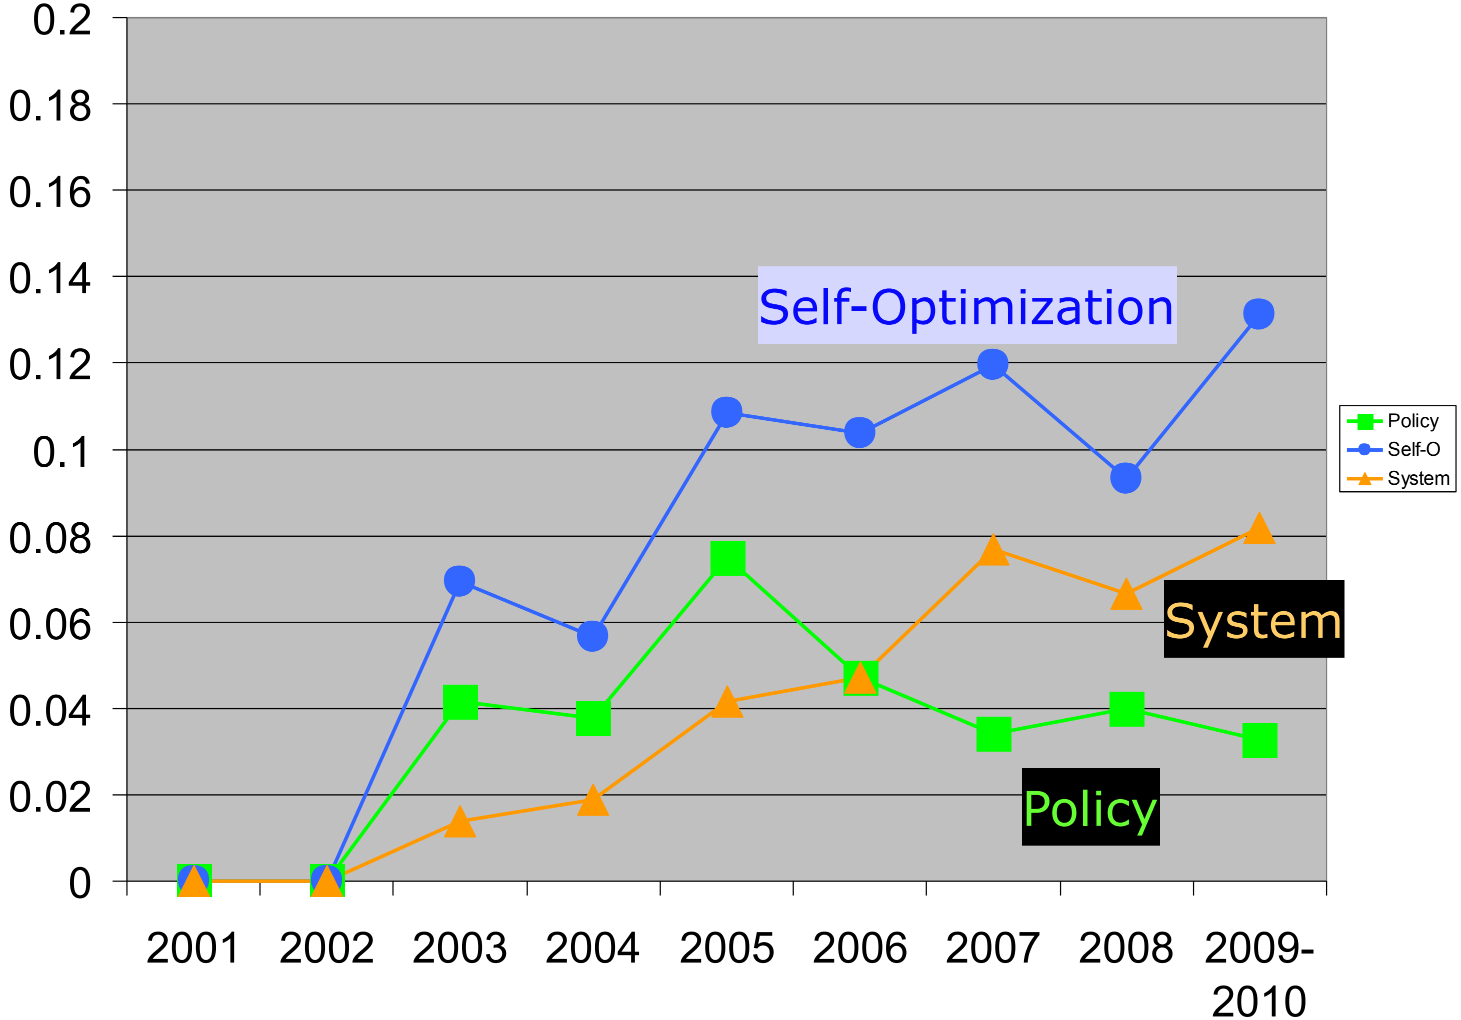

How to *represent* high-level policies?

Kephart and Walsh, Policy04

- Utility functions map any possible state of a system to a scalar value
- They can be obtained from
  - Service Level Agreement
  - preference elicitation
  - simple templates
- They are a very useful representation for high-level objectives
  - Value can be transformed and propagated among agents to guide system behavior

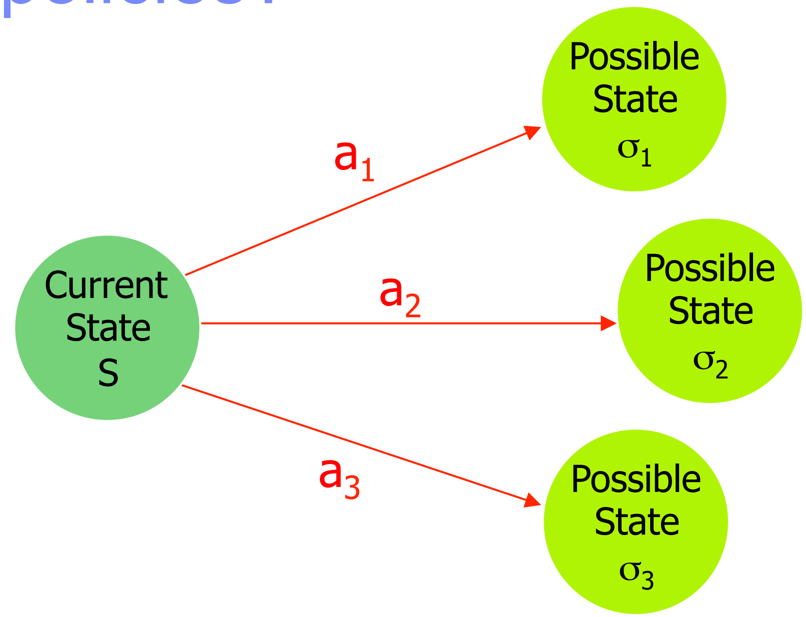

U(**RT**) =

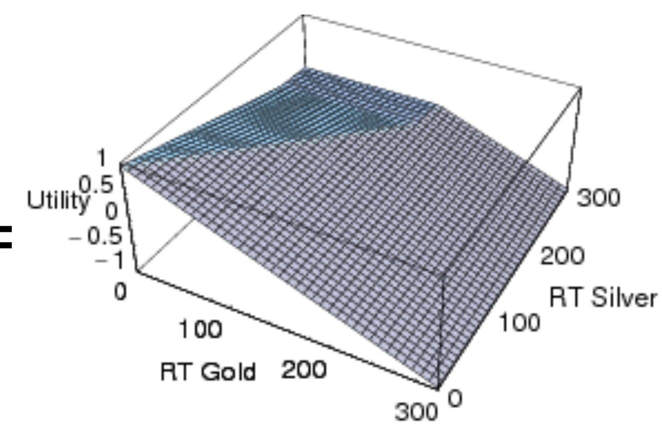

How to *manage* with high-level policies?

- **Elicit** utility function U(**S**) expressed in terms of service attributes **S**
- **Model** how each attribute S_i depends on controls **C** and observables **O**
  - Models expressed as **S**(**C**; **O**)
  - E.g., RT(routing weights, request rate)
  - Models from experiments, learning, theory
- **Transform** from service utility U to resource utility U' by substitution
  - U(**S**) = U(**S**(**C**; **O**)) = U'(**C**; **O**)
- **Optimize** resource utility. As observable **O** changes, set **C** to values that maximize U'(**C**; **O**)
  - **C*(O)** = argmax$_C$ U'(**C**; **O**)
  - U'***(O)** = U'(**C*(O)**; **O**)

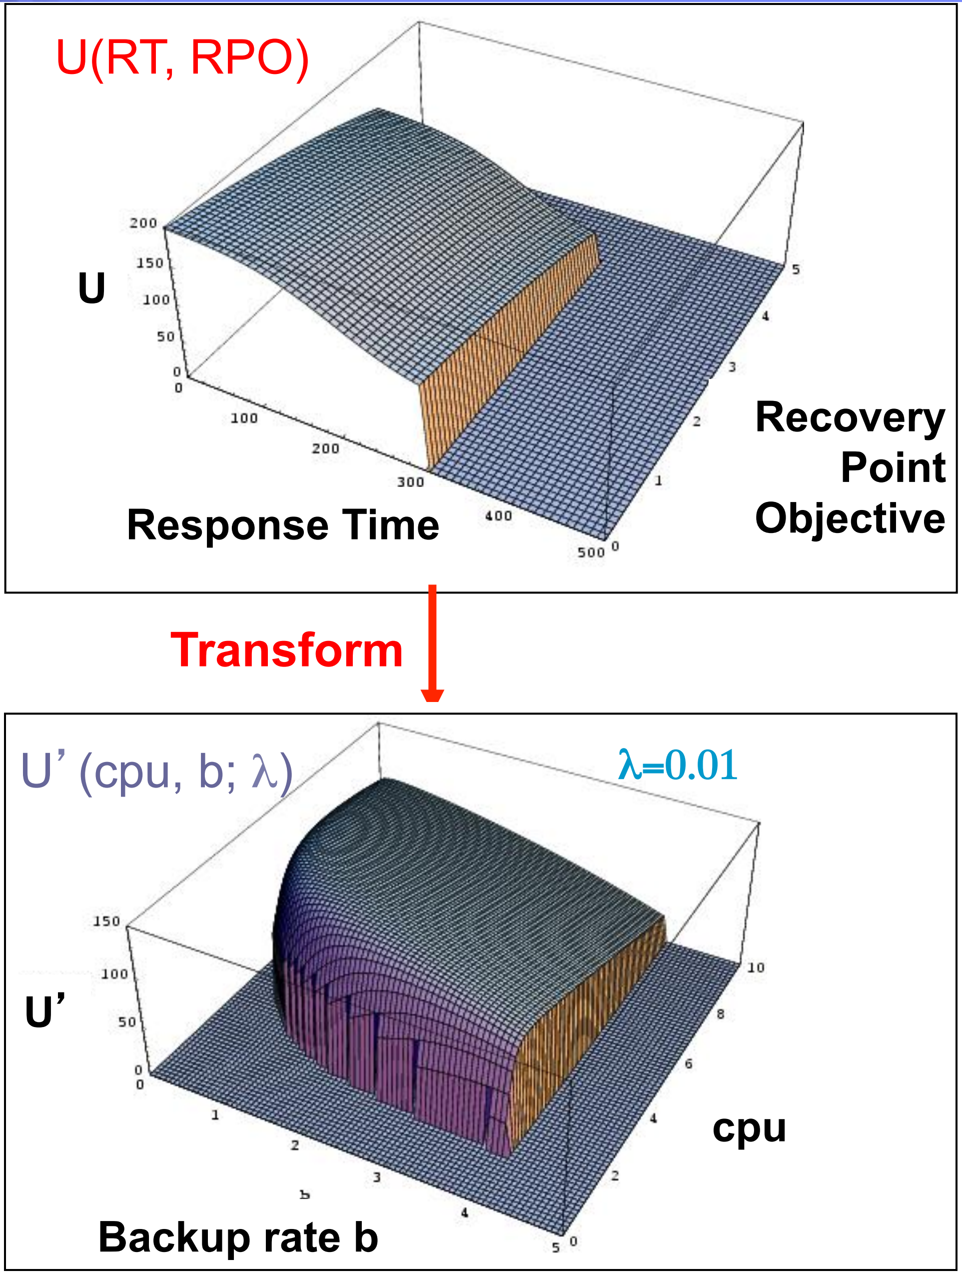

Unity Data Center Prototype: Experimental setup

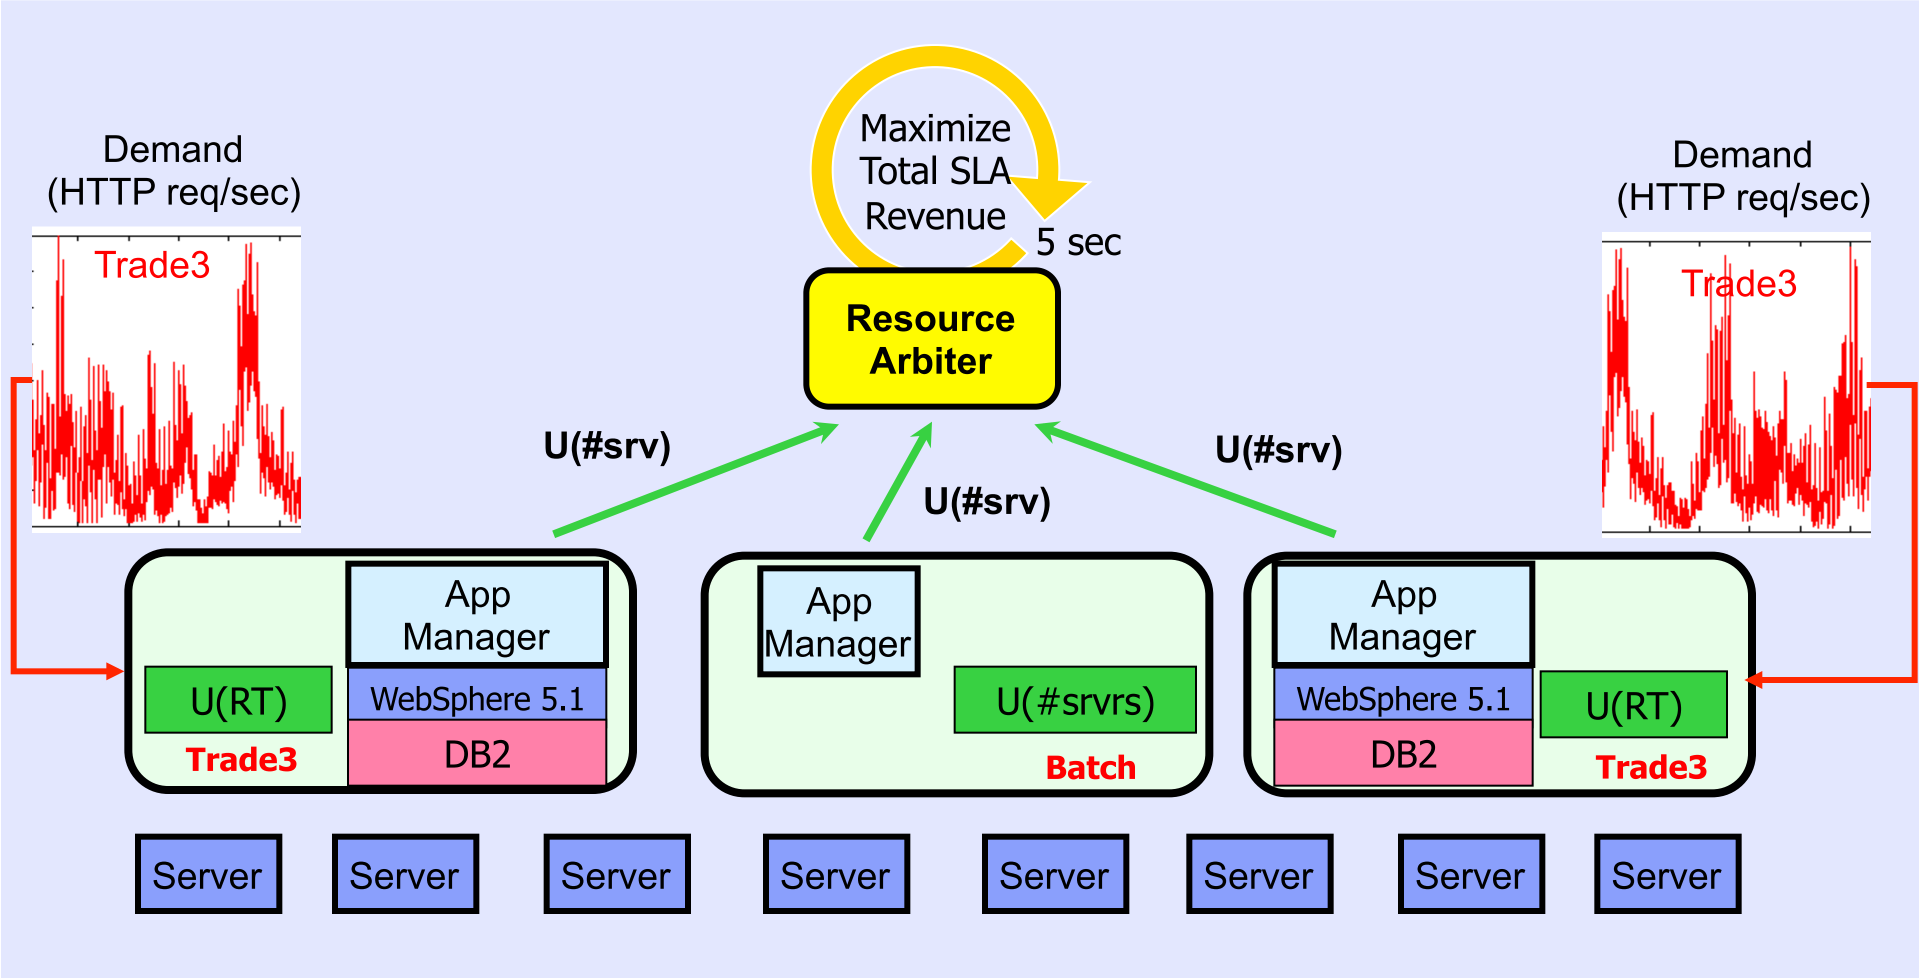

Chess, Segal, Whalley and White, *Unity: Experiences with a Prototype Autonomic Computing System*, ICAC 2004

How App Mgr computes its external resource utility

Alternative to generating full curve: utility elicitation

Patrascu, Boutilier et al. *New Approaches to Optimization and Utility Elicitation in Autonomic Computing*, AAAI 2005

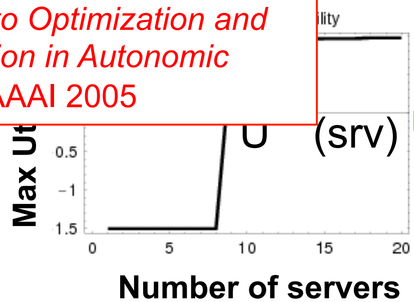

Resource Arbiter

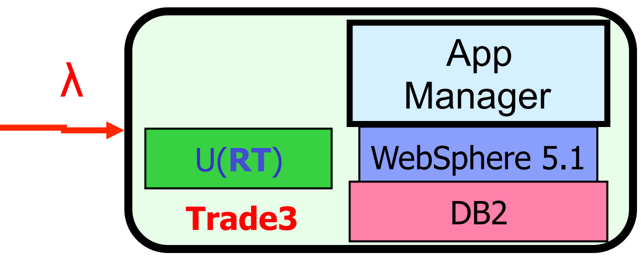

Elicit: U(**RT**) — Service-level utility

Model: U(**RT**(**C**; srv, λ))

My controls — Arbiter's controls — Observable

Transform: U' (**C**; srv, λ) = U(**RT**(**C**; srv, λ))

Internal resource-level utility

Optimize:

Optimal internal control settings: **C***(srv, λ) = argmax$_C$U' (**C**; srv, λ)

External resource-level utility: U'' (srv, λ) = U' (**C***(srv, λ); srv, λ)

Chess, Segal, Whalley and White, *Unity: Experiences with a Prototype Autonomic Computing System*, ICAC 2004

How the Arbiter determines optimal resource allocation

Decision problem:

Allocate resources

srv* = argmax$_{\mathbf{srv}}\Sigma U''_i(srv_i)$

Effectively maximizes $\Sigma U_i(\mathbf{S}_i)$

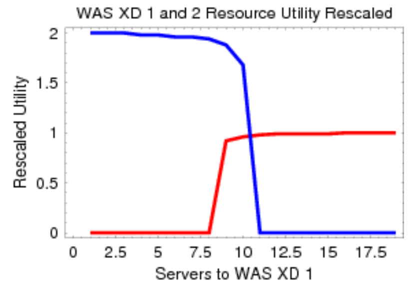

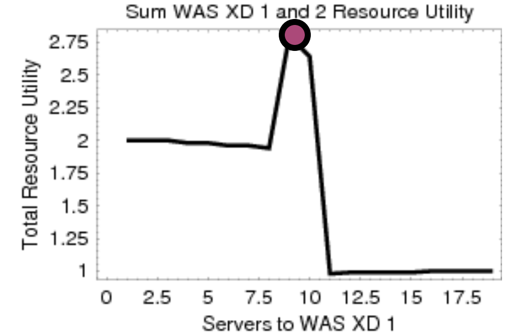

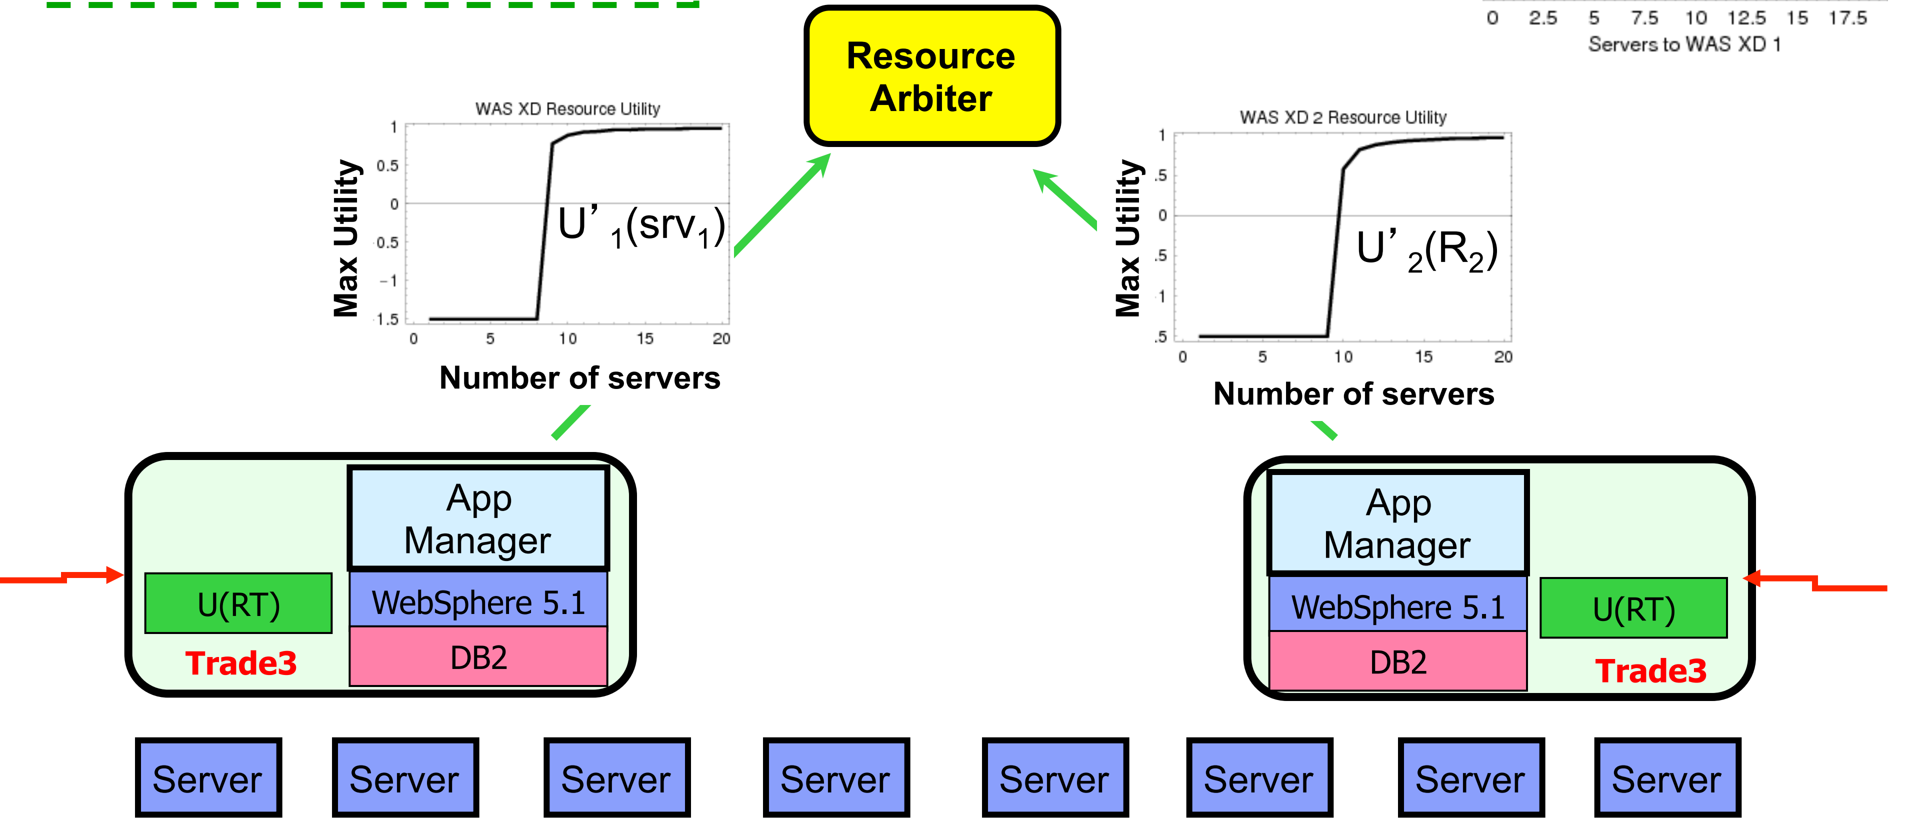

Policy and Systems: Status and Future

- We’ve made a good start on developing the utility-optimization design pattern
  - Theoretically well-grounded
  - Proven practical in several scenarios
- **But we need to push this work much further**
- **Establish that utility works on a grand scale in AC systems**
  - More than just a few agents and attributes
  - An economy, perhaps?
- **Utility elicitation from humans**
- **Need planning technologies to support goal policies**
  - More than just an engine
  - Tools for constructing planning domain descriptions

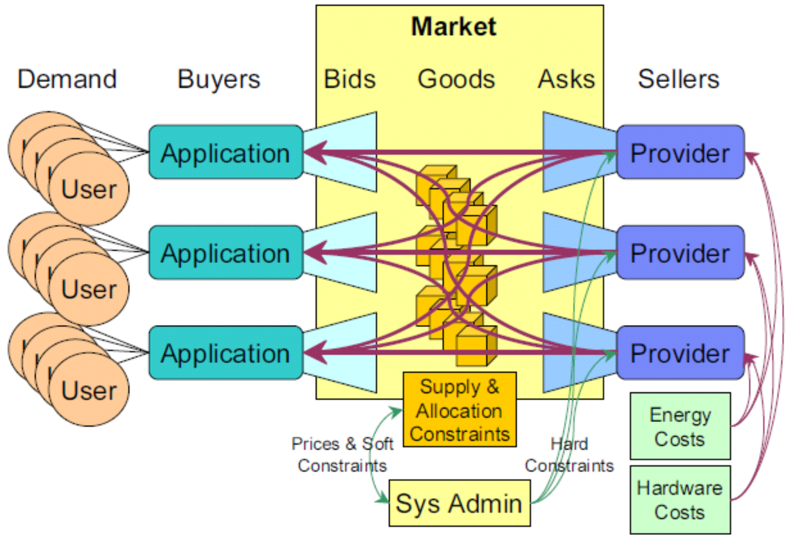

Figure 1: The Data Center Market Model

Lubin, Kephart, Das and Parkes. *Expressive Power-Based Resource Allocation for Data Centers*. **IJCAI 2009.** (Exploring market-based resource allocation for data centers.)

AC Paper Trends 2001-2010: Self-*, Benchmarks

- David Patterson warned us that we needed benchmarks for self-{C,H,P} in order to drive work in the field
- It appears that he was right
- **We need to revive the benchmark work**
- **We need more work on self-{C,H,P}**

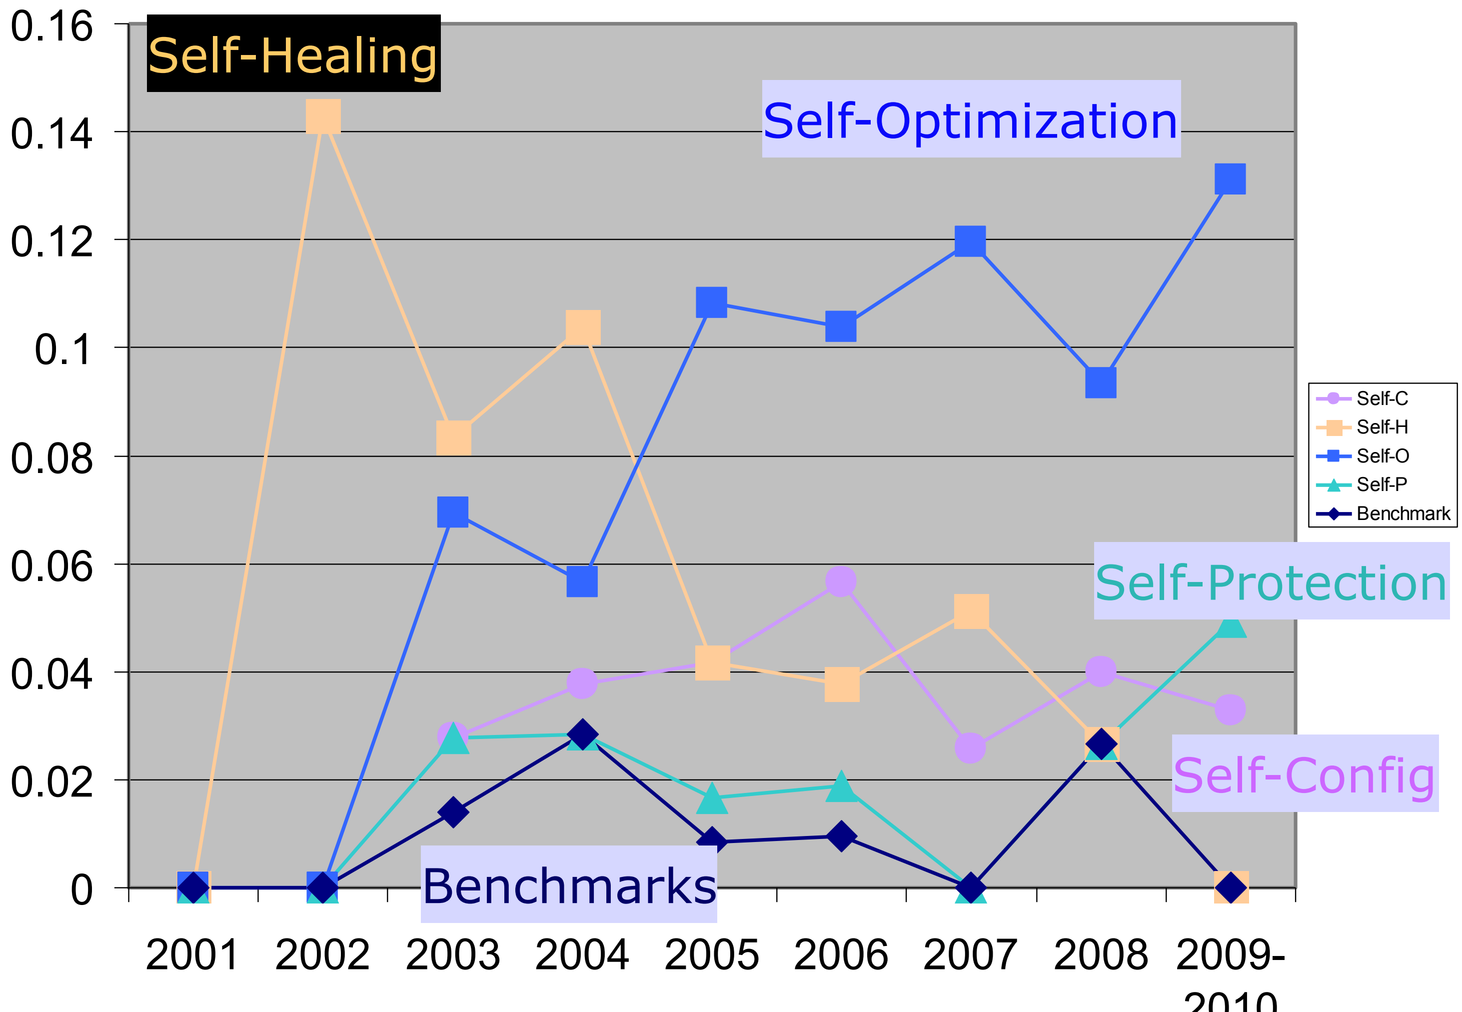

Benchmarks

- David Patterson noted that
  - Benchmarks drive innovation, but practically all are performance-related
  - Innovations pertaining to self-{C, H, P} require appropriate metrics

- Brown et al. developed benchmarks for configuration and healing
  - **Brown & Keller**. *A model of configuration complexity and its application to a change management system*. IM 2005.

- McCann et al. recommended metrics for adaptivity, robustness, autonomy, sensitivity, stabilization; suggested adapting existing benchmarks
  - **McCann & Huebscher**. *Evaluation issues in autonomic computing*. GCC 2004

- Other papers include
  - **Consens et al.** *Goals and benchmarks for autonomic configuration recommenders*. 2005
  - **K. Kanoun**. *Dependability benchmarking for computer systems*. 2008.

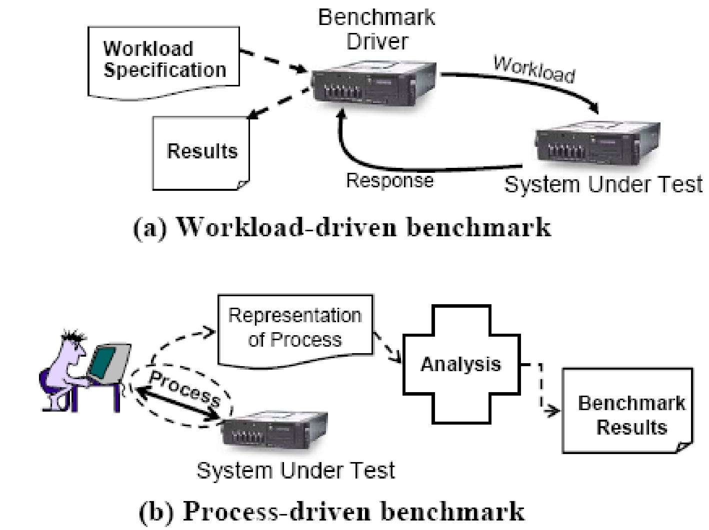

Figure 2. Workload- vs. Process-driven benchmarks.

Brown & Hellerstein. *Benchmarking Autonomic Capabilities: Promises and Pitfalls*. ICAC04

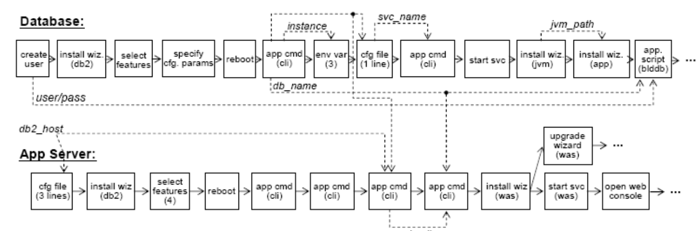

Figure 1. Representation of configuration process for Java-based web application. The figure shows the first few steps in the configuration process for a simple enterprise Java web application. Each box represents a configuration step, and dotted arrows represent information that the operator must pass between tasks.

After a promising start, work on autonomic computing benchmarks appears to have (mostly) stagnated.

AC Paper Trends 2001 – 2010: Relationships: WebServices/Grid

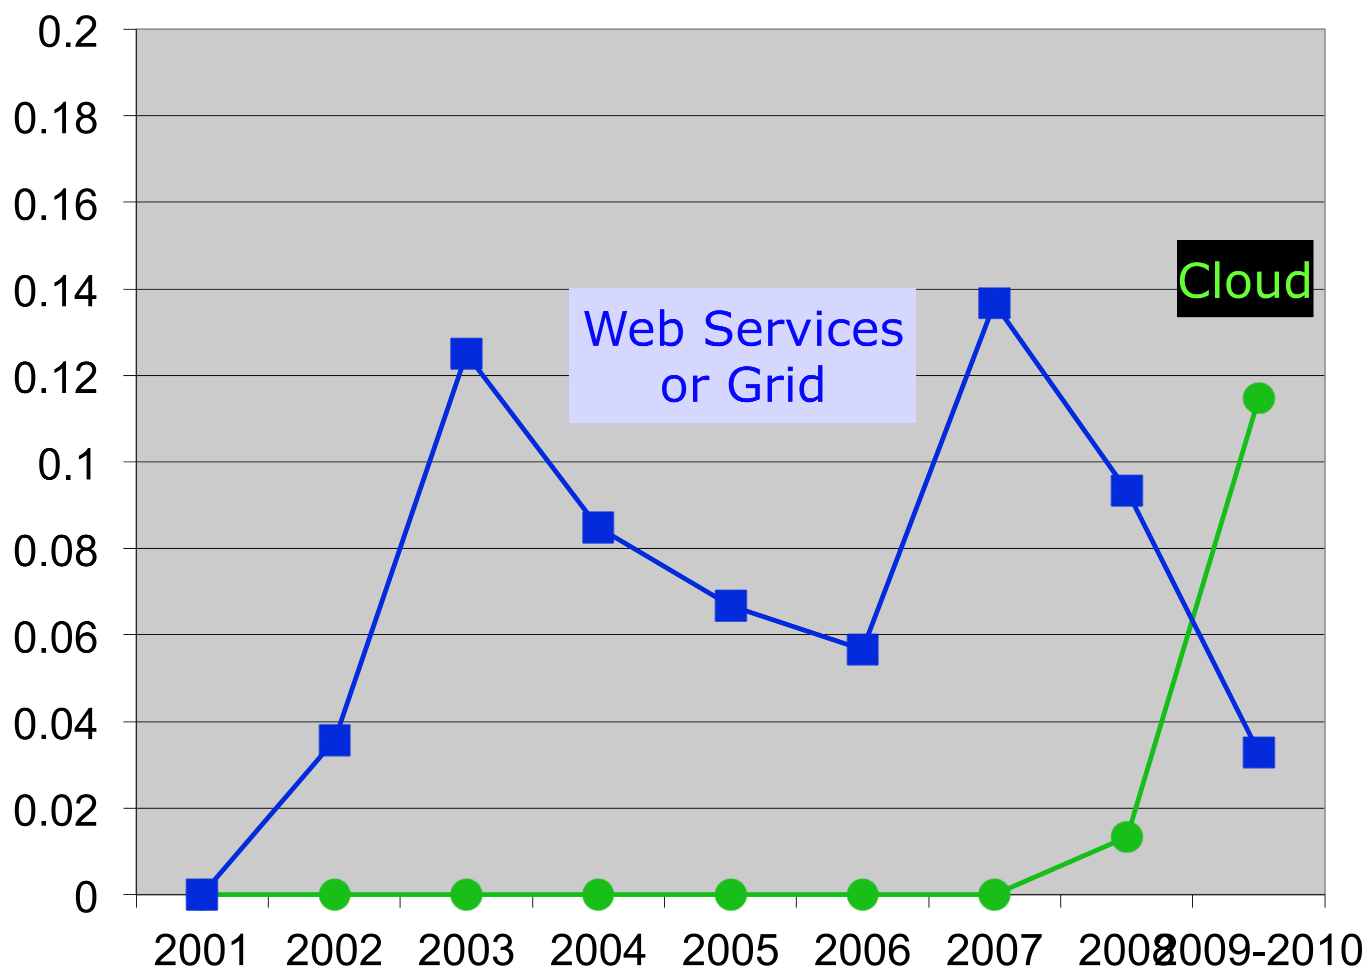

Relationships: WebServices/Grid

- Agent communication standards likely to derive from services

- Foresee convergence of autonomic computing, web services, grid interfaces

Web-Scale Workflow

Editor: M. Brian Blake • *mb7@georgetown.edu*

Agents and Service-Oriented Computing for Autonomic Computing

A Research Agenda

Frances M.T. Brazier • *Vrije Universiteit Amsterdam*
Jeffrey O. Kephart • *IBM T.J. Watson Research Center*
H. Van Dyke Parunak • *Tech Team Government Solutions*
Michael N. Huhns • *University of South Carolina*

Autonomic computing is the solution proposed to cope with the complexity of today's computing environments. Self-management, an important element of autonomic computing, is also characteristic of single and multiagent systems, as well as systems based on service-oriented architectures. Combining these technologies can be profitable for all — in particular, for the development of autonomic computing systems.

In recent years, computing environments' complexity has begun to grow beyond the limits of what human system administrators can manage. This increasing complexity has three sources. First, individual components of computing systems, such as workload managers and database management systems, are becoming more difficult to configure, manage, and maintain as each release includes ever more system, which regulates heart and respiratory rates, digestion, and other bodily functions, freeing the conscious brain to focus on higher-level goals. Similarly, autonomic computing systems are expected to free system administrators to focus on higher-level goals. Autonomic computing systems can perform the following functions without human intervention:

Article resulted from brainstorming session at Agents for Autonomic Computing workshop, ICAC 2008

Brazier, Kephart, Parunak, and Huhns, Internet Computing, June 2009

AC Paper Trends 2001-2010: AI Technologies

- Relatively small but sustained effort on AI technologies for autonomic systems

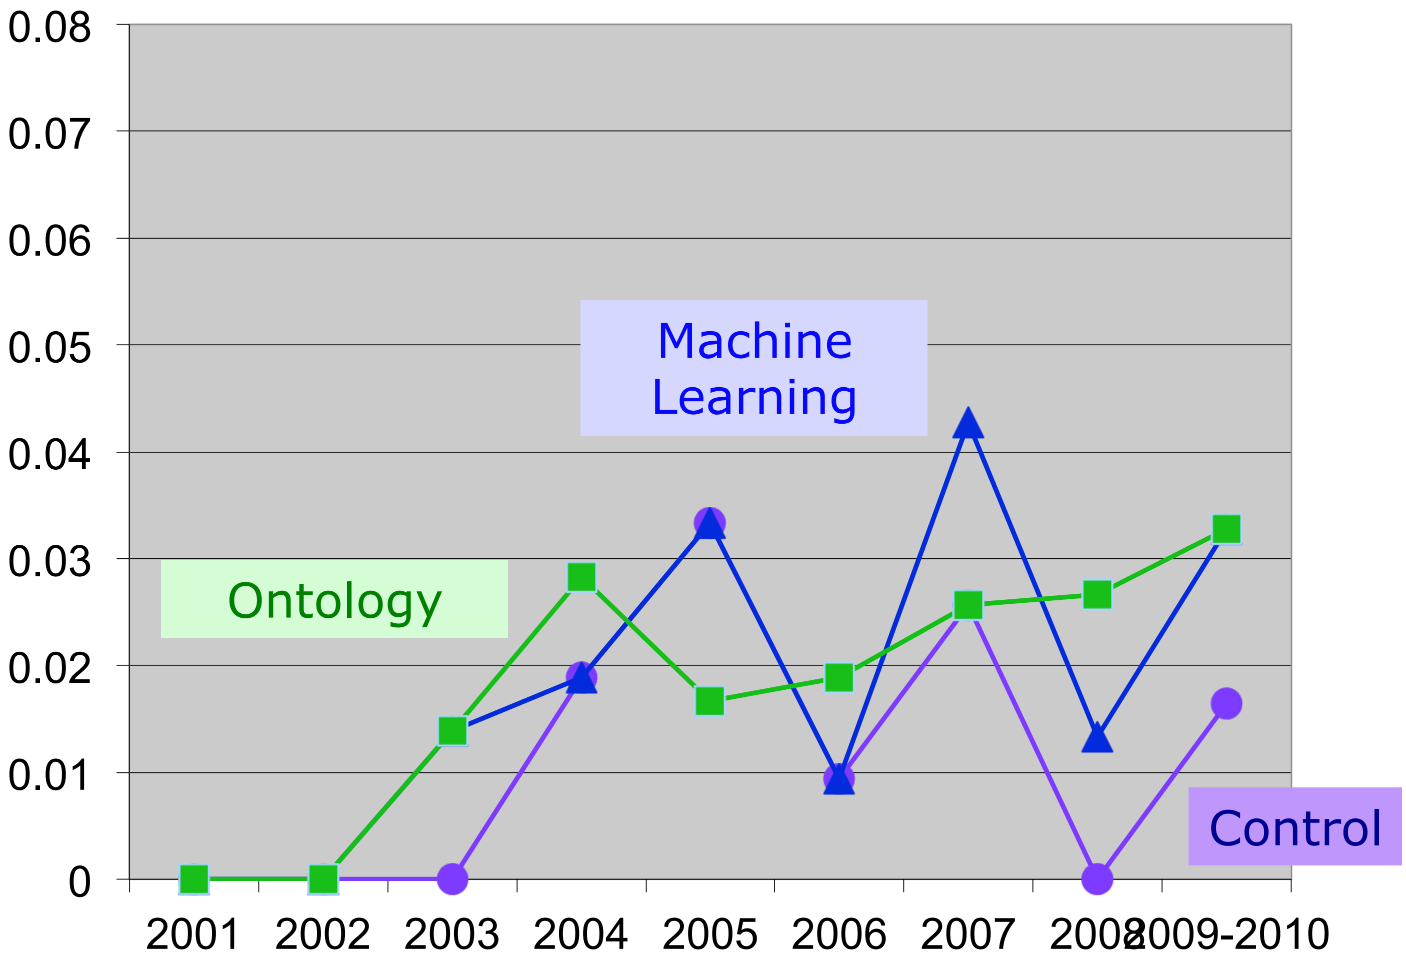

Machine Learning

- Good progress on learning models and policies
  - I. Cohen et al. *Correlating ….* OSDI04.
  - G. Jiang et al. *Discovering likely invariants of distributed transaction systems for autonomic system management.* ICAC06
  - G. Tesauro et al. *A hybrid ….* ICAC06

- **We still need to tackle multi-agent learning**
  - Several interacting learners
  - What are good learning algorithms for cooperative, competitive systems?
    - Stability and sensitivity characteristics
    - What is sensitivity to perturbations?
  - Opportunities for layered learning

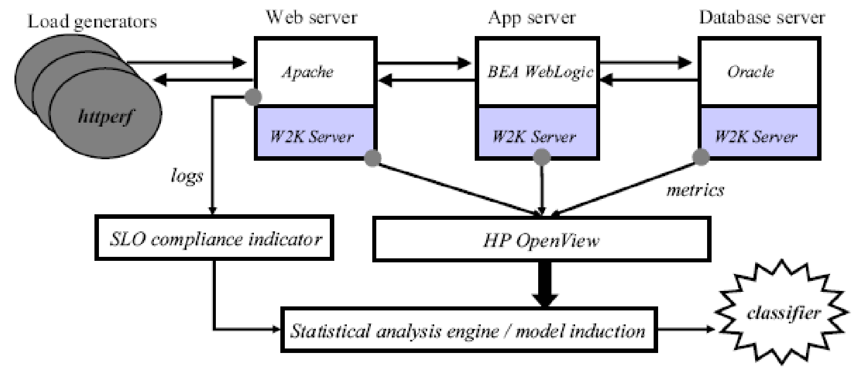

I. Cohen et al. *Correlating instrumentation data to system states: A building block for automated diagnosis and control*. OSDI04

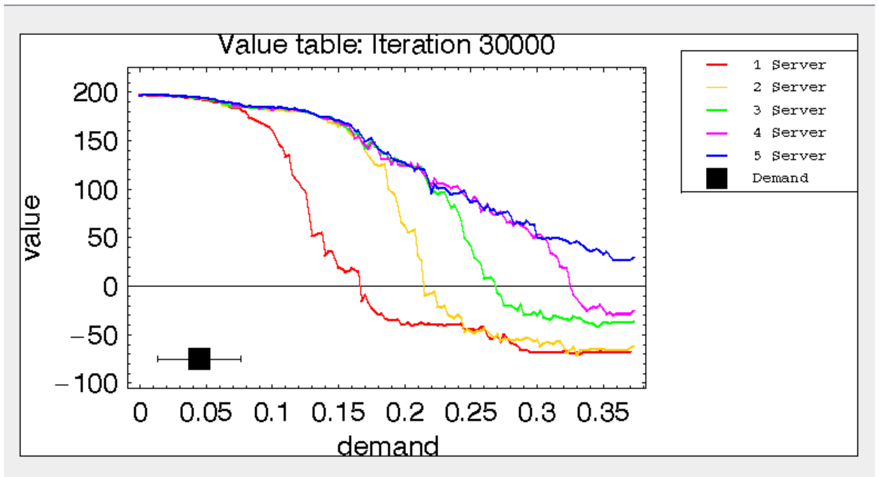

G. Tesauro et al. *A hybrid reinforcement learning approach to autonomic resource allocation*. ICAC06

Feedback control

- Good progress on applying feedback control to individual autonomic elements
  - Middleware including databases, application servers: Book and multiple papers by J. Hellerstein et al.
- Good progress on applying feedback control to clusters of compute resources , power and performance
  - Kusic et al.

- **We still need to understand and control the behavior of multiple interacting feedback loops**
  - Hierarchical and distributed
  - Some good early thoughts in P. Ranganathan. *No Power Struggles: Coordinated Multi-Level Power Management for the Data Center.*
  - ASPLOS08

- **Generally, we still need to understand emergent behavior much better**

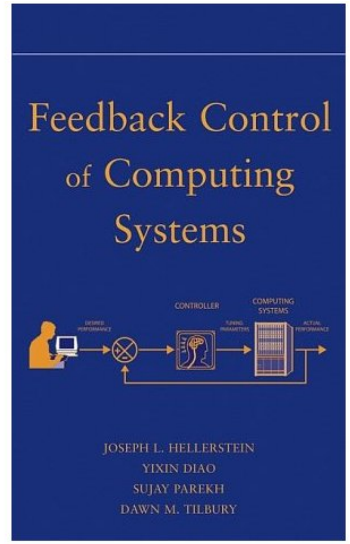

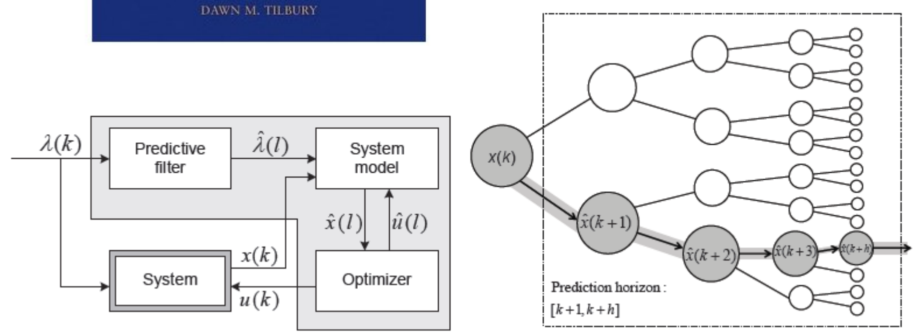

(a) (b)

D. Kusic et al. *Power and Performance Management of Virtualized Computing Environments via Lookahead Control*. ICAC07

Unanticipated trends, and their impact on AC

- Data centers and energy management
  - The physical infrastructure is complex, and needs to be autonomic, too!
  - New attributes: Energy and temperature
- Cloud Computing
  - Some vendors (Google, Amazon, Facebook) can get away with highly standardized and homogeneous environments
  - Outsourcing to the cloud means that fewer companies manage IT themselves
  - Perhaps it places a greater burden on cloud providers to implement AC
    - Lower costs
    - Places premium on easy configurability
    - Outages are more embarrassing and costly

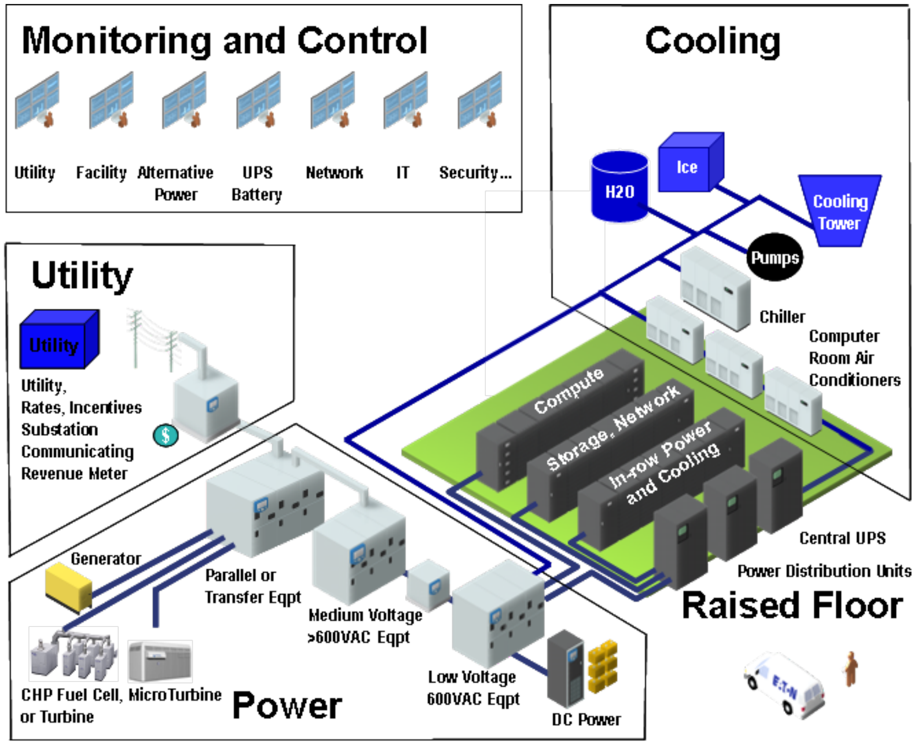

Conclusions

- Autonomic Computing is alive and well
  - Thousands of papers, 129 of them with at least 50 citations
  - Hundreds of conferences and workshops that touch on AC
- We have had a busy and fruitful first decade
  - Good balance of vision, architecture, new techniques, apps
  - We haven’t exploited the autonomic nervous system analogy – but that’s OK
  - Not much new theory, or system-level prototypes that address multiple facets
- Several serious engineering and science challenges remain
  - We need more work at the system level
    - Multi-agent learning, interacting feedback loops
    - Understanding/harnessing emergent behavior
    - Economic models should be pursued seriously
    - We need to build and experiment with prototypes and testbeds
  - We need to revive our development of benchmarks for Self-{C,H,P}
  - We need more focus on human interaction with autonomic systems; elicitation

Backup

ICAC 2004-2011

Year	Location	General Chairs	Program Chairs
2004	New York, NY	Jeff Kephart (IBM Research) Manish Parashar (Rutgers)	Rajarshi Das (IBM) Vaidy Sunderam (Emory)
2005	Seattle, WA	Jeff Kephart (IBM Research) Manish Parashar (Rutgers)	Karsten Schwan (Ga Tech) Yi-Min Wang (Microsoft Research)
2006	Dublin, Ireland	Karsten Schwan (Ga Tech) Yi-Min Wang (Microsoft Research)	Mazin Yousif (Intel) Omer Rana (Cardiff U.)
2007	Jacksonville, FL	Mazin Yousif (Intel) Omer Rana (Cardiff U.)	Jose Fortes (U. Florida) Kumar Goswami (HP Labs)
2008	Chicago, IL	Jose Fortes (U. Florida) Kumar Goswami (HP Labs)	John Strassner (Motorola) Simon Dobson (UCD Dublin)
2009	Barcelona	John Strassner (Motorola) Simon Dobson (UCD Dublin)	Manish Parashar (Rutgers) Onn Shehory (IBM Research)
2010	Reston, VA	Manish Parashar (Rutgers)	Renato Figueiredo (U. Florida) Emre Kiciman (Microsoft Research)
2011	Karlsruhe, Germany	Hartmut Schmeck (Karlsruhe, GE)	Joseph Hellerstein (Google) Tarek Abdelzaher (UIUC)

ICAC Steering Committee (2011)

- Jeffrey Kephart, IBM Research (Co-chair)
- Salim Hariri, University of Arizona (Co-Chair)
- Manish Parashar, Rutgers University
- Karsten Schwan, Georgia Tech
- Emre Kiciman, Microsoft Research
- Renato Figueiredo, University of Florida
- John Wilkes, Google

Autonomic computing paper impact
(from Harzing's Publish or Perish)

Papers:	998	Cites/paper:	30.06	h-index:	75	AWCR:	4494.42
Citations:	29999	Cites/author:	N/A	g-index:	140	AW-index:	67.04
Years:	11	Papers/author:	N/A	hc-index:	51	AWCRpA:	1881.16
Cites/year:	2727.18	Authors/paper:	2.78	hI-index:	25.45	e-index:	101.85
		hI,norm:	44	hm-index:	50.95		

Query date: 6/12/2011

Hirsch a=5.33, m=6.82

Contemporary ac=6.91

Cites/paper 30.06/11.0/2 (mean/median/mode)

Authors/paper 2.78/3.0/3 (mean/median/mode)

192 paper(s) with 1 author(s)

244 paper(s) with 2 author(s)

248 paper(s) with 3 author(s)

227 paper(s) with 4 author(s)

76 paper(s) with 5 author(s)

11 paper(s) with 6 author(s)

ic [New Search]

raphy "Ronald L. Rivest" asiacrypt (note the use of double quotes)

ype: General | Keywords | Related Conferences | Authors

Sort first 20 query mat

Prev | Next | Last

switch to textual view

21 to 40 c

Important Dates (Sort by: Deadline, Notification, Start, End)

Name	"Rating"	Match	Location
wac	3.0	1.00	
autonomics	3.0	0.99	
dasc	3.0	0.99	
atc	3.0	0.99	
soas	3.0	0.99	
ifip6	3.0	0.99	
bertinoro	3.0	0.99	
upp	3.0	0.99	
wowmom	3.0	0.98	Italy, Lucca
dsom	3.0	0.97	
ecbs	3.0	0.97	USA, Las Vegas
selfman	3.0	0.97	
eurongi	3.0	0.94	
im	3.0	0.94	
wrac	3.0	0.93	
ipom	3.0	0.92	
policy	3.0	0.92	
dexaw	3.0	0.92	
CLADE	3.0	0.89	USA, Chicago, IL

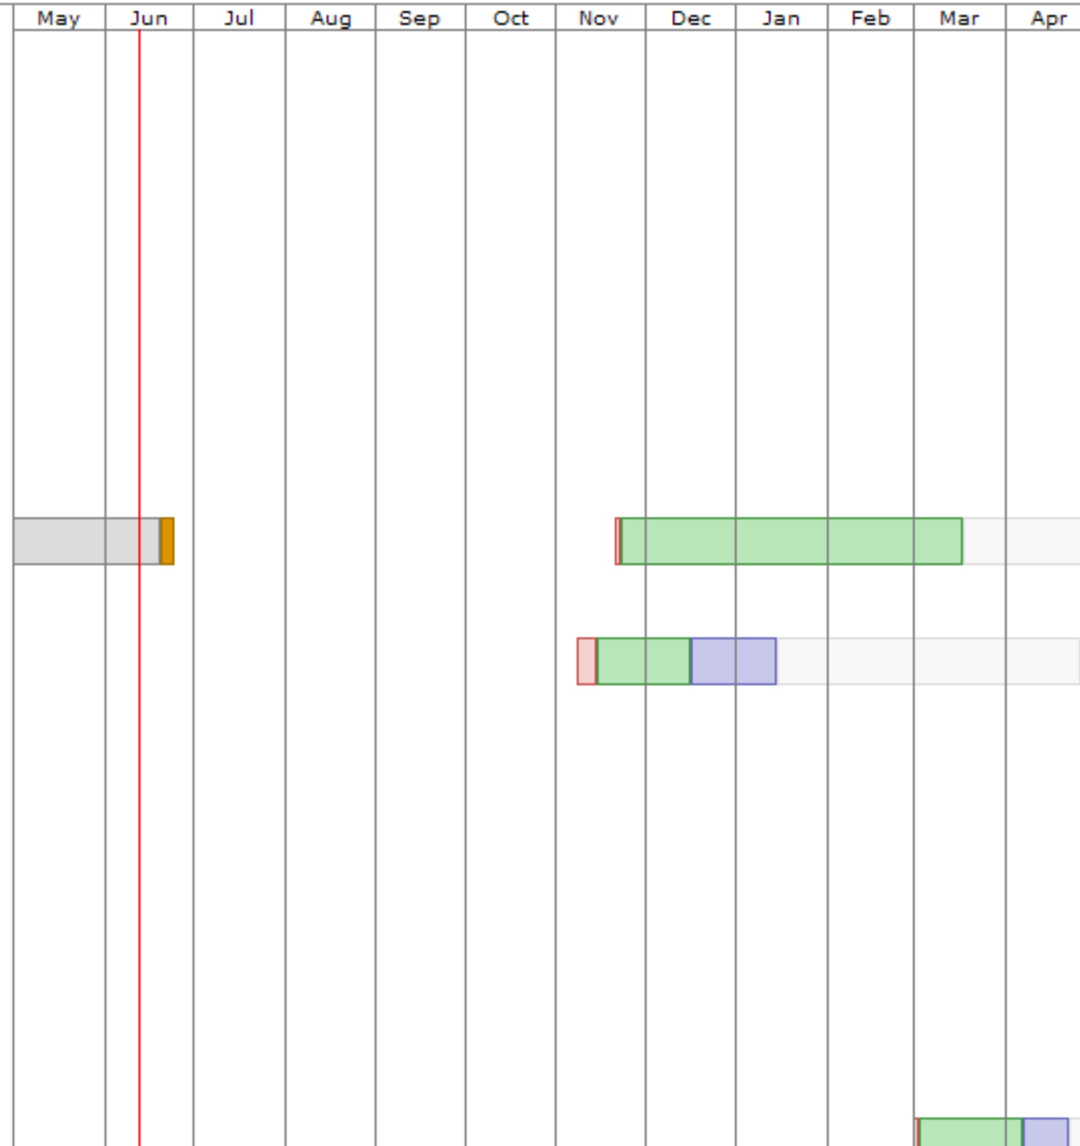

M

Wordle™ Home Create Gallery Credits News Forum FAQ Advanced

Edit Language Font Layout Color

Open in Window Print... Randomize Save to public gallery...

Not working?