



# Grid models: challenges and examples

**Cécile Germain-Renaud**

**Laboratoire de Recherche en Informatique**

**Université Paris Sud, CNRS, INRIA**

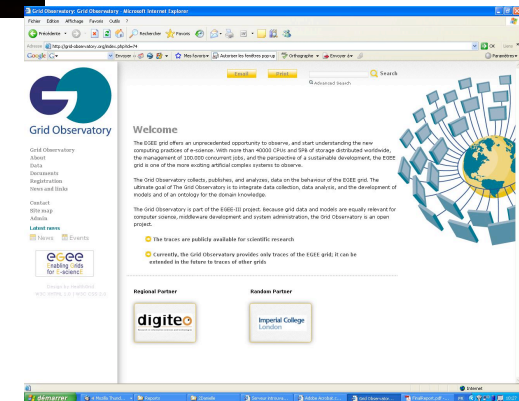
Results from the GO collaboration

# Outline

- ✓ Once upon a time...
- ✓ Scientific challenges
- ✓ Towards realistic behavioural models
- ✓ Conclusion and questions

# The Grid Observatory


- Digital curation of the behavioural data of the EGI grid: observe and publish
- Complex systems description
- Models, optimization, Autonomics
- This talk is about some lessons learned



# Remember tomorrow

*A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities*

Ian Foster, 1998



# Still a long way to go



How do we configure our grids?

James Casey talk @EGEE09

*Consistent*

*Dependable*

## Amazon's Cloud Crash Disaster Permanently Destroyed Many Customers' Data

Henry Blodget | Apr. 28, 2011, 7:10 AM | 77,816 | 76

Share

Tweet 1,297

Like 1K

Email

AA

71

In addition to taking down the sites of dozens of high-profile companies for hours (and, in some cases, days), Amazon's huge EC2 cloud services crash permanently destroyed some data.

The data loss was apparently small relative to the total data stored, but anyone who runs a web site can immediately understand how terrifying a prospect any data loss is.

(And a small loss on a percentage basis for Amazon, obviously, could be catastrophic for some companies).




Um...

# Remember tomorrow (5 years later)

*Systems manage themselves according to an administrator's goals. New components integrate as effortlessly as a new cell establishes itself in the human body*

J. Kephart and David M. Chess, the Autonomic Computing Manifesto, 2003



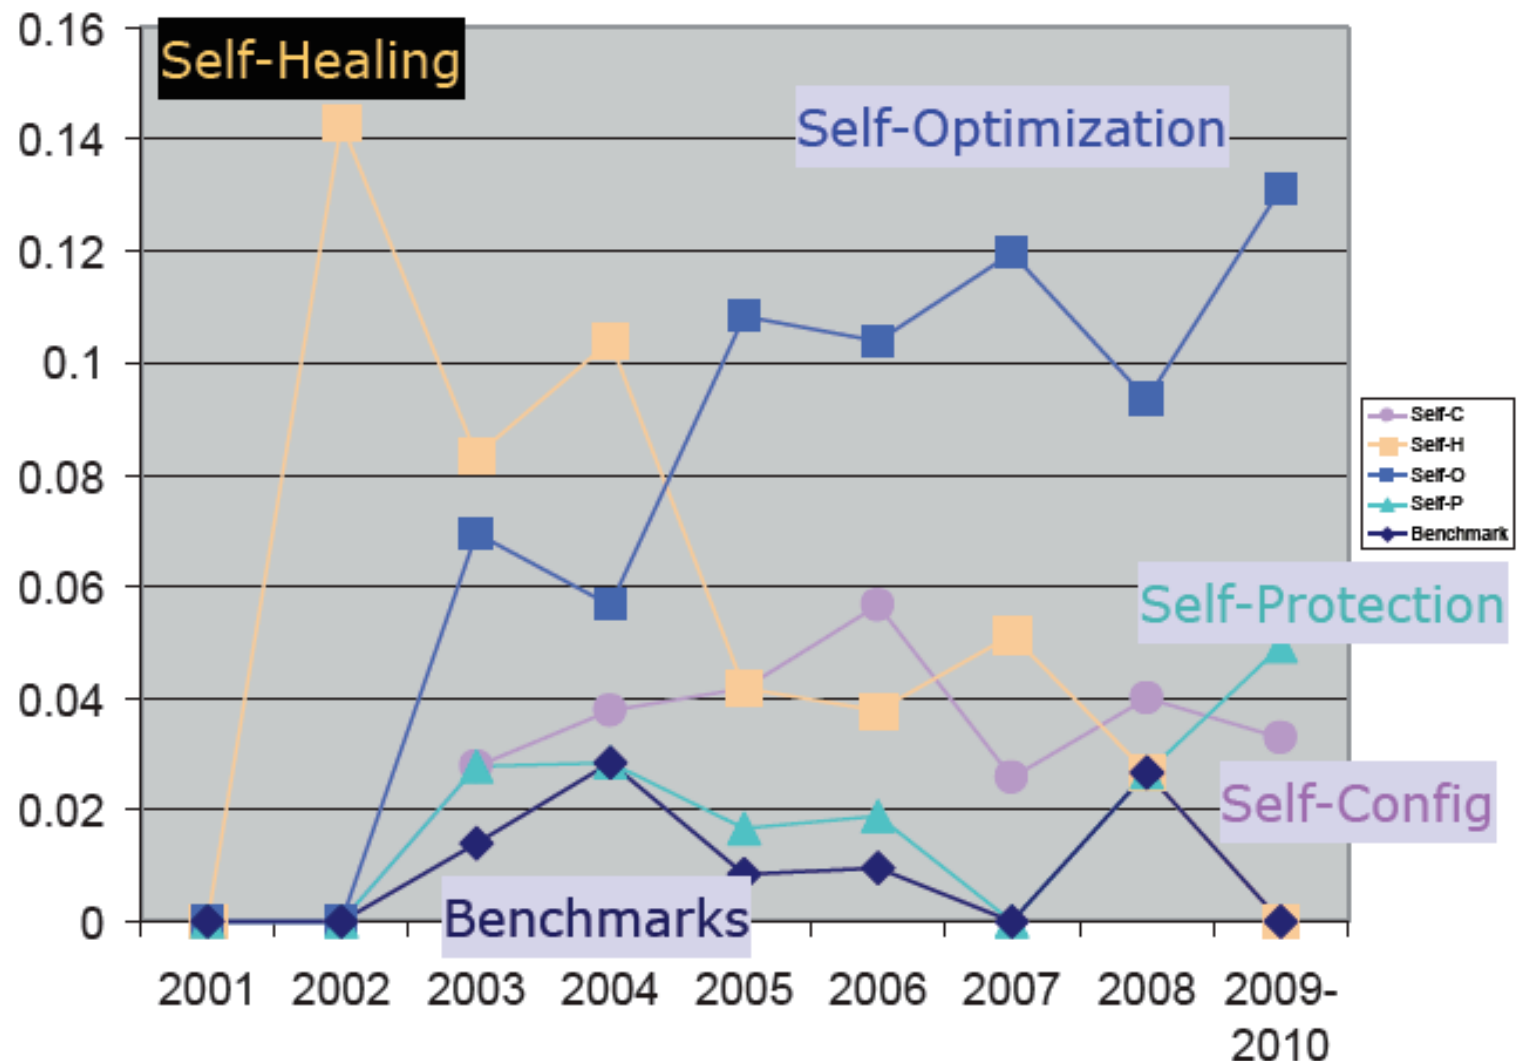
# The challenge

- How much of Computer Science, and specifically Autonomics research went to the real world?
- We need to show that the research has **verifiable** and positive impact on production systems




# 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}**





# A fairy tale



- 1979: Patterson went for a sabbatical at DEC. He worked on the implementation of the VAX machine on a single chip. This turned out to be impossible because of the extreme complexity of the microcode.
- With J. Hennessy he built the RISC I and **demonstrated** it better than its competitor from experience
- And a new era in Computer architecture began

# The challenge

- We need to show that the research has **verifiable** and positive impact on production systems
- Not Yet Implemented. Why? Sociology e.g. research vs production is a real, but minor issue
- This talk is about:

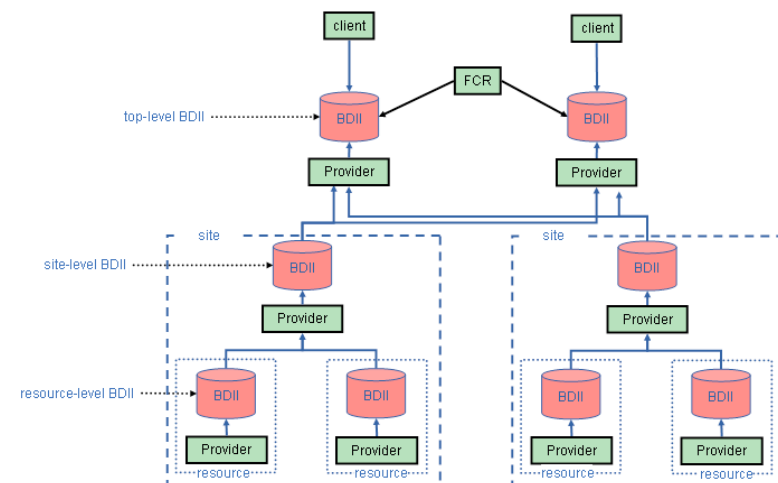
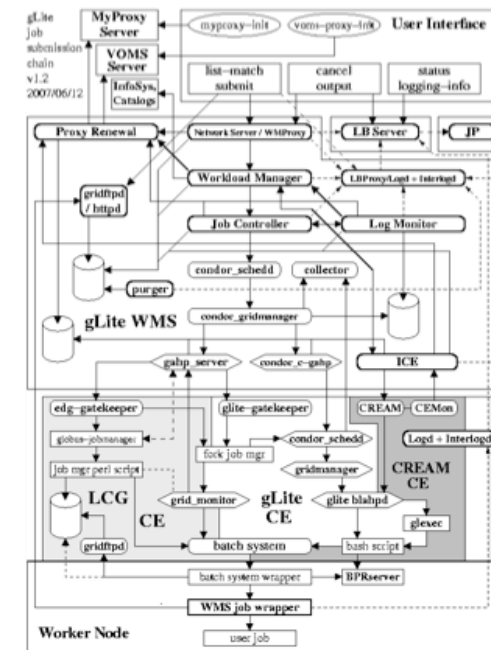
Demonstrating on complex systems  
raises serious scientific issues

# Outline


- ✓ Once upon a time...
- ✓ Scientific challenges
- ✓ Towards realistic behavioural models
- ✓ Conclusion and questions

# Issue I: Uncertainty

- As a dynamic(al) system
  - Entities change behavior as an effect of unexpected feedbacks, emergent behavior
  - Organized self-criticality, minority games,...
- Lack of complete and common knowledge – Information uncertainty
  - Monitoring is distributed too
  - Resolution and calibration



# Unexpected behavior



Predicting the response time

# Issue II: Fundamentals in statistics

- Statistical significance
- Which metrics?
- Are our systems stationary?



# Statistical significance

Extreme values may dominate the statistics

	Atlas		Biomed	
	ART	ERT	ART	ERT
Mean	1.33E3	2.74E4	3.01E2	2.66E2
Median	11	1	11	1
Std	1.09E4	7.41E4	4.33E3	5.99E3
RMSE	7.94E4		7.21E3	
<i>q</i> <sub>90%</sub>	1.35E2	1.16E5	25	4
Over. fraction	22%		3%	
Over. median	9.34E4		228	
Under. fraction	77%		96%	
Under. median	9.01E0		9.00E0	

# Statistical significance

Can we predict anything?  
 Maybe as difficult as earthquakes and markets

	Atlas		Biomed	
	ART	ERT	ART	ERT
Mean	1.33E3	2.74E4	3.01E2	2.66E2
Median	11	1	11	1
Std	1.09E4	7.41E4	4.33E3	5.99E3
RMSE	7.94E4		7.21E3	
$q_{90\%}$	1.35E2	1.16E5	25	4
Over. fraction	22%		3%	
Over. median	9.34E4		228	
Under. fraction	77%		96%	
Under. median	9.01E0		9.00E0	

# Metrics

Root Mean Squared Error is inadequate

	Atlas		Biomed	
	ART	ERT	ART	ERT
Mean	1.33E3	2.74E4	3.01E2	2.66E2
Median	11	1	11	1
Std	1.09E4	7.41E4	4.33E3	5.99E3
RMSE	7.94E4		7.21E3	
<i>q</i> <sub>90%</sub>	1.35E2	1.16E5	25	4
Over. fraction	22%		3%	
Over. median	9.34E4		228	
Under. fraction	77%		96%	
Under. median	9.01E0		9.00E0	

# Metrics

Should make sense for the end user

	Atlas		Biomed	
	ART	ERT	ART	ERT
Mean	1.33E3	2.74E4	3.01E2	2.66E2
Median	11	1	11	1
Std	1.09E4	7.41E4	4.33E3	5.99E3
RMSE	7.94E4		7.21E3	
<i>q</i> <sub>90%</sub>	1.35E2	1.16E5	25	4
Over. fraction	22%		3%	
Over. median	9.34E4		228	
Under. fraction	77%		96%	
Under. median	9.01E0		9.00E0	

# A few keywords



Heavy tail



Self-similarity

Long range dependency

Heteroskedasticity



Harold Edwin Hurst  
1880-1978



# Do naïve statistics make sense?

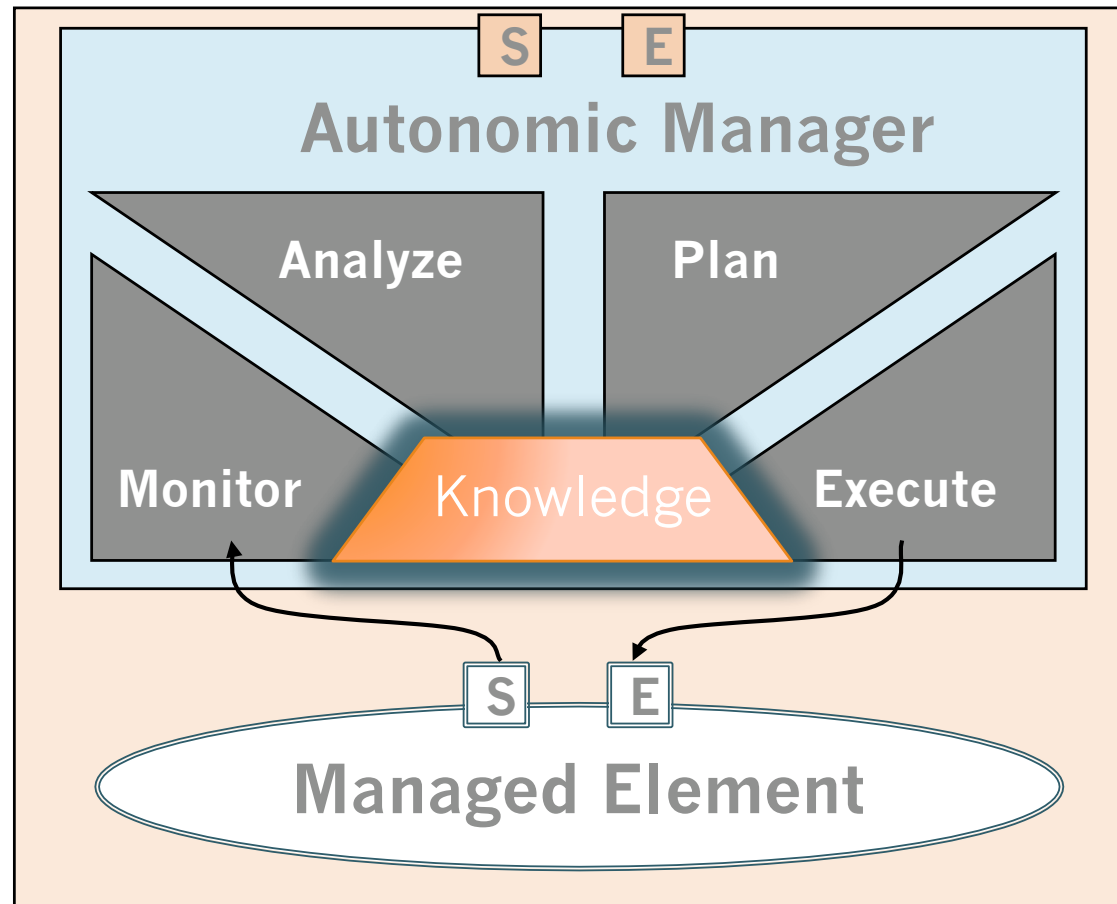
## Non-stationarity and heavy-tailedness can easily be confused

- The Hurst effect under trends. J. Appl. Probab., 20(3), 1983.
- Occasional structural breaks and long memory with an application to the S&P 500 absolute stock returns. J. Empirical Finance, 11(3), 2004.
- Testing for long-range dependence in the presence of shifting means or a slowly declining trend, using a variance-type estimator. J. Time Ser. Anal., 18(3), 1997.
- Long memory and regime switching. J. Econometrics, 105(1), 2001.

## NON-STATIONARITY IS A REASONABLE ALTERNATIVE




# Issue III: Intelligibility



# Issue III: Intelligibility

## How to build the knowledge?

- No Gold Standard, too rare experts
- Let's go to unsupervised learning, model-free policies eg Reinforcement Learning!
- Unfortunately, tabula rasa policies and vanilla ML methods are too often defeated




Exploration/exploitation  
tradeoff

# Outline

- ✓ Once upon a time...
- ✓ Scientific challenges
- ✓ Towards realistic behavioural models
- ✓ Conclusion and questions

# The ROC metrics: à la BQP

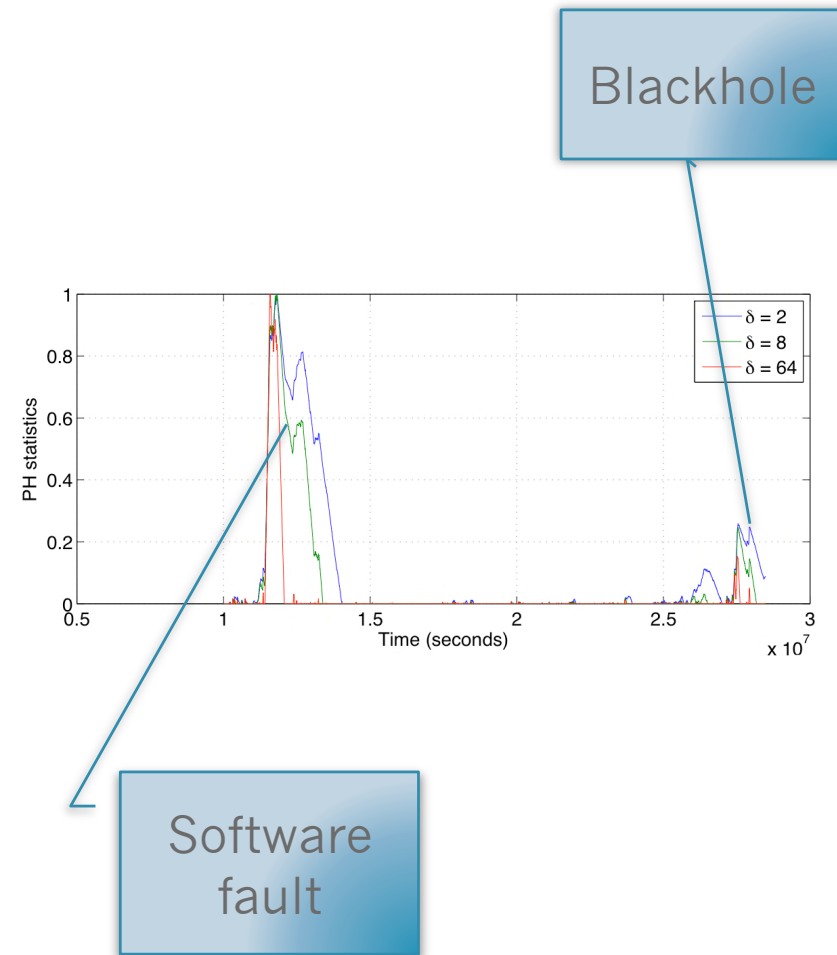


[C. Germain-Renaud et al. The Grid Observatory. CCGRID 2011]

# Dealing with non-stationarity

## 1. Statistical testing

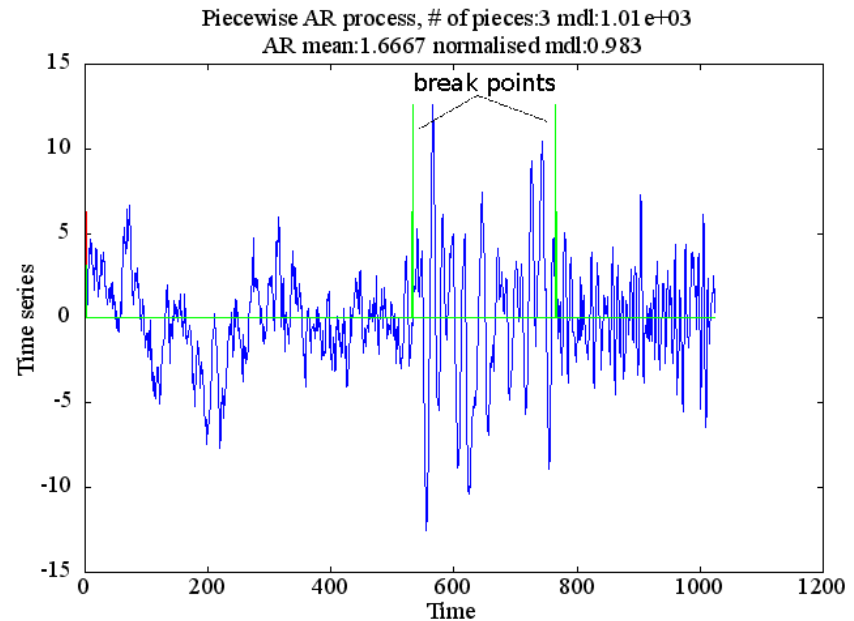
- Sequential jump detection
- Theoretical guarantees for known distributions
- Predictive, not generative
- Example: blackhole detection
- Calibration and Validation: by the Expert



# Dealing with non-stationarity

## 2. Segmentation

- Fit a piecewise time-series: infer the parameters of the local models and the breakpoints
- Model selection: AIC, MDL,... – based
- a priori hypotheses on the segment models: AR, ARMA, FARMA,...




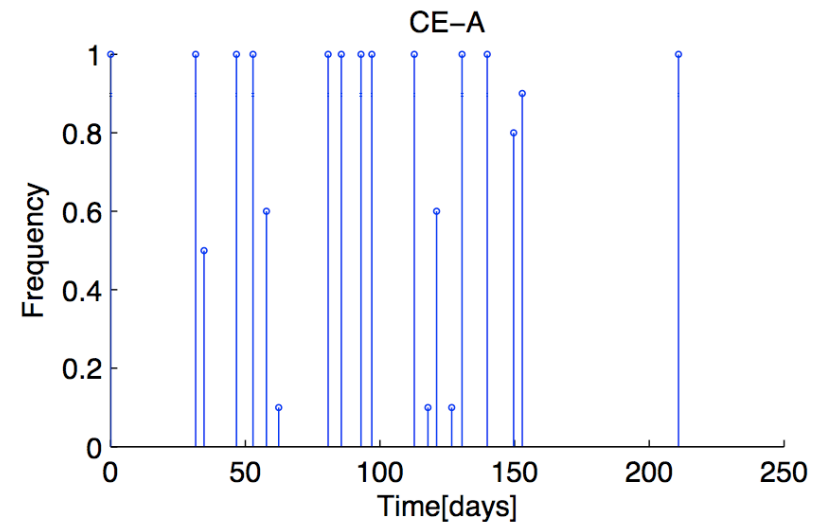
[T. Elteto et al. Towards non stationary Grid Models, to appear in JoGC Dec. 2011]



# Dealing with non-stationarity

## 2. Segmentation

- Mostly off-line and computationally expensive: generative, explanatory models
- Validation is not trivial
  - Fit quality
  - Stability: bootstrapping
  - Randomized optimization: clustering the results
- Hints at global behavior

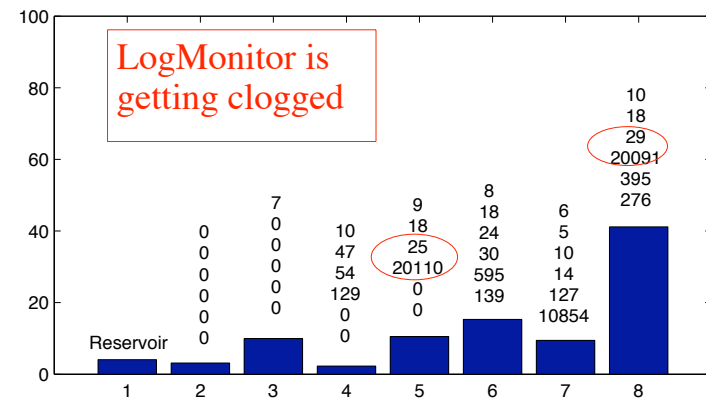
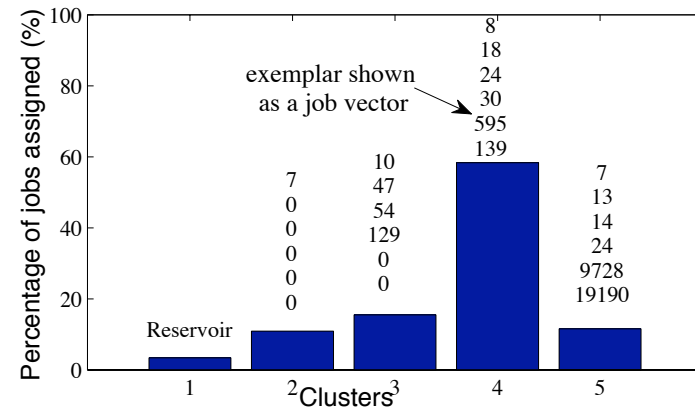


# Dealing with non-stationarity

## 3. Adaptive clustering:

- Clustering with Affinity Propagation (Frey & Dueck): the exemplars are the model
- Feedback from analysis to curation: feature selection instead of a priori

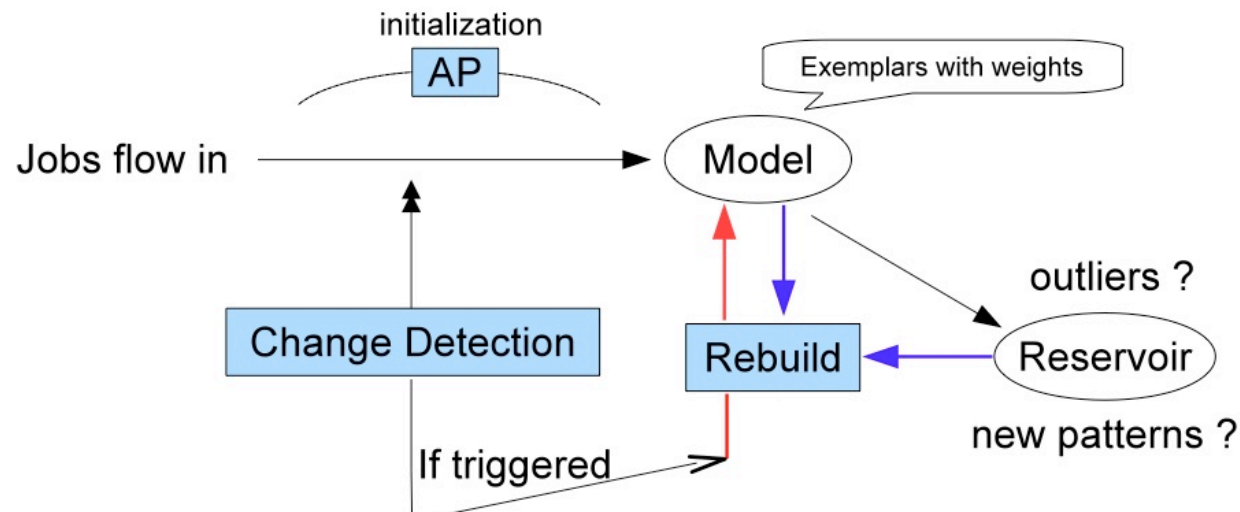
[X. Zhang et al. "Toward Autonomic Grids: Analyzing the Job Flow with Affinity Streaming". SIGKDD'2009]



# Dealing with non-stationarity

## 3. Adaptive clustering:

- Adaptive: on-line rupture detection
- Back to statistical testing, but on the model, not on the data



# Issue III: Intelligibility

- No Gold Standard, too rare experts: unsupervised learning, model-free policies eg Reinforcement learning
- Tabula Rasa policies, vanilla ML methods are too often defeated
- Exploration/exploitation tradeoff
- Ongoing work: infer causes from data as latent topics, in the spirit of text mining
- [Y. Kim et al. Characterizing E-Science File Access Behavior via Latent Dirichlet Allocation, to appear un UCC 2011]

# Conclusion



# The collaboration

- Born in EGEE-III, now a collaborative effort of
  - CNRS/UPS Laboratoire de Recherche en Informatique
  - CNRS/UPS Laboratoire de l'Accélérateur Linéaire
  - Imperial College London
  - France Grilles – French NGI of EGI
  - EGI-Inspire
  - Ile de France council
    - (Software and Complex Systems programme)
  - INRIA – Saclay (ADT programme)
  - CNRS (PEPS programme)
  - University Paris Sud (MRM programme)
- Scientific Collaborations
  - NSF Center for Autonomic Computing
  - European Middleware Initiative
  - Institut des Systèmes Complexes
  - Cardiff University

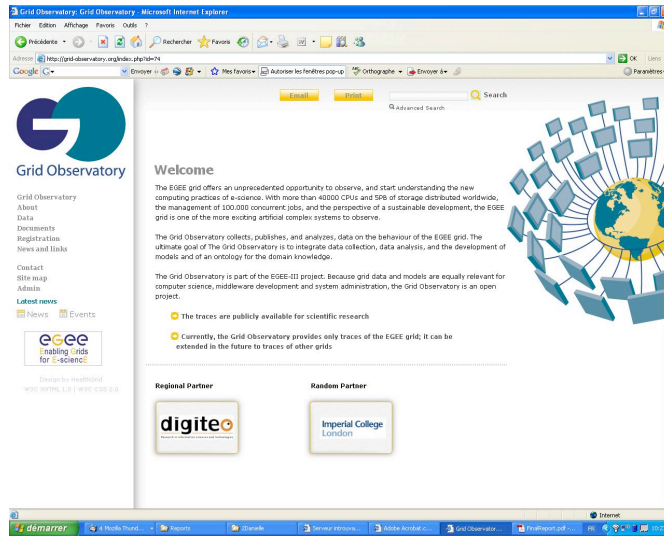




# The digital data

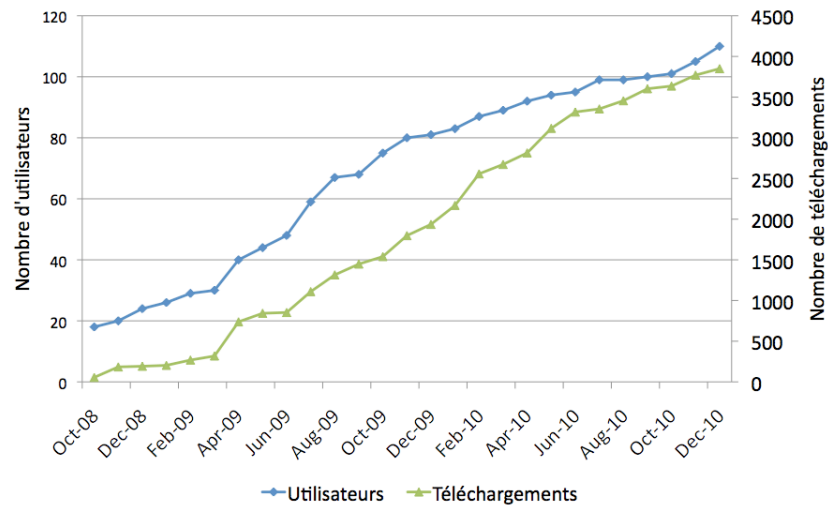
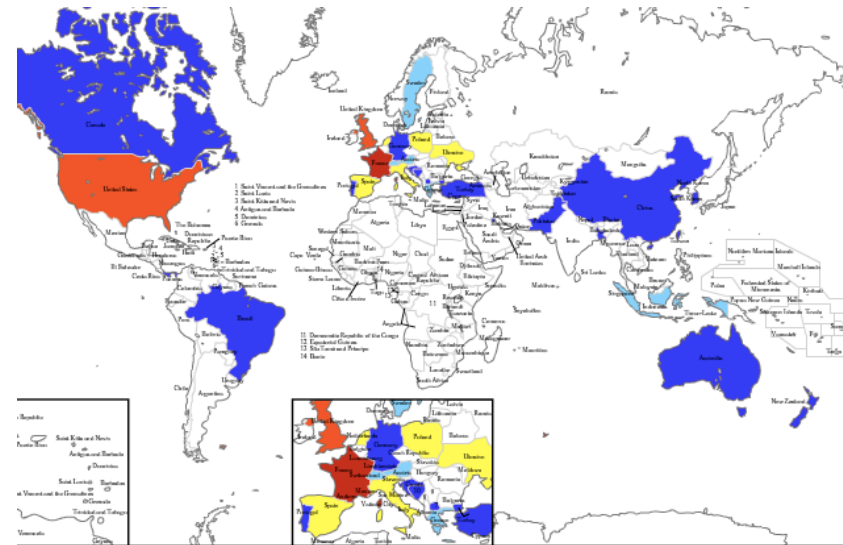
Component	Range	Scope	Format	Size
RTM	Comprehensive	gLite	Spec.	200MB
IS	Comprehensive	EGI	LDIF	300MB
L&B	Partial	gLite	SQL	2GB
Accounting	Local	gLite	PBS	6MB
CondorG	Partial	gLite	Spec.	15 KB
JobController	Partial	gLite	Spec.	40MB
LogMonitor	Partial	gLite	Spec.	70 MB
WorkloadManager	Partial	gLite	Spec.	70MB
GridFTP	Partial	EGI	Spec.	11MB

# The GO Portal



Traces available through the portal: no grid certificate

Global Impact



Use and users both increasing steadily



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

3.2.4180 Saturday, June 11, 2011 Done