VI Latin American Symposium on
High Performance Computing
HPCLatAm 2013
Mendoza, Argentina

Cloud Computing:
Helping Humanity to reach the next Final Frontier

Dr. José Luis Vázquez-Poletti

dsa-research.org

Distributed Systems Architecture Research Group
Universidad Complutense de Madrid

Research mentioned in this talk was funded by: MEDIANET (Comunidad de Madrid S2009/TIC-1468), ServiceCloud (MINECO TIN2012-31518) and MEIGA-METNET-PRECURSOR (AYA2009-14212-C05-05/ESP)

About me

José Luis Vázquez-Poletti

M.E. in Computer Science (2004) from Universidad Pontificia de Comillas and Ph.D. in Computer Science (2008) from Universidad Complutense de Madrid

Assistant Professor at Universidad Complutense de Madrid and Cloud Researcher at Distributed Systems Architecture Group

Directly involved in EU funded projects, such as EGEE and 4CaaSt

- 2005 – 2009: Application porting onto Grid Infrastructures (i.e. Fusion Physics and Bioinformatics) and training events
- Since 2010: Different aspects of Cloud Computing but always with applications in mind

@jlvazpol
http://dsa-research.org/jlvazquez/

About Frontiers

UNCuyo and about expanding frontiers

Battery recycling

Mobile plant that recycles batteries on the way.

Solar engine

Possible replacement for air conditioning units.

Biodiesel from algae

Attracted investments from US.

Hydrogen fuel

Collaboration with Centro Atómico de Bariloche.

Anti-earthquake systems

Structure reinforcement of legacy buildings and bridges.

…

HPCLatAm 2013

Continues to expand computational frontiers.

And…

UNCuyo and about expanding frontiers

… First HPC in the Cloud tutorial ever in Argentina!

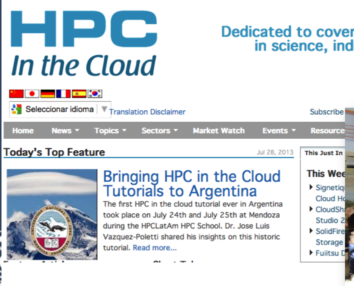

http://www.hpcinthecloud.com/

Computing as Humanity's tool

Konrad Zuse (1910-1995)

First working computer ever (Z3).

Yuri Gagarin (1934-1968)

First man in Space (thanks to a computer).

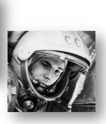

Neil Armstrong (1930-2012)

First man on the Moon (thanks to the Apollo Guidance Computer).

Don Estridge (1937-1985)

First PC (IBM).

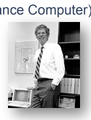

European Union (since 1951)

First reference multinational grid computing infrastructure (LCG, EGEE, EGI).

Expanded to Latin America (EELA).

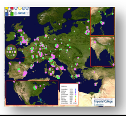

Various (since ???)

Cloud computing.

What is Cloud Computing?

Cloud

- A cloud is a visible mass of microscopic droplets of water or frozen crystals suspended in the atmosphere.
- *(or…)*

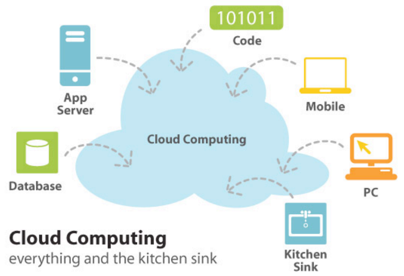

Cloud Computing layers

IaaS (Infrastructure as a Service)

- Externalization of computing servers, data warehouses, …
- Example: Amazon EC2

PaaS (Platform as a Service)

- Provision of application development and implementation platforms
- Example: Google Apps

SaaS (Software as a Service)

- Distribution of software and associated services (support)
- Example: Google Docs

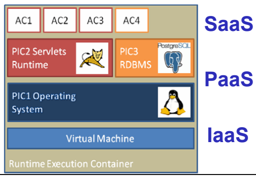

Cloud infrastructure types

Private Cloud

- Own physical resources
  - Limited
  - Can be distributed (virtual infrastructure manager is needed)
- Maximum management power
- Maintenance costs

Public Cloud

- 3rd Party resources
  - "Unlimited"
  - Location transparent to the user
- Only VMs are administered
- *Pay-as-you-go* basis

Hybrid Cloud

- Combination of previous types
- Virtual infrastructure manager is needed
- Example: Extra computing power when local resources are overrun

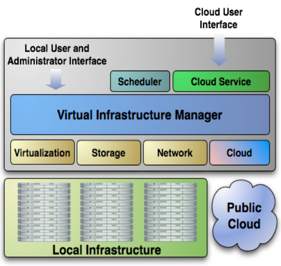

About Mars

Some Martian facts

4th planet from the Sun in Solar System

Iron oxide prevalent on surface gives red color

Rotational period: 24h40' hours

Half the radius of Earth

- 38% Earth's gravity

Thin atmosphere

- 95% carbon dioxide, 3% nitrogen, 1.5% argon, traces of oxygen and water
- Methane detected (volcanic, cometary impacts and/or microbial?)
- No magnetosphere

Once had large-scale water coverage

- Now only in poles and mid-latitudes

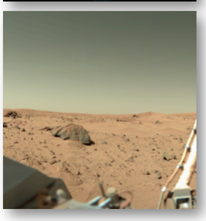

Why go to Mars?

Getting out of Earth to discover our origins

Comparative planetology can combine the in-sights gained from Mars and Earth (i.e. Martian climate development).

Exploring the nature of (possible) extraterrestial life to understand ours.

Revaluing the Scientific profession

The unknown territory and the limited predictability of events together with the long delay of signals (10' to 45') require an unique flexibility and talent to improvise.

NASA places its bets on astronauts and scientists for accomplishing the complex exploration of Mars.

Why go to Mars?

“Peace and Prosperity”

Mars missions include technological, economical, political and cultural aspects which would push the limits of technological innovations, promote a peaceful cooperation and a sense of global unity.

Mars in times of crisis

Mars missions create interesting jobs (more than 500.000 people were involved in the Apollo program) and provide motivation for scientific education.

Cost-efficient exploration missions would require a well-balanced combination of manned and unmanned activities complementing each other. Achievements can be applied to other fields (i.e. GPS, velcro, teflon, smoke detectors, diapers, …).

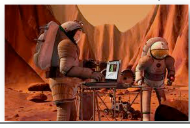

Missions to Mars

Mission	Country	Launch	Results
Marsnik-1	USSR	10/10/1960	Exploded before reach terrestrial orbit
Marsnik-2	USSR	10/14/1960	Exploded before reach terrestrial orbit
Sputnik 29	USSR	10/24/1962	Exploded in terrestrial orbit
Mars 1	**USSR**	**11/01/1962**	**Passed by Mars 200,000 Km.**
Sputnik 31	USSR	11/04/1962	Failure in terrestrial orbit
Zond 1	USSR	06/04/1964	Failure before reach terrestrial orbit
Mariner 3	USA	11/05/1964	Entered in Sun orbit
Mariner 4	**USA**	**11/28/1964**	**First Mars photos (21)**
Zond 2	USSR	11/30/1964	Communications failure
Zond 3	USSR	07/18/1965	Destroyed in terrestrial orbit
Mariner 6	USA	02/24/1969	Photograhies. Passed by Mars 3,215 Km.
Mariner 7	USA	03/27/1969	Photograhies. Passed by Mars 3,516 Km.
Mars 1969 A	USSR	03/27/1969	Launch failure
Mars 1969 B	USSR	04/02/1969	Launch failure
Mariner 8	USA	05/08/1971	Launch failure
Cosmos 419	USSR	05/10/1971	Launch failure
Mars 2	USSR	05/19/1971	Second artificial satellite of Mars. Surface module destroyed
Mars 3	USSR	05/28/1971	Third artificial satellite of Mars. Surface module sent signals for 20 seconds

Missions to Mars

Mission	Country	Launch	Results
Mariner 9	**USA**	**05/30/1971**	**First artificial satellite of Mars (7,329 Photos)**
Mars 4	USSR	07/21/1973	Passed by Mars 9846 Km.
Mars 5	USSR	07/25/1973	Operative 9 days in Martian orbit (60 Photos)
Mars 6	USSR	08/05/1973	Surface module sent data during the descent but crashed.
Mars 7	USSR	08/09/1973	Surface module passed by Mars 1,500 Km.
Viking 1	**USA**	**08/20/1975**	**First surface data. Operative during several years**
Viking 2	USA	09/09/1975	Second successful module. Operative during several years
Phobos 1	USSR	07/07/1988	Communications failure approaching Mars
Phobos 2	USSR	07/12/1988	Contact lost during obtaining Phobos photos
Mars Observer	USA	09/25/1992	Contact lost approaching Mars
Mars Global Surveyor	USA	11/07/1996	Operative until November 2006
Mars-96	Russia	11/16/1996	Failure leaving terrestrial orbit
Mars Pathfinder	**USA**	**12/04/1996**	**First robotic vehicle. More than 160,000 photos**
Nozomi	**Japan**	**07/04/1998**	**Failure before entering in Martian orbit**
Mars Climate Orbiter	USA	12/11/1998	Lost before entering in Martian orbit (1999 September 23th)
Mars Polar Lander	USA	01/03/1999	Lost landing

Missions to Mars

Mission	Country	Launch	Results
Mars Odyssey	USA	04/07/2001	Still operative
Mars Express	**ESA**	**06/02/2003**	**Orbital module operative. Surface module lost (Beagle 2)**
Mars Exploration Rover: Spirit	USA	06/10/2003	Operative for 7 years
Mars Exploration Rover: Opportunity	USA	07/07/2003	Still operative
Mars Reconnaissance Orbiter	USA	08/12/2005	Still operative
Phoenix	USA	08/04/2007	Operative for 5 months
Mars Science Laboratory	USA	26/11/2011	Still operative

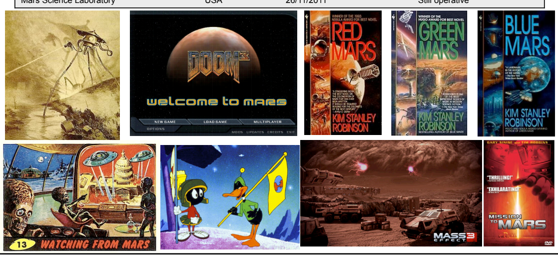

Ok, but…
Where are the Clouds?

NASA

Nebula: NASA’s Private Cloud (2008)

Emerged at NASA Ames Research Center and now supported by staff and infrastructure at Ames and at NASA Goddard Space Flight Center

“Computing Container as a Service”

- Open-source cloud computing project and service developed to provide an alternative to the costly construction of additional data centers whenever NASA scientist or engineers require additional data processing.
- Each shipping container data center can hold up to 15,000 CPU cores or 15 petabytes of storage while proving 50% more energy efficient than traditional data centers.

- Virtualization technologies
  - XEN and KVM hypervisors
  - Eucalyptus and then OpenStack virtual infrastructure manager

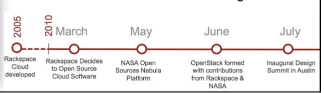

NASA

Some recent applications

SERVIR

- Integrates satellite observations, ground-based data and forecast models
- Monitors environmental changes and improves response to natural disasters

SPoRT Center

- Transitions unique NASA satellite observations and modeling capabilities to NOAA's National Weather Service
- Improves the analysis and prediction of weather events occurring within a 0-48 hour time-frame.

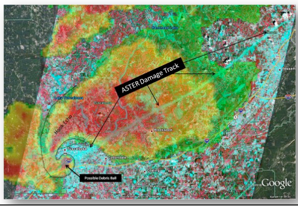

NASA

NASA uses Amazon EC2 (December 2010)

ATHLETE (All-Terrain Hex-Limbed Extra-Terrestrial Explorer)

- High resolution satellite images
- Process needed for guidance

Application: Polyphony

- Evaluates cloud capabilities
- Delivered to the Mars Science Laboratory

"AWS's resources completed the work in less than two hours on a cluster of 30 Cluster Compute Instances. This demonstrates a significant improvement over previous implementations."

"(Polyphony) allowed us to process nearly 200,000 Cassini images within a few hours under $200 on AWS."

amazon web services™

1 year before, in the Old World…

Mars MetNet

Atmospheric Science Mission to Mars initiated and defined by the Finnish Meteorological Institute (FMI).

Put together by FMI, Lavochkin Association (LA), the Russian Space Research Institute (IKI) and Instituto Nacional de Técnica Aeroespacial (INTA).

Universidad Complutense de Madrid (UCM) participates within the MEIGA Project (Science Team).

DSA Research Group from UCM collaborates by bringing Cloud Computing where it would benefit the Mission's applications and systems.

Finnish Meteorological Institute
Russian Space Research Institute
Lavochkin Association, Russia
INTA, Spain
METNET
The Next Generation Lander Mission for Martian Atmospheric Science

Mars Environmental Instrumentation for Ground and Atmosphere
MEIGA

Mars MetNet

The MetNet mission to Mars is based on a new type of semi-hard landing vehicle called MetNet Lander (MNL).

The scope of the MetNet Mission is eventually to deploy several tens of MNLs on the Martian surface.

The basic ideas of MetNet were cast by the FMI-team already in late 1980s. The development work started in the year 2000.

The first step in the MetNet Mission is to have a MetNet Mars Precursor Mission (MMPM) with a few MNLs deployed to Mars.

Mars MetNet

Mars MetNet

Application #1

Starting issues

- Probe landing location is needed for onboard instrument Calibration.
- Exact landing coordinates are unknown.
- Landing area (wide set of coordinates) is only determined 1h30' before entry procedure.
- Compass is useless and GPS is not offering service (yet) on Mars.

Approach

- Orientate through well known celestial objects

Mars MetNet: Phobos

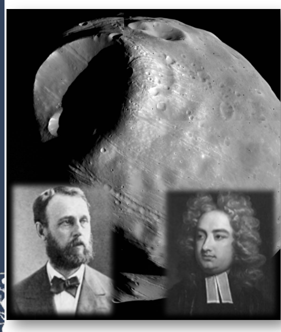

Phobos: The biggest Martian moon.

Discovered in 1877 by Asaph Hall or… by Jonathan Swift in his book "Gulliver's Travels" (1726).

The nearest moon to its planet in all the Solar System (6000 Km)

Orbit: 3 times/day aprox.

Diameter: 28x20 Km

Gravity: 0.00067g

Escape velocity: 25 Km/h

Very interesting effects on the Martian surface, **specially eclipses**.

Mars MetNet: Phobos

Phobos tracing application developed within the MEIGA project.

Phobos tracing will help in:

1. Obtaining probe landing coordinates
2. Data analysis

   (Eclipses = low radiation)

P. Romero, G. Barderas, J.L. Vázquez-Poletti and I.M. Llorente: *Chronogram to detect Phobos Eclipses on Mars with the MetNet Precursor Lander*. **Planetary and Space Science**, vol. 59, n. 13, 2011, pp. 1542-1550.

Mars MetNet: Phobos

Simulation in 1 area = 800 years in 1 coordinate

11 days 8 hours (without parallelism)

Using VMs from **Amazon EC2**

¿Which Interval/Task? ¿Which type/number of VMs?

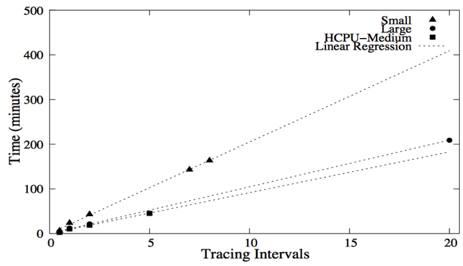

A Public Cloud Example: Amazon

Machine Type	Cores	C.U.	Memory	Storage	Platform
Standard On-Demand Instances					
Small (Default)	1	1	1.7GB	160GB	32bit
Large	2	2	7.5GB	850GB	64bit
Extra Large	4	2	15GB	1,690GB	64bit
High CPU On-Demand Instances					
Medium	2	2.5	1.7GB	350GB	32bit
Extra Large	8	2.5	7GB	1,690GB	64bit

amazon webservices™

Execution

Machine Type	Price in USA
Standard On-Demand Instances	
Small (Default)	$0.10/hour
Large	$0.40/hour
Extra Large	$0.80/hour
High CPU On-Demand Instances	
Medium	$0.20/hour
Extra Large	$0.80/hour

Mars MetNet: Phobos

Model for best strategies by means of **Time** and **Cost**

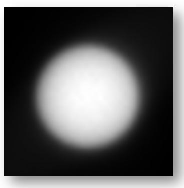

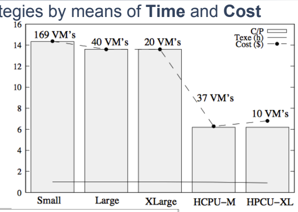

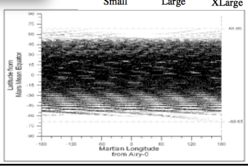

J.L. Vázquez-Poletti, G. Barderas, I.M. Llorente and P. Romero: *A Model for Efficient Onboard Actualization of an Instrumental Cyclogram for the Mars MetNet Mission on a Public Cloud Infrastructure.* **PARA2010: State of the Art in Scientific and Parallel Computing**, Reykjavík (Iceland), June 2010. Proceedings published in Lecture Notes in Computer Science (LNCS). Volume 7133, pp. 33-42, 2012. Springer Verlag.

Mars MetNet: Phobos

More in detail…

Possible solution: 37 HighCPU Medium Machines (**1h** ~ **$7.50**)

What is the price of a similar cluster for the chosen solution?

Example:

HP ProLiant DL170h G6 Server - **$4,909** x **37 nodes** = **$181,633**

What about administration? Electricity? Physical Security?

How many times will it be used at full power? Amortization?

Viking Landers

Application #2

Starting issues

- 2 landers (1976-1980) with different sensors (temperature, pressure).
- Data taken in different intervals and some got corrupted.
- Measurements will help MetNet.

Approach

- Some of the “inconsistent” data could be provoked by Phobos (or Deimos) eclipses.
- Other may indicate a temporal malfunction of sensors.

A.-M. Harri, W. Schmidt, P. Romero, L. Vazquez, G. Barderas, O. Kemppinen, C. Aguirre, J.L. Vazquez-Poletti, I.M. Llorente and H. Haukka: *Phobos Eclipse Detection on Mars, Theory and Practice*. **Finnish Meteorological Institute Research Report 2012:2**, Finland, 2012.

Viking Landers

Application #2

- Framework accommodating weather models for Mission production.
- Optimum execution in mind (Cost/Throughput now).
- Reuse previous work (below) but now some tasks have different sizes.

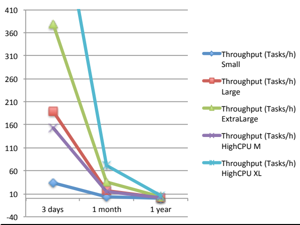

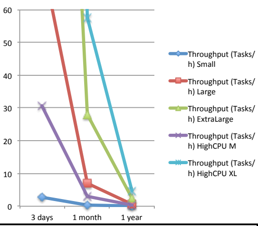

J.L. Vázquez-Poletti, J. Perhac, J. Ryan and A.C. Elster: *THOR: A Transparent Heterogeneous Open Resource Framework*. **Workshop on High Performance Computing on Complex Environments (HPCCE) on the IEEE International Conference on Cluster Computing 2010**, Heraklion, Crete (Greece), September 2010. Proceedings published online by IEEE Computer Society Press (DOI: 10.1109/CLUSTERWKSP.2010.5613099).

Mars MetNet

Application #3

Starting point

- Mars Science Laboratory (Curiosity)
- Phobos eclipses prediction
  - Sol 37 (13/09/2012)
  - Sol 41 (17/09/2012)

Results

- On-site validation
  - <1 second precision!!!
- The application is ready for its use on the Mars MetNet probes

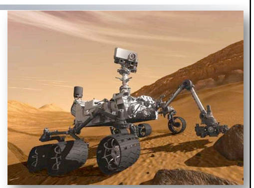

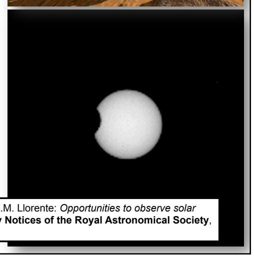

G. Barderas, P. Romero, L. Vazquez, J.L. Vazquez-Poletti and I.M. Llorente: *Opportunities to observe solar eclipses by Phobos with the Mars Science Laboratory*. **Monthly Notices of the Royal Astronomical Society**, 2012, Volume 426, Number 4, pp. 3195–3200. Wiley.

Mars MetNet

Application #4

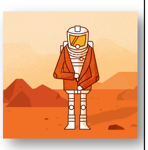

Starting point

- Need of a Martian meteorological model

Computational workplan

- Cost optimization of terrestrial meteorological models
  - Will apply to Martian models later
- Validation of proposed Martian models
  - Huge amount of data process

Mars MetNet

Application #4

WRF

- Weather Research & Forecasting Model (numeric mesoscale)
- Models:
  - **GFS:** free data, sometimes imprecise. 2 level computation.
  - **ECMWF:** restricted access data. 1 level computation.
- Amazon EC2 Machines:
  - **Cluster Compute Eight Extra Large** 60.5 GiB memory, 88 EC2 Compute Units, 3370 GB of local instance storage, 64-bit platform, 10 Gigabit Ethernet
  - **High-Memory Quadruple Extra Large** Instance 68.4 GiB of memory, 26 EC2 Compute Units (8 virtual cores with 3.25 EC2 Compute Units each), 1690 GB of local instance storage, 64-bit platform

amazon webservices™

dsa-research.org

Mars MetNet

Application #4

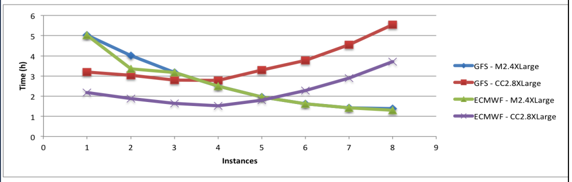

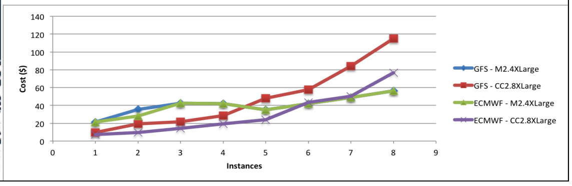

dsa-research.org

Another tool for expanding Humanity's limits

Cloud Computing is taking the role of ancient computing systems which brought Humanity to Space.

Mars represents the nearest frontier to be broken by Humanity.

As in the rest of scenarios, achievements will directly affect technology and general knowledge. Cloud Computing won't be an exception.

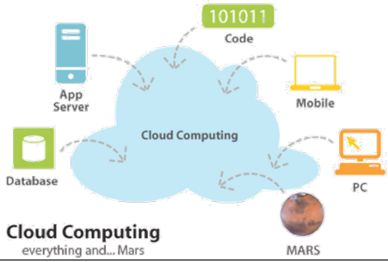

Research mentioned in this talk was funded by: MEDIANET (Comunidad de Madrid S2009/TIC-1468), ServiceCloud (MINECO TIN2012-31518) and MEIGA-METNET-PRECURSOR (AYA2009-14212-C05-05/ESP).