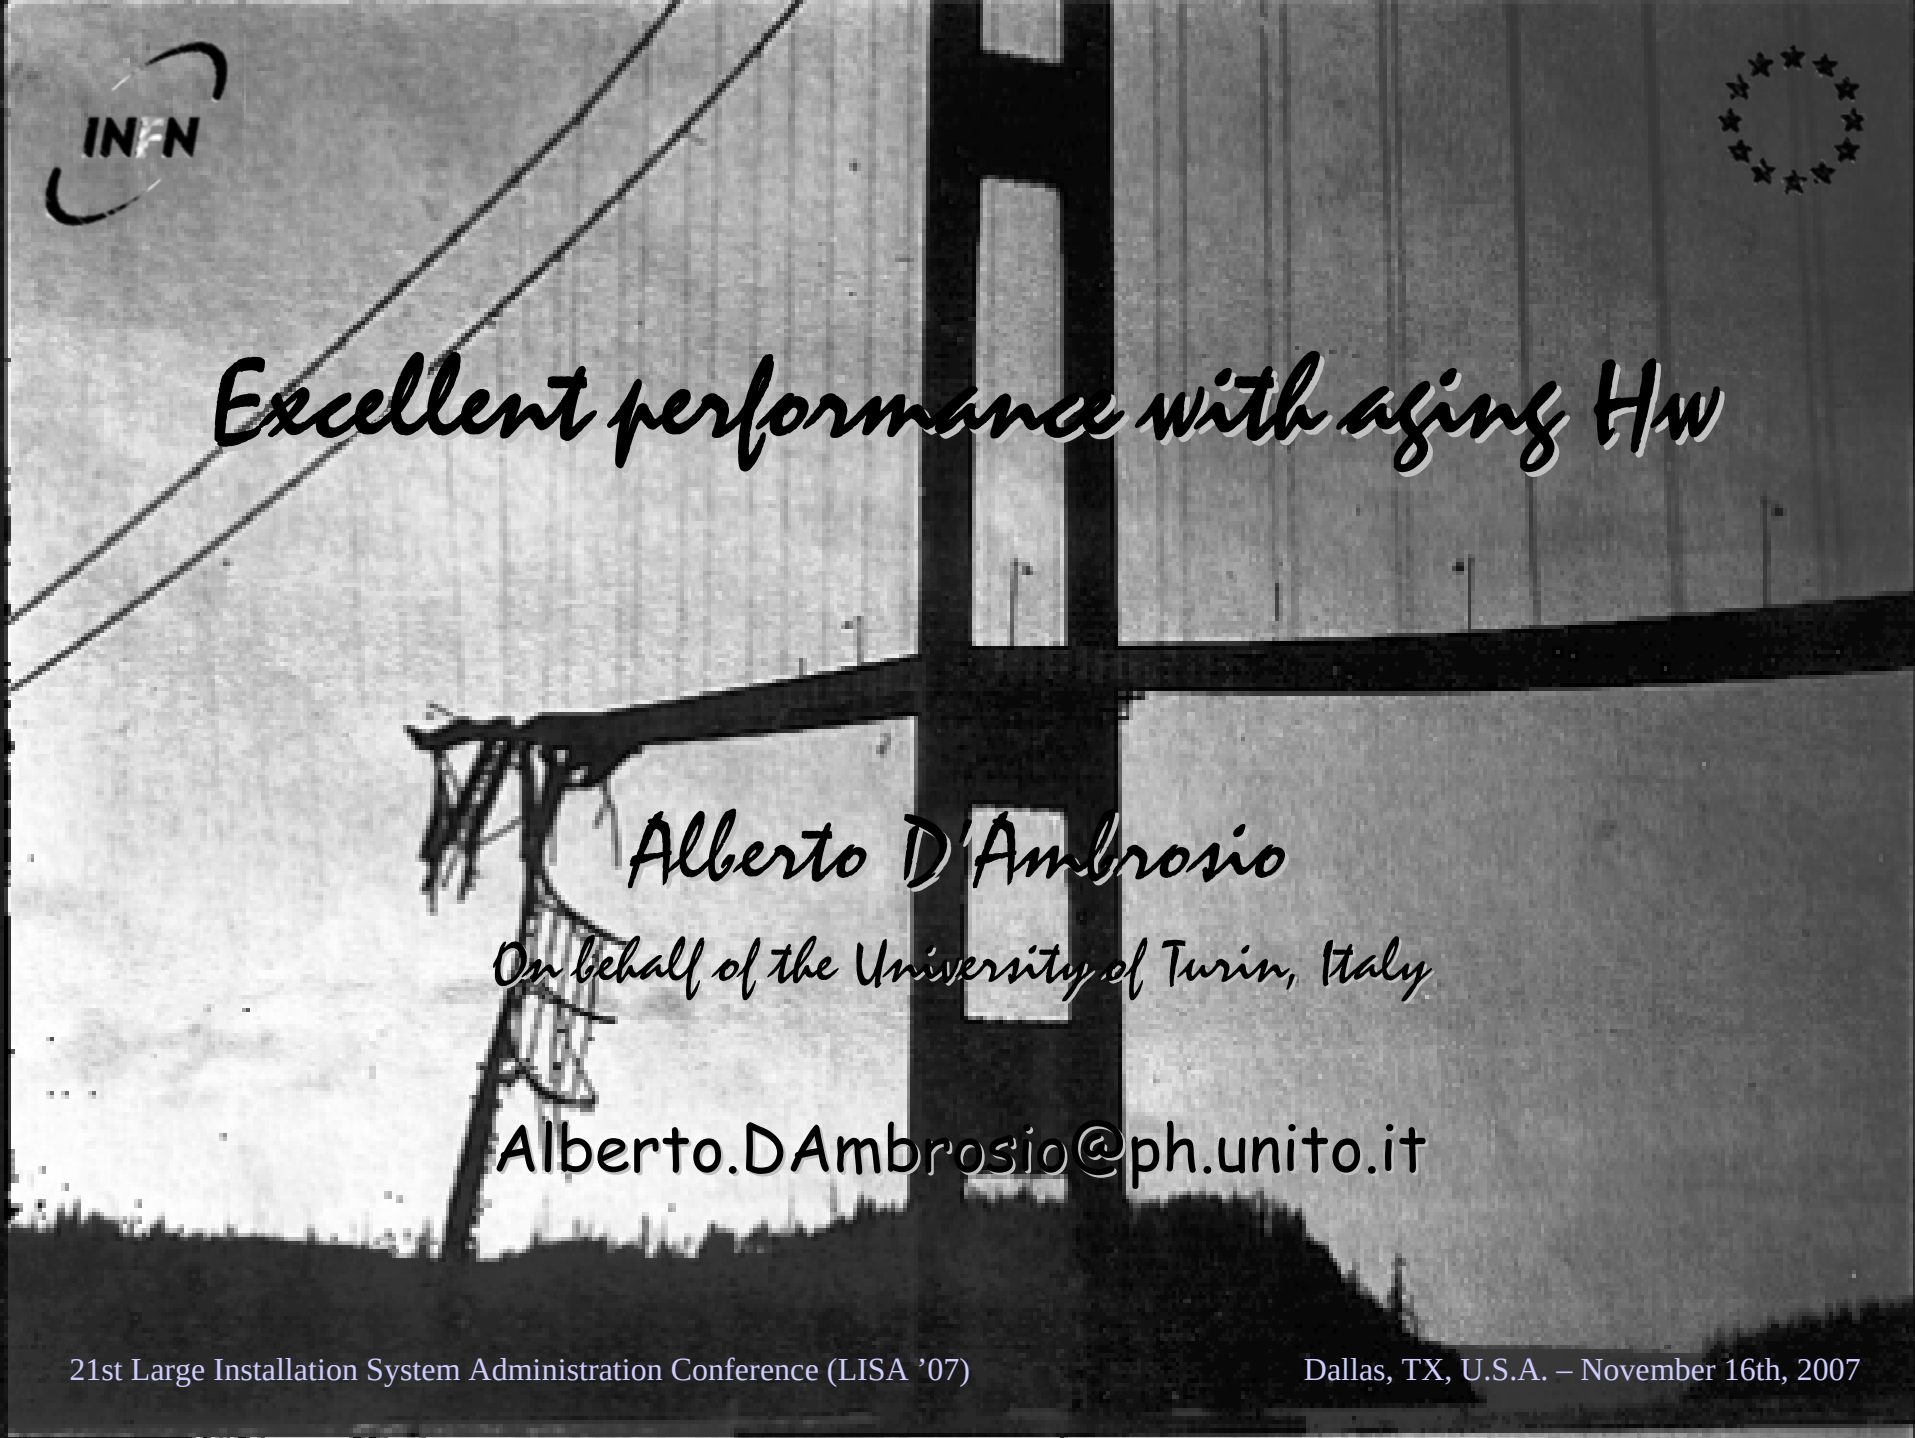

Excellent performance with aging Hw

Alberto D'Ambrosio

On behalf of the University of Turin, Italy

Alberto.DAmbrosio@ph.unito.it

THE SAFETY FACTOR
(From "Programming Pearls", by Jon Bentley)

The Brooklyn Bridge is the only suspension bridge of its era still standing:

- John Roebling had sense enough to know what he *didn't know*.
- He designed the bridge six times as strong as it would have needed to be.
- He built a good bridge by employing a huge safety factor to compensate for his ignorance.

THE FACTS OF LIFE

Unfortunately, due to factors you can't change (budget restrictions, politics, etc..), you often have to administer a system designed by someone else.

MY CASE STORY

For our departmental mail-server I was given :

- 2 powerful workstations (instead of real servers)
- A storage subsystem no longer produced a couple of months after the purchase
- Disk space barely sufficient for our initial needs

FIXES, TRICKS, PATCHES AND WORKAROUNDS From 2002 till' 2004

Everything went well for a couple of years:

- Brand-new mail-server processing 5 Kmsgs/day
  - The system was able to manage this workload
- Installed anti-spam filter
  - The workload increased; system still able to manage it.
- 10 Kmsgs/day
  - The workload increased; system still able to manage it.
- Enabled Bayesian correction for the spam-filter
  - The workload increased; system sometimes overloaded.

FIXES, TRICKS, PATCHES AND WORKAROUNDS From 2004 till' 2005

First problems, first fixes:

- Installed E-Mail content-scanner with anti-virus filter
  - Too much swapping activity
  - Frequent SMTP-OUT timeouts (lots of user complaints)
  - FIXES:
    - **Increased RAM (twice the initial value)**
- 15 Kmsgs/day
  - System often overloaded, mainly in the rush hours.
  - "disk full" error messages
  - FIXES:
    - **Enabled REJECT of SPAM messages**
    - **One machine added to the cluster**
    - **Storage increased (10 times the initial global value!)**

FIXES, TRICKS, PATCHES AND WORKAROUNDS From 2005 till' 2006

More problems, more fixes:

- One machine added to the cluster
  - Cluster less "reactive" due to *SCSI starvation*
  - FIXES:
    - **All 3 cluster members plugged to a SCSI arbiter**
- 20 Kmsgs/day
  - Disk I/O timeout error messages
  - FIXES:
    - **Swap area moved from shared (slow) storage to local (fast) disk**
    - **Temporary area of E-Mail content-scanner moved from shared (slow) storage to local RAM disk**

FIXES, TRICKS, PATCHES AND WORKAROUNDS From 2006 till' 2007

The last fixes:

- 40 Kmsgs/day
  - Locking issues on file-based Bayesian DB
  - FIXES:
    - **Bayesian single-file DB moved to MySQL server**

- 60 Kmsgs/day
  - Very frequent SMTP-OUT timeouts (a deluge of user complaints)
  - FIXES:
    - **SMTP-OUT moved to an external (caching) machine**

THE WHOLE STORY (2002÷2007)
From 5000 to 65000 msgs/day!

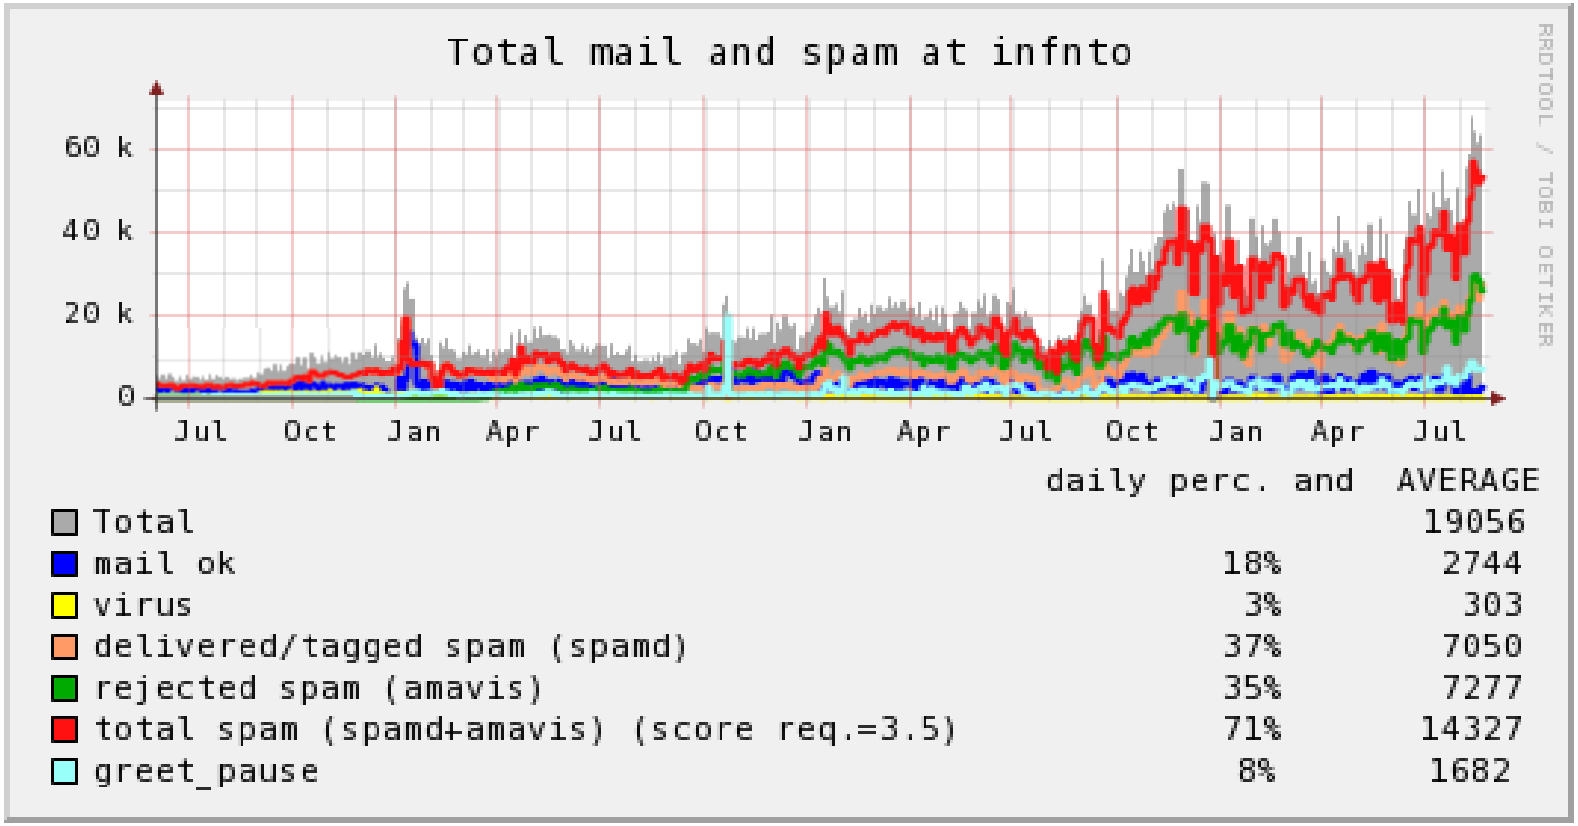

	daily	perc.	and	AVERAGE
Total				19056
mail ok		18%		2744
virus		3%		303
delivered/tagged spam (spamd)		37%		7050
rejected spam (amavis)		35%		7277
total spam (spamd+amavis) (score req.=3.5)		71%		14327
greet_pause		8%		1682

FINAL ACHIEVEMENTS AND SUCCESSIVE ACTIVITY Fall 2006 and mid 2007

Last glories and rising from the ashes:

- December 2006: Performance satisfied request; with this Hw we had our **LISA '06** paper accepted:
  - **http://www.usenix.org/events/lisa06/tech/cecchini.html**

- August 2007: Replaced with **new** Hw, and recycled as interactive machine:
  - ***IT'S STILL ALIVE & KICKIN'!***

CONCLUSIONS – The lesson (not) learned…

- A system designer should have sense enough to know what he doesn't know.
- He will be able to design a good system only by employing a huge safety factor to compensate for his ignorance.
- Because of several factors (budget restrictions, politics, etc..) you will hardly ever have a well-designed system. In fact…
- **Our new system is just sufficient for our current needs. Maybe the story will repeat…**

ACKNOWLEDGMENTS
Thanks to...

- Giampiero Passarino, the fifth floor financer.
- Lodovico Riccati, who made some dreams come true.
- Albert Werbrouck, wonderful *draft reviewer*.
- Silio d'Angelo, spiritual IT guide.
- Giorgio & Franca, very patient colleagues.

and, last but not least:

- Lorenzo & Elisa, at home!

Excellent performance with aging Hw

Alberto D'Ambrosio

On behalf of the University of Turin, Italy

Alberto.DAmbrosio@ph.unito.it

http://www.to.infn.it/~gandalf/