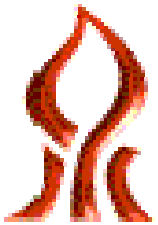

Self-stabilizing Host Systems

SHOS…

Shlomi Dolev and Reuven Yagel,
Computer Science Department Ben-Gurion University of the Negev, Beer-Sheva, Israel

dolev,yagel@cs.bgu.ac.il

Virtualization@SYSTOR
29 Oct. 2007

Can a guest exhaust a VMM?

- A soft-error moves a KVM guest into VMX-root
  - Not so rare: "a potential error every five hours" [Mastipuram &Wee '04]
- A KVM bug?
  - Many lines of new code (> 10KSLOC)…
- Linux security/robustness issues
  - An unaware user runs a faulty program with high privilege
- In our design the host can reach any state and still eventually provide predefined guarantees

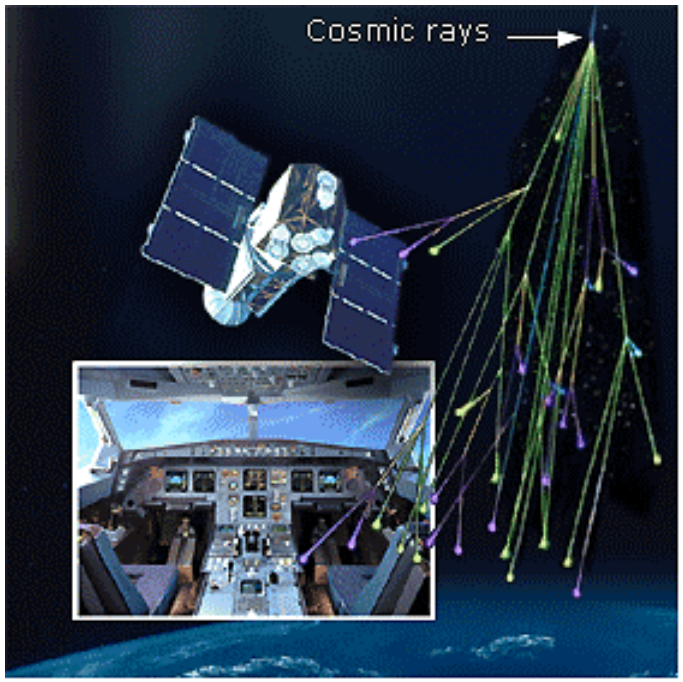

Self-Stabilization

- The combination and type of faults cannot be totally anticipated in on-going systems
- Elegant fault tolerant approach
  - Alternative to the various duplication techniques
  - Started at any state, the system converges to a desired behavior

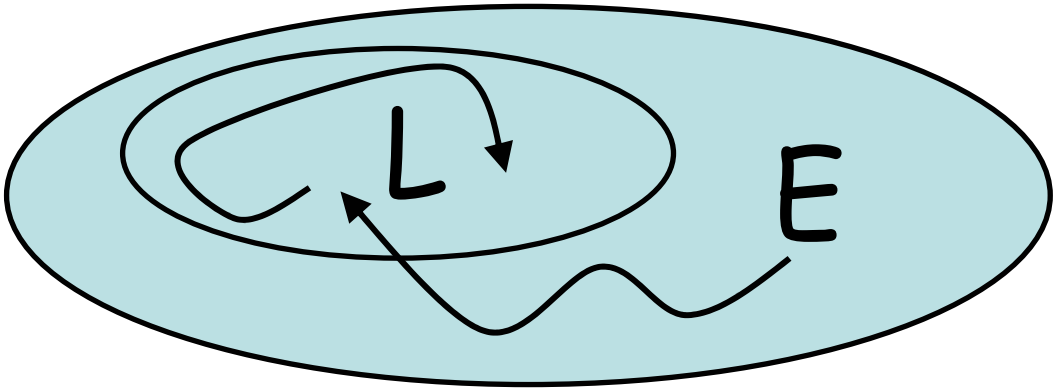

- "Self-Stabilization in Spite of Distributed Control" [Dijkstra '74]
- Self-Stabilization [Dolev '00]

A Self-stabilizing Stack

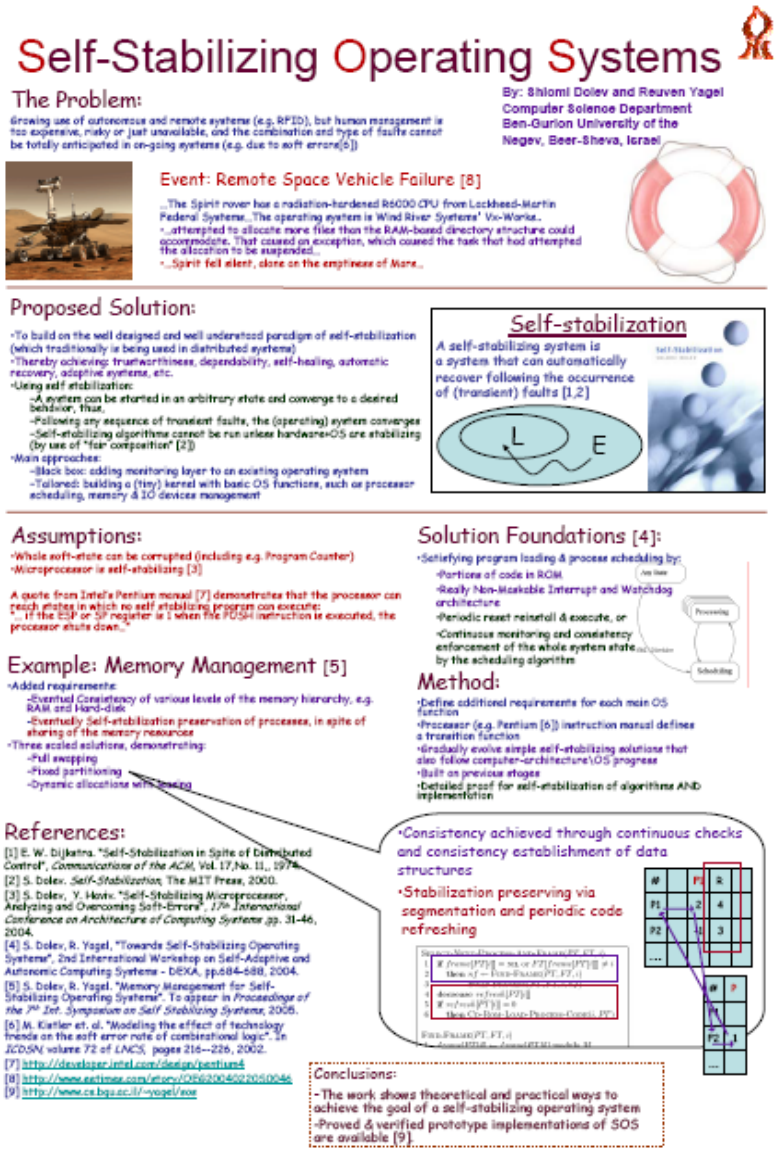

Stabilizing Reputation & Trust

- Give a chance to change reputation, both ways…
- Maybe the reputation history is corrupted
- Constantly fading old reputation and accumulating new reputation
- Thus, stabilization of reputation and trust

Main SHOS Concepts

- Secure booting of minimal TCB
- Offline Byzantine behavior detectors
- Runtime anti-Byzantine contract enforcers
- Stabilizing trust and reputation
- Self-stabilization

➢ Stability of non-Byzantine programs is preserved from any state

Conclusion

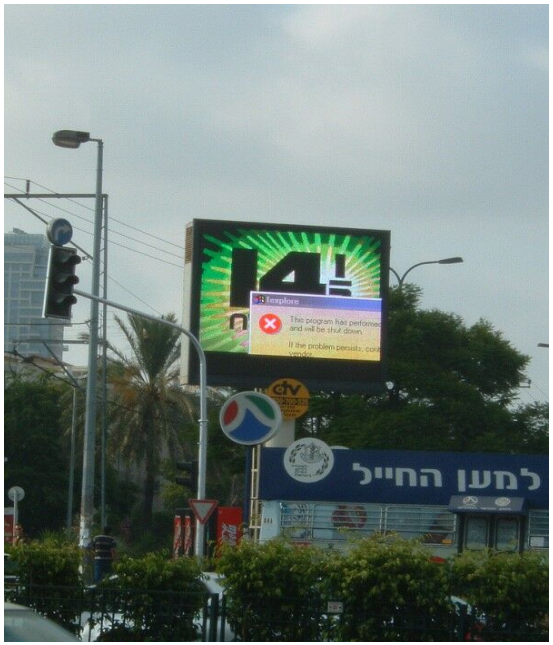

- We presented a general design for stabilizing host systems
- Virtualization can benefit from the foundation of self-stabilization

http://www.cs.bgu.ac.il/~yagel/sos

Thank you