

Elements of the Self-Healing System Problem Space

Phil Koopman
Carnegie Mellon University

WADS, May 2003



Overview

- x **“Self-Healing” – it’s getting attention, but what does it mean?**
 - This talk is based on observations from the most recent Workshop on Self-Healing Systems (WOSS’02)

- x **Description of some general problem elements of Self Healing research**
 - Fault models – what is an “injury”?
 - System responses – what is “healing”?
 - System incompleteness – what’s unknown?
 - Design context – what injuries are beyond healing?

- x **Two challenges:**
 1. *Fault Tolerant Computing*: broaden perspectives with SH ideas
 2. *Self Healing*: don’t waste time reinventing existing FT ideas

Fault Model – “injury”

- x **First question in fault tolerant computing is:**

“What is the fault model?”

- x **Reasons for a fault model**
 - Need to know expected faults to measure fault tolerance coverage
 - Not all faults are equal in time, space, severity
- x **Some challenges:**
 - Is Injury == Fault ????
 - Is a software defect an injury?

Self-Healing Fault Model Issues

x **Fault duration:**

- Permanent / intermittent / transient

x **Fault manifestation:**

- Fail silent / Byzantine / correlated faults
- Impaired: run-time, reserve capacity, brittleness, resource consumption

x **Fault source:**

- Wear-out / design defects / reqts. defects / environment change / malicious

x **Granularity:**

- One designer's "system" is the next level designer's "component"
- Transistor failure / ... node failure ... / system failure

x **Fault profile expectations:**

- No faults / historically known faults / foreseen faults / unforeseen faults
- Random+independent / random+correlated / expected / predicted

System Response – “healing”

- x **After an injury, what happens?**

- x **Fault tolerant system responses include:**
 - Diagnosis / identification
 - Isolation / containment
 - System reconfiguration
 - System reinitialization

- x **Does “healing” mean something additional?**
 - Or is it a difference at a different level?

Self Healing System Responses

x **Fault Detection:**

- Self-test / pairwise checking / peer checking / supervisor checking
- Self-injected faults to ensure detection is working?

x **Degradation during & after healing:**

- Fail-operational / degraded performance / fail-fast+ fail-safe

x **Response:**

- Fault masking / failover / reconfiguration
- Optimize for: safety / reliability / availability / ...
- Preventative (periodic reboot) / Proactive (diagnosis-based) / Reactive

x **Recovery of state:**

- Hot swap / restore quiescent state / warm boot / cold boot
- Rollback / recovery block / control gain changes / rollforward / run-while-reconfiguring
- What about recovering component state?

x **Time constants:**

- Most faults are transient
- Important that system response time constant be faster than injury arrival rate

x **System Assurance:**

- After injury / during healing / after healing

System Completeness – *What do we know and when?*

x **System self-knowledge**

- How much self-knowledge is required for healing?
- How should healing knowledge be abstracted?
- How do we deal with not knowing how much the system doesn't know?

x **Designer knowledge**

- Not all systems are complete when design is “done”
- Even if complete, we won't know everything about all components
- How do we deal with not knowing how much we don't know?

Self Healing System Completeness

x **Architectural Completeness:**

- Proprietary & known / open & regulated / extensible

x **Designer Knowledge:**

- Component knowledge (especially COTS components)
- Faulty behavior characterizations
- How do you heal after suffering a component behavior that is “unspecified”?

x **System Self-Knowledge:**

- How complete is system’s self-model? (idea of reflection)
- Is healing an intentional or emergent behavior?

x **System Evolution**

- Configuration changes & usage changes
- Are outages random / predictable / schedulable?

Design Context – *What are the scope limits?*

- x **The real world is a messy place – what assumptions are made?**
 - Homogeneous system?
 - “Perfect” components (e.g., perfect healing management software?)
 - ...

- x **What is the size of the system?**
 - A single software module?
 - A complex software system?
 - A person plus a computer system?
 - The North American power grid?
 - The Internet?

- Does teaching users to press CTL-ALT-DEL achieve “self-healing” of the user+computer “system”?

Self Healing Design Context

- x **Abstraction Level:**
 - Implementation / design / architecture / ...

- x **Component Homogeneity:**
 - Can any software component run in any node?
 - Perfect configuration homogeneity / plug-compatible / heterogeneous

- x **Predetermination of system behavior:**
 - Specific design / rule-based system / service discovery / emergent behavior

- x **User Involvement in healing:**
 - User direction / user-provided hints / user ability to tune / invisible to user

- x **System Linearity:**
 - Linear+composable / monotonic / mildly discontinuous / arbitrary
 - Single operating mode / mode changes

- x **System scope:**
 - Component / computer system / computer+person / enterprise / society

Conclusions

- x **“Self-Healing” potentially encompasses a lot of ground**
 - Smaller than expected intersection of research assumptions at WOSS02
 - Consensus will take a while

- x **Some of this has been done before!**
 - Fault models – well known in FT, don’t reinvent without good reason
 - System responses – how different are they from FT?
 - System incompleteness – FT usually assumes relative completeness
 - Design context – plenty of room for novelty in both FT & SH

 - **But there is plenty of room for more good research**

- x **A final thought:**
 1. *Fault Tolerant Computing*: broaden perspectives with SH ideas
 2. *Self Healing*: don’t waste time reinventing existing FT ideas
even better: articulate the novelty of approaches