



Ultra-reliable Communication Architectures

Alfons Geser

Formal Methods Team, National Institute of Aerospace
Hampton, VA

Based on joint work with Paul S. Miner, Jeff Maddalon, Lee S. Pike
(NASA Langley Research Center, Hampton, VA)

SEAMS Workshop 2004, Charleston, SC



Overview

- The National Institute of Aerospace
- Fault-Tolerance
- The SPIDER Project

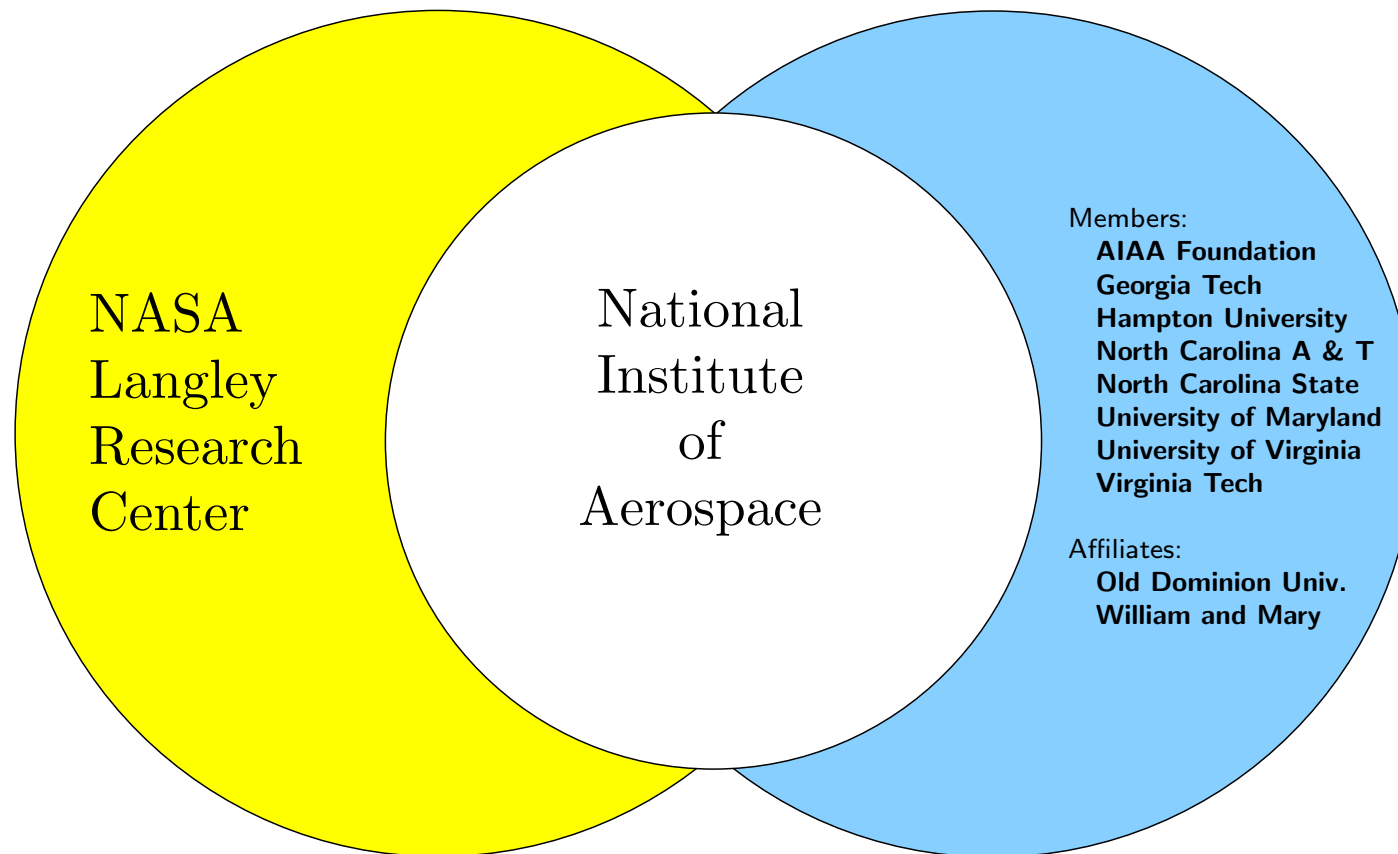


The National Institute of Aerospace (NIA)

- An Independent, Non-profit, 501(c)(3), Research and Graduate Education Institute formed in 2002 by a Consortium of Six Universities and the AIAA Foundation
- Established to Serve as a Collaborative Partner for NASA Langley Research Center
- Operates Under a Long-term Cooperative Agreement with LaRC (Five Years, plus 3 Five Year Options)
- Conducts Collaborative Research with NASA, University Faculty, Resident Staff, Industrial Partners, other Government Agencies, and other Non-profit Institutes
- Offers Full-time and Part-time Resident Graduate Education in Engineering and the Sciences from Member Universities



The NASA-NIA Collaborative Network





NIA Graduate Education Program

- M.S. and Ph.D. offered from all six founding universities
 - Enrollment at one of the member universities
 - All residency requirements fulfilled in Hampton, Virginia
 - Up to 50% of coursework from other member schools
- Outstanding Courses
 - Best courses from member universities
 - Coordinated offerings from schools in advanced ‘niche’ Ph.D. courses
 - New team-developed and -taught courses in emerging areas
- Exceptional Students
 - Top students nationwide attracted by Rising Star Fellowships: generous financial package, research at NASA, courses at NIA
- Vibrant Graduate Environment Planned
 - 10-20 Faculty (Langley professors, NIA Professors, Liaison Professors)
 - 50-60 full-time grad students
 - Active seminar, colloquia and short course program



Major Collaborative Research Thrusts

Large Space Structures Assembly

Univ. of Maryland, NIA resident staff, Systems Integration Branch

Shuttle ET Foam Investigations

NIA resident staff, Advanced Materials and Processing Branch

Fatigue Crack Propagation in Al Alloys

AFRL, Univ. of Virginia, Metals and Thermal Structures Branch

Flight Crew Operations within Simulated DAG-TM

NIA resident staff, small business, Crew Systems and Operations Branch

Rotorcraft Aeromechanics Research

ARL, NIA resident staff, Aeroelasticity Branch

Inflatable Large Space Structures

DARPA, Virginia Tech, NIA graduate students, Structural Dynamics Branch

Aerocapture and Planetary Flight Vehicle Engineering


NIA resident faculty and graduate students, Vehicle Analysis Branch

Software Certification via Formal Methods

NIA resident staff, NIA summer visitors, Assessment Technology Branch



Embedded Systems



embedded: automated, without human supervision

dependable: precisely defined, reliable services

hard real-time: every service has a strict deadline



Faults and Failures

- *fault*: damaged hardware

Example: worn-out bearing (*permanent* fault),
electromagnetic interference (*transient* fault),
loose contact (*intermittent* fault)

Faults are rare but unavoidable (CPU ca. $10^{-6}/h$)

- *failure*: the delivered service differs from the specified service
may be catastrophic (e.g., plane crash)

Catastrophic failures must be avoided



Ultra-Reliability

Assuming a fleet of 100 aircraft, each flying 3000 h/a over 33 years lifetime (thereby accumulating 10^7 flight years), and 10 catastrophic failure conditions, no loss may happen over the lifetime of the fleet. Such systems must have a failure rate $< 10^{-9}/h$ (*Rushby 93*).

A system having a failure rate $< 10^{-9}/h$ is called *ultra-reliable*.

Just how little is $10^{-9}/h$?

- $10^9 h \approx 114,000$ years
- rate of being killed by lightning in US: $3.5 * 10^{-11}/h$ (93 per year)
- rate of dying in a car accident in US: $1.5 * 10^{-8}/h$ (40,000 per year)



Fault Tolerance

$$P(\text{no failure}) = P(\text{no fault}) + \\ P(\text{no failure/some faults}) * P(\text{some faults})$$

A system is *fault tolerant* if $P(\text{no failure/some faults}) > 0$.
Fault tolerance is achieved by replication (redundancy).



How to Prove Fault Tolerance

Determine a *fault hypothesis*(FH) such that

- $FH \Rightarrow$ no failure
- $P(\neg FH) < 10^{-9}/h$.

Example: FH = “at most one fault at a time”.



Fault Containment Unit (FCU)

Life-testing of ultra-reliable systems is infeasible (*Butler/Finelli 93*).

We need to mathematically derive the system's reliability bound from the reliability of its components.

fault containment unit (FCU): all hardware in the system that a single fault can physically reach.

Example: fire-resistant cabinets.

A tyre of a Concorde is *not* an FCU.

FCUs fail independently.

Given the fault rate of an FCU, the fault rate of the system can be calculated using Markov Chains.



Checks

Some bad messages can be detected at the receiver.


Example: CRC check, timing window

message *accepted*: not detectably bad

value fault: accepted bad message



Redundant Control




Let V denote the accepted messages among $\{v1, v2\}$.
If $V \neq \emptyset$ then let $v' \in V$.

A value fault may prevail.



Triple Modular Redundancy



Let V denote the accepted messages among $\{v1, v2, v3\}$.
Let v' be the majority in V if one exists.

Example: $\text{maj}(1, 2, 1) = 1$.

If the good nodes agree and are in the majority then v' is good.



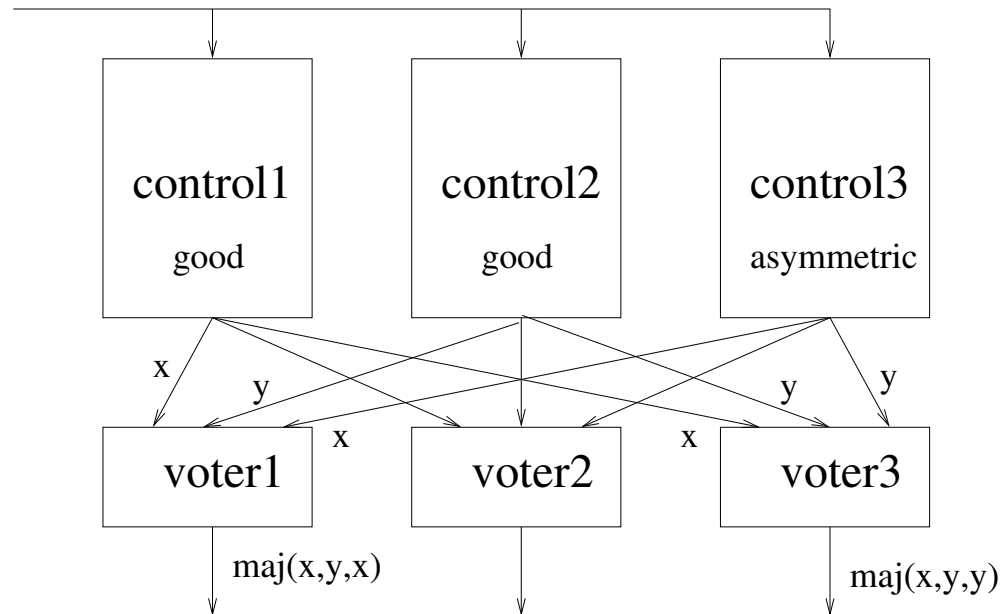
Single Point of Failure

An FCU the failure of which may entail system failure.

Example: anything, including voters, cords, and power supplies



Replicated Voters



An asymmetric node may prevent good voters from agreement.
An asymmetric node may prevent good voters from synchronizing.



The SPIDER Project

Scalable **P**rocessor-**I**ndependent **D**esign for **E**lectromagnetic **R**esilience

- a general-purpose fault-tolerant broadcast architecture
- research platform to explore recovery strategies for HIRF/EMI induced faults

Project Goals

- develop a case study application of DO-254
- provide feedback on problem areas
- provide material suitable for DO-254 training



Concept

Palumbo's fault-tolerant processing system (U.S. Patent 5,533,188).
Part of Fly-by-Light/Power-by-Wire project

Related Designs

- FTTP, FTP, FTMP (Draper Labs)
- MAFT (Allied Signal)
- SIFT (SRI)
- TTA (TTTech)
- SAFEbus (Honeywell)



The ROBUS


Reliable Optical BUS

- SPIDER hardware
- connects processing elements (PEs): computers, sensors, actuators
- offers basic fault-tolerant services

Effectively, the system around the ROBUS may assume less severe fault modes.



ROBUS Architecture






Two Laboratory Prototypes



(1) by NASA Langley



(2) by Derivation Systems, Inc.



The ROBUS Protocols

- Message Passing
- Clock Synchronization
- Fault Diagnosis
- Rejoin
- Startup
- Restart (Never Give Up)

All protocols except Restart depend on

- synchronous clocks (responsibility: clock sync)
- FH (responsibility: diag and rejoin)



PVS Code

Formal Verification (= machine-checked rigorous mathematical proofs)
in SRI's PVS prover tool.

Current *version2* contains full PVS proofs of 3 protocols.

- 6686 lines of PVS code
- 62183 lines of proof script



Conclusion

- Fault tolerance: sticky business; issues understood
- Ultra-reliability: impossible to validate by simulation and testing; formal verification necessary

Future Work

- Fault-tolerance library in PVS
- Prove correctness of rejoin in PVS

On the Web

NIA Formal Methods: <http://research.nianet.org/fm-at-nia/>