An Immune System Paradigm for the Design of Fault Tolerant Systems

Algirdas Avizienis

A. Avizienis and Associates Inc.
2711 Washington Avenue, Santa Monica, CA 90403 U.S.A.
aviz@cs.ucla.edu

An in-depth assessment of the implementation of fault tolerance in contemporary "off-the-shelf" computing systems [1] leads us to conclude that hardware defenses are not adequately exploited for the assurance of dependability. In the search for a fundamentally better solution we have looked at the self-protection (i.e., fault tolerance) mechanisms of the human being. We use two analogies [2]: (1) the body is analogous to hardware; and (2) the cognitive processes of the mind are analogous to software. The immune system of the body is a major protective mechanism that is completely independent of the cognitive processes. It functions from conception until death of the body and protects the body of an unconscious or sleeping human equally well as that of a conscious one.

The solution that we have proposed is to insert into a given "host" computing system a hardware subsystem called the FTI that is analogous to the immune system of the human body [3]. We call this approach to building dependable systems the "immune system paradigm" (ISP). The ISP is a set of design principles for a software-independent and fully fault-tolerant implementation of the FTI.

To develop the ISP we identify the key properties of the human immune system and from them derive the attributes that the FTI must have in order to satisfy the analogy with the immune system. There are four attributes of the immune system that are especially relevant [4]:

1. It functions (i.e. detects and reacts to threats) continuously and autonomously, independently of cognition.
2. Its elements (lymph nodes, other lymphoid organs, lymphocytes) are distributed throughout the body, serving all its organs.
3. It has its own communication links – the network of lymphatic vessels.
4. Its elements (cells, organs, and vessels) themselves are self-defended, redundant and in several cases diverse.

The properties that the FTI must possess to justify the immune system analogy are:

1. The FTI consists of hardware and firmware elements only.
2. The FTI is independent of (requires no support from) any software of the host platform, but can communicate with it and support its recovery.
3. The FTI supports (provides protected decisions algorithms for) multichannel computing of the host platform, including diverse hardware and software channels to provide design fault tolerance for the host platform.
4. The FTI is compatible with (i.e., protects) a wide range of host platform components, including processors, memories, supporting chipsets, discs, power supplies, fans and various peripherals.

5. Elements of the FTI are distributed throughout the host platform and are interconnected by their own autonomous communication links.
6. The FTI is fully fault-tolerant itself, requiring no external support. It is not susceptible to attacks by intrusion or malicious software and is not affected by natural or design faults of the host platform.

An FTI that possesses the above attributes and supports host platforms that use Intel's P6 family of processors has been described in [3]. It is a generic, hierarchical, fault-tolerant (f-t) hardware infrastructure that serves as a software independent innermost defense for error detection and recovery of a platform that may also employ various other fault tolerance and security techniques.

The hierarchical structure within the FTI is as follows: an f-t set of S^3 (Startup, Shutdown, Survival) nodes protects an f-t set of M (Monitor) nodes, which in turn protects f-t A (Adapter) and D (Decision) nodes that are connected to and protect the components (C-nodes) of the host platform (see figures 2 and 4 of [3]). A new concept that is being explored is a "hierarchy of infrastructures:" an FTI is installed within each chip of the host platform, another FTI protects one board, and still another FTI protects the entire platform. The on-chip FTI performs the A-node and D-node functions for the board FTI, while the board FTI does the same for the platform FTI. The principal constraint in developing a hierarchy of FTIs is the need for dedicated and protected communication links within each FTI and between FTIs. This constraint makes the further extension of FTI hierarchy to clusters and LANs of platforms relatively costly to implement.

The goal of the Immune System Paradigm is to use hardware more extensively and more effectively than it is being done currently in providing fault tolerance for very dependable high-performance platforms. A benefit of the FTI is the ability to simplify higher-level defenses that require software participation. The presence of an effective FTI simplifies the error detection and recovery requirements for system software.

In concluding we predict that adoption of the FTI in platform designs will lead to a better structured and more cost-effective overall dependability assurance architecture, since the other levels of protection will be supported by hardware that is missing in today's designs.

References

1. A. Avizienis and R. Avizienis. An immune system paradigm for the design of fault-tolerant systems. Presented at Workshop 3: Evaluating and Architecting Systems for Dependability (EASY), in conjunction with DSN 201 and ISCA 2001, Goteborg, Sweden, Jul 1, 2001. Available at: http://www.crhc.uiuc.edu/EASY/easy01-program.html
2. A. Avizienis. Toward systematic design of fault-tolerant systems. Computer, 30(4)51-58. April 1997.
3. A. Avizienis. A fault tolerance infrastructure for dependable computing with high-performance COTS components. In Proceedings Of the Int. Conference on Dependable Systems and Networks (DSN 2000), June 2000, pp. 492-500.
4. G.J.V. Nossal. Life, death and the immune system. Scientific American, 269(33)52-62, September 1993.