

Security Without Compromise by Steve Schuster

INTRODUCTION

Pervasive network connectivity increasingly forces engineers to confront and solve a wide range of technical security and reliability issues. Very few product classes remain that are *not* racing to feature Internet accessibility, and the resulting vulnerability has become an inevitable challenge. The more embedded devices operate in open network environments, the more security becomes a horizontal technology issue for virtually every embedded design. But embedded system developers have classically faced steep cost, and time-to-market compromises in exchange for ensuring consistently reliable operation and security.

This paper uses automotive telematics to illustrate the need for embedded security and reliability assurance, and explains how the blending of hardware and software technologies from Analog Devices and Green Hills Software eliminate the previously inevitable compromise between security, cost, and time-to-market. The result is a high reliability, high security, memory-protected RTOS running on a high performance, low cost, low power, memory-protected processor complemented by a development environment that facilitates fast time-to-market.

Bringing Reliability and Security to the Open Car

For embedded systems, security is directly coupled with system reliability, as opposed to the data integrity and privacy issues relevant to enterprise computing applications. An unsecured embedded device is inherently an unreliable device (and vice versa), resulting in consequences that range from minor nuisance to grave physical danger. And as systems and networks grow in complexity, techniques used by those who would aim to compromise security also become more sophisticated. In a typical modern car, for example, there can be on the order of 50 embedded control units interconnected via possibly five different communication bus technologies (e.g., SAEJ1850 and CAN), all of which are fair game for errant code or malicious tampering. Worse is that with wireless networks, physical contact with the internal buses or the car itself is not even required to involve externally generated corruption, either unintentional or malicious.

Automotive telematics, which are intrinsically embedded systems, are inevitably evolving to integrate a wide range of communication, information, navigation, and entertainment functionality. Such feature-rich systems include telecommunication functions that originate or end inside automobiles,

effectively exposing critical control systems to the outside world. A vision of a network-connected car has naturally motivated innovators, but as a potential target for wrongdoers or sloppy data transmissions, the networked vehicle is a life-threatening nightmare.

To address these risks, next-generation vehicles demand underlying hardware and software that can provide the most rigorous security, operational reliability, and real-time performance. Anything less compromises the safety of these vehicles. But even with quality and safety front and center for carmakers, the automotive marketplace is a cost-driven environment where bill of materials and profit margins are key concerns for manufacturers, sometimes even at the expense of innovation. In the 1980s, for example, initial sales of optional airbags were so poor that General Motors pulled them off the market. Later, however, costs came down and price-elastic consumer demand changed, and now every new vehicle has airbags.

Clearly there can be no waiting around for cost miracles to occur before cars and other safety-critical embedded systems advance to become uncompromisingly reliable and secure. The potential for catastrophic system failures is otherwise simply too great.

Protect Memory to Protect the System

The risks are not only severe, but also wide in scope. In addition to passenger cars, telematics systems are being implemented for a wide variety of intelligent systems, including trucks and buses, off-highway and off-road vehicles, and passenger and cargo trains. Other applications in other markets that are at risk due to network connectivity include maritime electronics, aircraft and aerospace electronics, factory automation, industrial machine control, lifts and escalators, building automation, medical equipment and devices, and nonindustrial control and non-industrial equipment. As previously stated, the security and safety implications of having so many devices open to possibly unauthorized access or software errors are enormous.

Intruders must be kept out, and software must be compartmentalized to segregate software bugs from doing broad, system-wide damage. Improperly segregated, for example, an automotive infotainment system could erroneously broadcast control-signal garbage onto a vehicle's CAN bus from which vital control modules take their marching orders.

Green Hills' INTEGRITY RTOS helps solve that problem by isolating resources at the software level. In turn, the Blackfin® processor's performance and memory management capabilities combine to allow the INTEGRITY RTOS to be applied in automotive applications while ensuring the cost is low enough for market viability. A true convergent processor that executes microcontroller and DSP processes equally well, the Blackfin processor brings empirically exceptional performance/price and performance/power. Equally crucial is the Blackfin processor's excellent suitability for hosting a secure and reliability-centric RTOS like the INTEGRITY RTOS.

The INTEGRITY RTOS is a memory-protected operating system with safe and secure access controls built in from the ground up. Meanwhile, the Blackfin processor's ability to establish protected memory spaces perfectly complements the INTEGRITY RTOS's conservative, permission-based ownership approach that only allows program objects to be shared between protected address spaces on demand. The resulting compartmentalization of program objects facilitates a highly reliable systems environment that shuts off inadvertent (e.g., a wayward address pointer) or deliberate (i.e., hackers) attempts by one segment of software to detrimentally affect another. The INTEGRITY RTOS also allows embedded systems developers to map performance-critical subsets of code onto the Blackfin processor's L1 cache, resulting in an order-of-magnitude faster context switch time for real-time processes.

The Blackfin processor contains a page-based memory management unit (MMU), which provides control over cacheability of memory ranges and management of protection attributes at a page level. The MMU uses a memory protection format that, when coupled with the core's User and Supervisor modes, can support a full RTOS. The RTOS runs in Supervisor mode and partitions blocks of memory and other system resources for the actual application software to run in User mode. Thus, the Blackfin MMU offers an isolated and secure environment for robust systems and applications. Honing the combination to perfection, the breakthrough low price point of the Blackfin processor (around \$5.00) for 400 MHz of performance allows the deployment of highly advanced embedded applications at market-supportable material costs that were previously unattainable.

To complete the picture, the MULTI integrated development environment (IDE) from Green Hills gives embedded software developers a direct view into the INTEGRITY RTOS target on the Blackfin processor. This allows quick and effective debugging that leads to far quicker (and cheaper) time-to-market schedules. The MULTI IDE also contains key tools and configurations needed for efficient design. These include an RTOS-aware source-level debugger, run-time error checker, version control system, performance profiler, real-time event analyzer, and the simulated ability to develop and test Blackfin processor code on a PC, or Linux, or UNIX workstation in the absence of actual target hardware.

A Leg Up on Security

The powerful combination of the Blackfin processor, the INTEGRITY RTOS and the MULTI IDE form an unprecedented three-legged platform upon which developers can bring reliable and secure embedded systems to market very quickly and at low cost. The outcome is the end of the classic compromise—now product developers can have security, time-to-market, and low-cost at the same time. And for users of the resulting products, the world becomes a safer and more reliable place that remains within financial reach.