Position Paper: Deeply Embedded Survivability

Philip Koopman, Jennifer Black, Theresa Maxino
Carnegie Mellon University
{koopman, jenm, maxino}@cmu.edu

Abstract

This position paper identifies three significant research challenges in support of deeply embedded system survivability: achieving dependability at the enterprise/embedded interface gateway, finding a viable security patch approach for embedded systems, and surviving run-time software faults.

1. Introduction

Deeply embedded systems consist of one or more embedded systems connected to an enterprise system or to the Internet (e.g., [3]). To be survivable, such systems must continue to function in the face of faults, whether accidental or malicious, and whether the faults are caused by design errors or unexpected operating conditions. Embedded system survivability can be more challenging than enterprise survivability because embedded systems may not be able to perform frequent reboots, incorporate weekly patches, transfer large amounts of data, or be cared for by trained system administrators. Beyond this, the different natures of embedded control vs. enterprise systems present fundamental limitations to applying known techniques from either area to the other. [1]

2. Fundamental limitations

2.1 Time triggered to event triggered interfaces

A fundamental limitation to achieving deeply embedded system survivability is the inherent mismatch between time triggered and event triggered systems.

Embedded systems are often "time triggered," meaning that they perform periodic computations and messaging in support of hard deadlines (e.g., [2]). Because of the dramatically different needs of real time control systems compared to desktop computing, they often use specialized network protocols such as CAN that provide low-cost, but low-bandwidth solutions optimized for very short messages (often 100 bits or fewer per message with network speeds on the order of 1 Mbit/sec).

Enterprise systems, in contrast, are usually characterized as "event triggered" systems with much larger, sporadic events, and typically have orders of magnitude more CPU power and network bandwidth.

The interface between the embedded and enterprise sides of a deeply embedded system is usually in the form of a "gateway" that provides a bidirectional transition between the time triggered and event triggered worlds. Given sufficient resources, each computing paradigm can be made to simulate the other. Event triggered systems can schedule events periodically to simulate time triggered operation. Time triggered systems can schedule periods so fast that they don't miss events. But, those approaches only work in the fault-free case.

Deeply embedded system gateways will encounter fundamental limitations when attempting to map faults and responses in one computing paradigm into the other computing paradigm. For example, what happens when event triggered messages are clumped in transit, and arrive faster than the minimum inter-arrival rate assumed by the time triggered side of the gateway? Queues in the gateway provide only a partial solution, and can cause problems when the system encounters queue overflow or system instability as a result of queue lag time.

In the other direction, time triggered messages that contain too much value jitter can defeat whatever low pass filters are in place at the gateway and can potentially flood the enterprise system with messages. Leaky buckets and other throttling methods can provide some relief, but are not necessarily able to do the right thing in those cases where an event shower is representative of a true emergency situation rather than a fault or attack.

Despite a lack of understanding of these fundamental issues, deeply embedded system gateways are already being deployed, sometimes in critical systems.

2.2 Limits to the patch mentality

The approach of using security patches to address emergent attacks is pervasive in the desktop computing environment. Embedded systems have fundamentally different constraints that make patching difficult.

Safety critical systems must be recertified each time critical software is updated. Doing so is usually a costly and time-consuming process. Quick-turnaround security patches are currently impracticable if they affect critical code. Unfortunately, many embedded systems are designed in such a way that all their code is effectively critical (i.e., any change to the code might affect critical properties, so it must all be assumed to be critical).

Strategies to isolate critical from non-critical software on the same CPU are still a subject of research.

An additional issue with patching embedded systems is that many of them have a zero down-time requirement. Maintenance reboots and physical operator intervention are simply unacceptable in many unattended applications.

Finally, patching approaches typically assume that the owner of a system is trustworthy. This is often not the case in embedded systems. For example, it is relatively common for sports car owners to install engine controller software that circumvents pollution emission and fuel economy controls as a way to get more performance.

2.3 Limits to the perfect software mentality

Much research in computer science is based on the laudable goal of creating perfect software. Industry practices also employ the assumption that "perfection" (or a close approximation thereof) can be achieved by identifying all the "important" bugs and removing them.

In the real world, very few application domains have the time and resources to deploy low defect rate software. Getting the highest software quality possible within time and budget is certainly important. But, spending exponentially increasing resources to chase down the last few bugs is usually impractical. Instead, it might make more sense to spend a small fraction of available resources providing ways to survive bugs that will inevitably be encountered, rather than throwing all resources at an attempt to achieve absolute perfection.

3. Research challenges

There are several research challenges that stem from the limitations just discussed. They are:

Understand what goes into the embedded/ enterprise gateway. While some combination of queues and message filters can work in the fault-free case, mapping fault manifestations and survivability mechanisms across the time triggered to event triggered interface provides fundamental research challenges.

Make patching of critical embedded software viable. Patching of unattended, critical embedded systems provides fundamental challenges that aren't encountered in most desktop systems. Creating patching approaches that maintain system integrity promises to be difficult.

Increase system survivability by tolerating inevitable software defects. Software defects are inevitable in most fielded systems. In some cases these defects will result in security vulnerabilities. In others they will result in failures to maintain critical system properties. Making software faults more survivable could offer improved cost effectiveness and reduced system fragility.

4. Promising innovations and abstractions

4.1 Safety invariants

Safety invariants, which are formal expressions of critical system properties that must hold true, offer new promise for increasing system survivability. Traditionally, analysis and testing are used to ensure the invariants are never violated. But, these techniques only work for the systems that are modeled (which are usually fault-free systems). One could also check safety invariants at run time to detect when a fault has occurred that is severe enough to compromise system safety. Safety invariant checks could act as failure detectors that activate recovery or safe shutdown mechanisms.

4.2 Graceful degradation

The term graceful degradation encompasses several meanings. The term was coined to describe modular redundancy in fault tolerant computing, and later evolved to encompass failover strategies and functional diversity. More recently, the term has been used to describe performability tradeoffs in Quality of Service research. The notion of providing systems that can partially work rather than only be fully working or fully failed is essential to achieving cost-effective survivability.

5. Possible Milestones

Survivability is an emerging research area, with the current emphasis more on understanding fundamental problems rather than on comprehensive solutions. Long-term milestones should include discovering fundamental tradeoffs, impossibility results, and workarounds applicable to realistic systems. Short term research milestones should emphasize characterizing practical limitations and exploring techniques to offer near-term improvement to system builders.

6. Acknowledgements

This work had been funded by the General Motors Collaborative Research Laboratory at Carnegie Mellon University and the Pennsylvania Infrastructure Technology Alliance.

7. References

[1] Koopman, P., Morris, J. & Narasimhan, P., "Challenges in Deeply Networked System Survivability," *NATO Advanced Research Workshop On Security and Embedded Systems*, August 2005

[2] Kopetz, H., *Real-Time Systems: Design Principles for Distributed Embedded Applications.* Kluwer, 1997.

[3] Tennenhouse, D., "Proactive Computing," *Comm. ACM*, 43(5): 43-50, May 2000.

Challenges In Deeply Networked System Survivability

Philip Koopman

February 2007

koopman@cmu.edu

http://www.ece.cmu.edu/~koopman

Electrical & Computer ENGINEERING

Carnegie Mellon

1

Overview

- **Brief introduction to the world of embedded control**
  - To a first approximation, desktop CPUs are 0% of the market
- **High Level look at two issues**
  - Embedded / Internet Gateways
  - An example threat: household thermostats

2

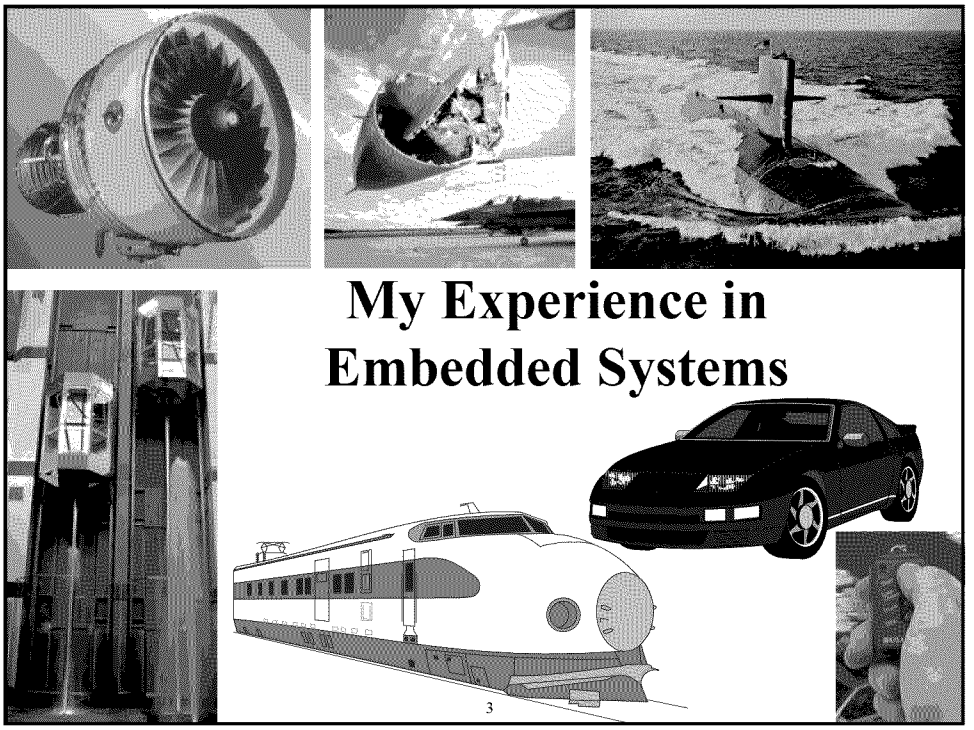

How Many CPUs In A Car Seat?

- **Car seat photo from Convergence 2004**
  - Automotive electronics show

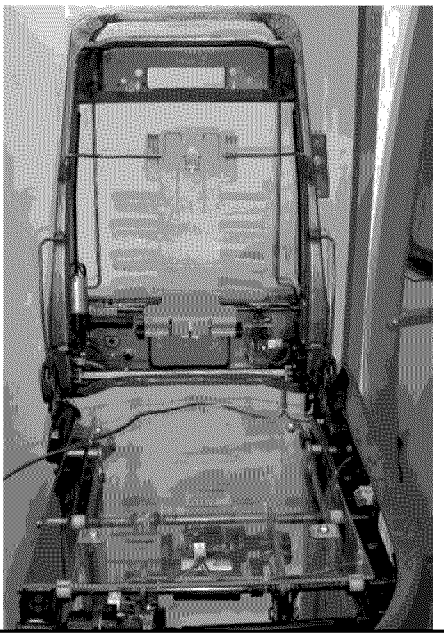

Car Seat Network (no kidding)

- **Low speed LIN network to connect seat motion control nodes**
- **This is a distributed embedded system!**
  - Front-back motion
  - Seat tilt motion
  - Lumbar support
  - Control button interface
  - Connects to body controls network beyond seat for per-driver customization

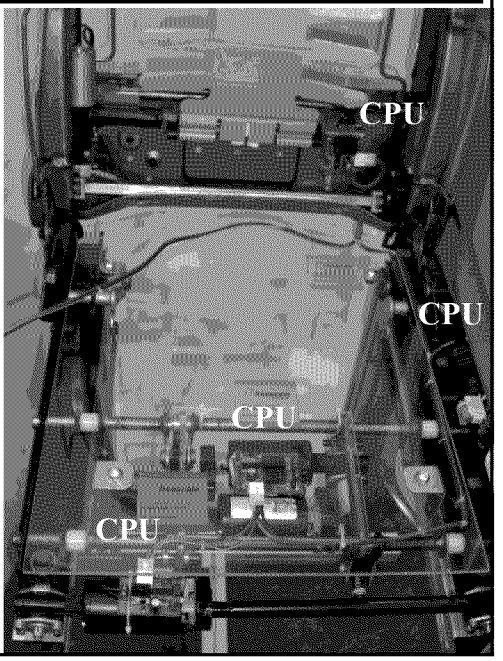

Microprocessor Unit Sales

All types, all markets worldwide

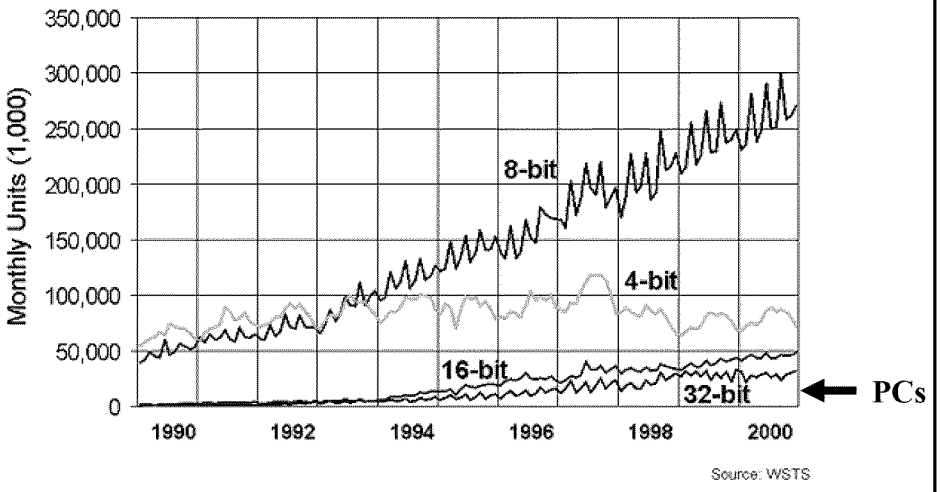

15 Million PCs per month in 2004 (15,000 on this graph)

Trend: External Connectivity

- **Safety critical subsystems will be connected to external networks (directly or indirectly)**
  - German proposal:
    wireless networks control car's max. speed
  - E-enabled aircraft architecture (next slide)

[Airbus 2004] A-380 scheduled to enter service in 2006

7

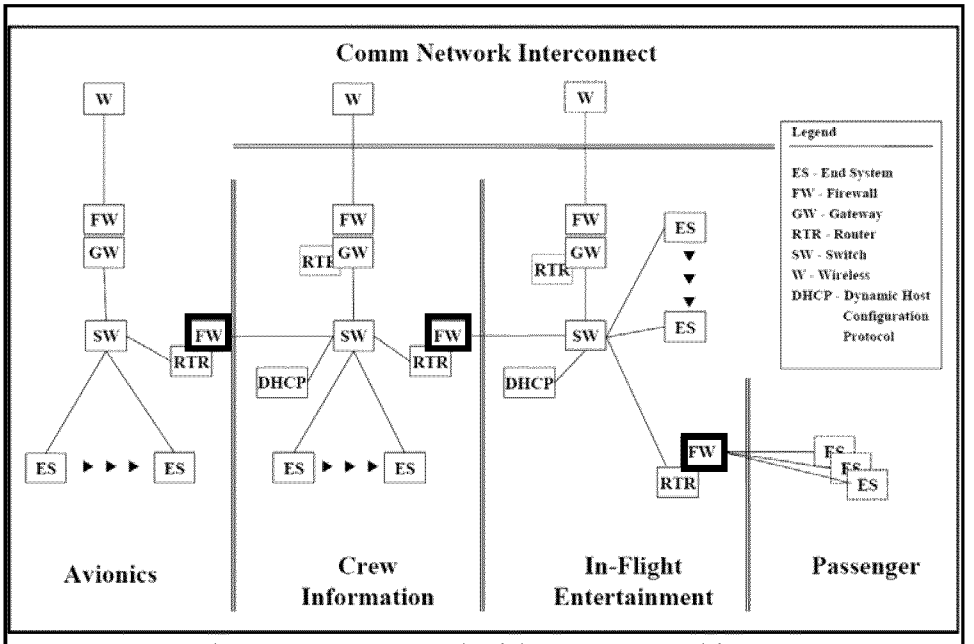

Wargo & Chas, 2003, proposed Airbus A-380 architecture

Passenger laptops are 3 Firewalls away from flight controls!

Deeply Embedded System Gateway

Enterprise system + Embedded System =
"Deeply Embedded System"

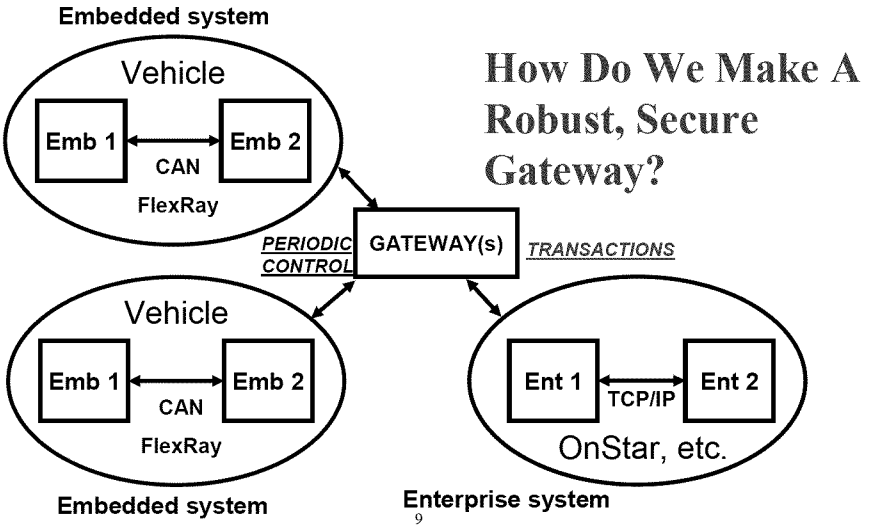

9

Research Area: Embedded/Internet Gateway

- **What happens at the embedded/internet interface?**
  - Fault propagation across the gateway presents fundamental challenges

Embedded Side	GATEWAY	Enterprise Side
Control-oriented		Transaction-oriented
Time Triggered		Event Triggered
Continuous		Discrete
Real Time		Mostly not Real Time
Periodic Messages		Aperiodic Messages
Short Messages		Longer messages
Roll-forward		Rollback
Lower cost		Higher cost

10

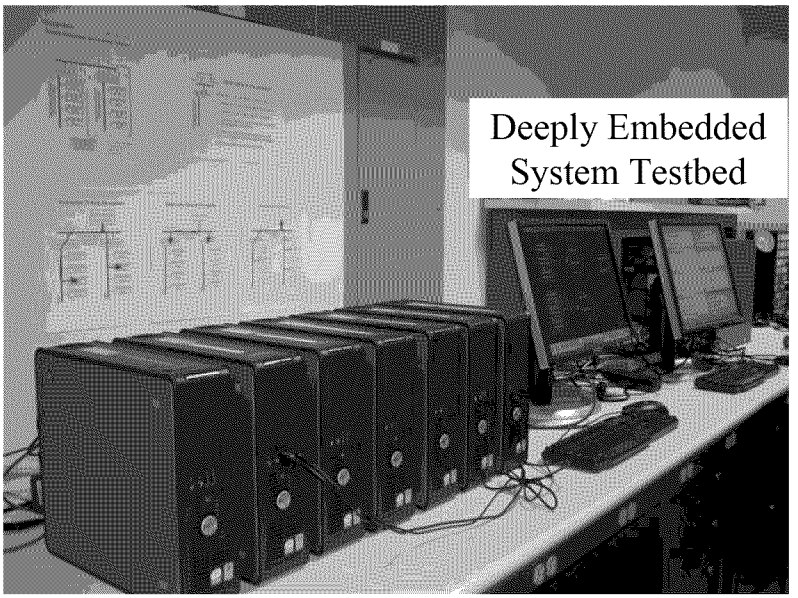

Initial Experiment: Queue overflow

- **How having a long queue can cause you to operate on stale data**

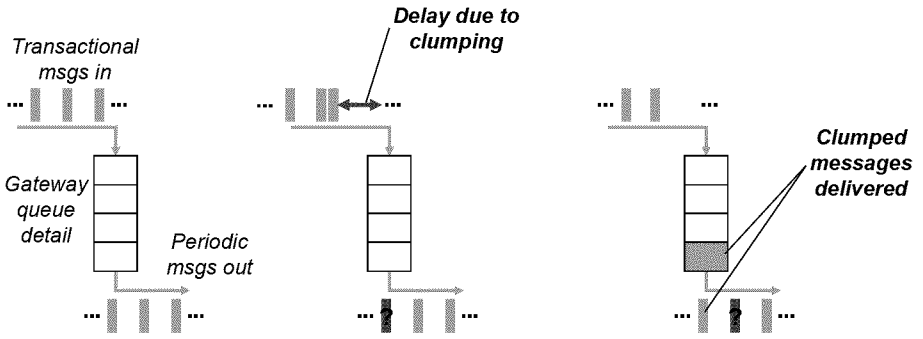

Ideal case: Queue is empty in the steady state

Clumping delay leads to missed deadline for periodic messages

Need a policy for dealing with having no message to send.

Two messages delivered, so a message is stuck in the queue.

Now all the messages delivered to the embedded system are 1 period old.

12

Deeply Embedded Scary Scenario

- **Consider the lowly thermostat**
  - Koopman, P., "Embedded System Security," *IEEE Computer*, July 2004.
- **Trends:**
  - Internet-enabled
  - Connection to utility companies for grid load management
- **Proliphix makes an Internet Thermostat**
  - (But it we're not saying that system has these vulnerabilities!)
  - Somebody else makes one almost exactly like this, deployed July 2005

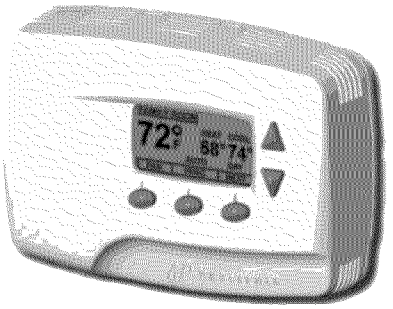

13

Waste Energy Attack

- **"I'm coming home" function**
  - Ability to tell thermostat to warm up/cool down house if you come home early from work, or return from a trip
  - Save energy when you're gone; have a comfy house when you return
  - Implement via web interface or SMS gateway
- **Attack: send a false "coming home" message**
  - Causes increase in utility bill for house owner
  - If a widespread attack, causes increased US energy usage/cause grid failure
  - Easily countered(?) – if designers think to do it!
    - Note that playback attack is possible – more than just encryption of an unchanging message is required!

14

Discomfort Attack

- **Remotely activated energy saver function**
  - Remotely activated energy reduction to avoid grid overload
  - Tell house "I'll be home late"
  - Saves energy / prevents grid overload when house empty
- **Attack: send a false "energy saver" command**
  - Will designers think of this one?
  - Some utilities broadcast energy saver commands via radio
    - In some cases, air conditioning is completely disabled
    - Is it secure??
  - Consequences higher for individual than for waste energy attack
    - Possibly broken pipes from freezing in winter
    - Possibly injured/dead pets from overheating in summer

15

Energy Auction Scenario

- **What if power company optimizes energy use?**
  - Slightly adjust duty cycles to smooth load (pre-cool/pre-heat in anticipation of hotest/coldest daily temperatures)
  - Offer everyone the chance to save money if they volunteer for slight cutbacks during peak times of day
  - Avoid brownouts by implementing heat/cool duty cycle limits for everyone
- **You could even do real time energy auctions**
  - Set thermostat by "dollars per day" instead of by temperature
    - More dollars gives more comfort
  - Power company adjusts energy cost continuously throughout day
  - Thermostats manage house as a thermal reservoir

16

Energy Auction Attacks – Naïve Version

- **What if someone broke into all the thermostats?**
  - Set dollar per day value to maximum, ignoring user settings
    - Surprise! Next utility bill will be unpleasant
  - Turn on all thermostats to maximum
    - Could overload power grid
  - Pulse all thermostats in a synchronized way
    - Could synchronized transients destabilize the power grid?

17

Energy Auction Attacks – Scary Version

- **What if someone broke into the energy auction server?**
  - If you set energy cost to nearly-free, everyone turns on at once to grab the cheap power
  - Guess what – enterprise computer could have indirect control of thousands of embedded systems!
  - Someday soon, almost "everything" will be "embedded," at least indirectly

18

"Unique" Embedded System Requirements

Embedded systems:

- **Are actually supposed to work**
  - Do you want to perform a workaround for your water heater?
  - Often have 24x7 requirements – zero down time
- **Often are safety critical**
  - Have you ever ridden in a fully automated train/peoplemover? (or an elevator?)
- **Are very cost sensitive & resource constrained**
  - A $0.50 CPU can't run a "big" OS with full security features
- **Don't have a sysadmin**
  - Who's the sysadmin for your DVD player?
  - The owner is often negligent, or even a malicious attacker

19