OPTIONS FOR UPGRADING LEGACY AVIONICS SYSTEMS

Russ Duren, Naval Postgraduate School, Monterey, CA

Abstract

This paper examines a comprehensive set of options for upgrading legacy avionics systems in tactical aircraft. The options have been developed as part of study involving the avionics system of a current military aircraft. Multiple options are considered ranging from minor changes to the legacy system to complete replacement of computers and communications buses with commercial off-the-shelf (COTS) hardware and new software. The pros and cons of each potential solution are described and guidelines are given for choosing the appropriate solution for a given program. The results of this study may be used when evaluating upgrade options for a wide range of embedded computer applications.

Introduction

The avionics system in the majority of military aircraft is based on a federated architecture [1]. In a federated architecture, functionality is partitioned among multiple subsystems. Each subsystem performs a particular function such as inertial measurement, communications, or radar. The subsystems are interconnected and share data over standardized serial data buses, which are predominantly MIL-STD-1553. Overall control of the avionics system is delegated to a central computer typically referred to as the mission computer.

The current inventory of military aircraft contains avionics systems that are 20 to 30 years old [2]. The lifetime of many of these aircraft has already exceeded their original design lifetimes and may be extended another 20 years or more into the future. As these systems age, they require subsystem replacement or upgrades to remain viable. These modifications are required as a result of problems that are common to other "legacy" embedded computer systems, including:

- **Memory Limits** – The software that runs in each computer is referred to as the operation flight program or OFP. The computer may be out of memory to support OFP growth. OFP growth in military avionics is driven by a number of factors including the addition of new mission-related subsystems such as weapons, sensors, and data links; changes in tactics requiring additional functional modes to be implemented in software; and the addition of equipment to meet changing requirements for operation in civil air space.
- **Processing Speed** – Avionics subsystems are real-time systems required to meet strict deadlines. As the OFP grows, the processors may not be fast enough to meet the required deadlines.
- **I/O capability** – The addition of new subsystems such as sensors or weapons may require new I/O capability.
- **Communications bandwidth** – As functions are added to each of the federated subsystems, increasing demands are placed on the communications network. This is exacerbated by the rapid growth in the capabilities of related systems including high resolution sensors and video displays, weapons with image-based guidance systems, and high-speed data links. While the MIL-STD-1553 data bus excels for communicating control, status, and small amounts of data, it was not designed to be used for high data rate applications such as digital video.
- **Communications addressing limits** – MIL-STD-1553 allows at most 32 terminals to be interconnected. Each terminal can define 960 unique 16-bit data words to be transmitted or received. These restrictions limit data transmission

between subsystems and can hinder the addition of modern weapons and new avionics subsystems to an aircraft.

- **Display Limitations** – The displays on many legacy aircraft were designed for text and monochrome graphics with significantly lower resolution than required by modern applications. As capabilities have been added to existing aircraft, these displays have reached limits in resolution, memory, graphics processing speed, and even the writing speed of the beam in stroke type displays. Updating to high resolution, high bandwidth color displays may require a large increase in the computation throughput and communications bandwidth of the associated computers.
- **Parts Obsolescence** – The hardware may need repair or it may need to be replaced as a result of reaching the end of its mechanical life resulting in an increasing and unacceptable failure rate. If the original components are no longer in production and cannot be replaced by any other means, then it will be necessary to redesign the hardware using new components.
- **Software Supportability** – Software support agencies are finding it increasingly difficult to hire and retain personnel to support OFPs written in lesser-used languages including Ada and assembly languages. Outdated software development tools add to this problem.

In response to this problem a study has been performed to assess available options for upgrading legacy avionics systems. Multiple options have been identified that vary in their cost and in the amount of improvement they offer. The choice of which options are best is highly dependent on the circumstances of the particular avionics system. This paper examines the both available options and the factors having the greatest influence on the final solution. Examples are provided to illustrate how the best choice is influenced by these factors.

Constraints

Upgrades to legacy systems face different constraints than the design of a new system. These constraints have a large influence on the final decision of how to upgrade the system. They include:

- **Budgetary Constraints** – The initial design and procurement of an avionics system is very expensive, representing up to 40% of the procurement cost of a modern tactical aircraft. The budget available to update a legacy avionics system is a typically small in comparison, limiting the scope of any upgrade. An upgrade program has to compete for budget with ongoing maintenance and operational needs. If too much money is spent on upgrading the computers, there may not be enough money left to pay for integrating new capabilities to take advantage of the upgraded computational performance.
- **Software Investment** – Software development can represent up to 80% of the acquisition cost of an avionics subsystem [3]. Software maintenance occurring after the initial procurement represents 60% to 80% of the life cycle cost of the software. The magnitude of this investment deters upgrade options that require the software to be replaced.
- **Performance Verification** – Any modification that impacts software, including the replacement of a processor with an allegedly binary-compatible upgrade from the same manufacturer, can require significant testing to verify that the change has not had a negative impact on safety-critical and mission-critical subsystems.
- **Minimize Wiring Changes** – It is often cost prohibitive to run new cables in a tactical aircraft. Therefore any changes to an avionics subsystem will often be constrained to work with the existing wiring. Electromagnetic compatibility concerns and the need to work with legacy communications protocols such as MIL-STD-1553 limit upgrade options.

- Minimize Impacts to Associated Legacy Systems – Even new designs must be compatible with their operational and maintenance environment. However, legacy upgrades may find these requirements more restrictive, as the legacy environment will typically impose limits, such as bandwidth restrictions, that eliminate certain options.

Options

A wide range of options is available for upgrading a legacy system. The options considered in this study ranged from minor modifications of existing subsystems, to replacement of existing subsystems with new components. The study

addressed the addition of new displays and new subsystems, but only considers effects the additions would have on the computational and communications requirements of the avionics system. The upgrade options can be classified into two categories. Figure 1 summarizes upgrade options for increasing the computational capabilities of individual subsystems. Figure 2 summarizes options for upgrading the MIL-STD-1553 communications system. In both figures, the cost impact tends to increase from the top to the bottom of the figure. However, the relative costs of various options are very sensitive to the details of the individual aircraft and avionics system being considered.

Figure 1. Options for Addressing Processing Limitations

Addressing Legacy Processing Limitations

Options for addressing processing limitations vary considerably. If only limited improvements in processing throughput or memory are needed for the projected lifetime of the aircraft, they may be obtainable through limited reengineering of the

existing OFP. Reengineering can be applied to an individual system to reclaim throughput and memory lost to inefficiencies inherent in the software maintenance process.

Every time an OFP is modified the program structure is weakened. Code may be left in the

program that supports functions that are no longer needed such as obsolete weapons and tactics. Reengineering the OFP provides the opportunity to eliminate obsolete functions and optimize the

program structure for the present functionality. If reengineering is done on a small scale, the cost may be acceptable.

Figure 2. Options for Addressing Communications Limitations

In addition to reengineering within a single OFP, reengineering can be applied to repartition functions between OFPs. Software maintenance may have impacted one subsystem more than others. If one subsystem has reached processing limits, it may be possible to shift some tasks to more lightly loaded subsystems. When tasks are repartitioned between processors, care must be taken not to adversely impact the communications system. Tasks that require significant data sharing should not be separated.

In some redundant systems, two or more processors execute identical copies of an OFP in order to provide complete redundancy. In such systems it is possible to split the functionality between the redundant computers. As not all tasks are performed on each computer, both memory and processing time are reclaimed. To provide graceful degradation, each box typically retains a minimal subset of redundant functionality. This has two negative effects on the overall system. One is that loss of any computer will cause loss of some functionality. The other is that the repartitioning of software functionality usually requires increased communication between the computers, which

increases the demands on both the processor and the data bus interconnecting the computers.

Reengineering legacy software may provide the benefit of avoiding changes to hardware. However, due to the large investment in legacy software it may be more economical to find options that do just the opposite, upgrading the hardware while leaving the legacy software unaffected. In some subsystems the capability exists to add additional legacy processors and memory, although this is not typical. However it is possible to redesign one or more circuit cards to provide both a faster version of the processor and more memory. Several companies involved in legacy avionics offer services to perform such upgrades. CPU Technology, Inc. in Pleasanton, California develops application specific integrated circuit (ASIC) based compatible processors, dual instruction set processors, and complete system-on-a-chip (SOC) products [4]. Titan Systems Corporation in San Diego, California is investigating field programmable gate array (FPGA) based emulations of legacy processors [5]. VP Technologies, Inc. in Marietta, Georgia offers both FPGA and ASIC based processors [6]. These companies offer a wide

array of services including updated software development tools and software reengineering services in addition to compatible processors.

ASIC and FPGA-based processors can be classified as hardware emulators. Software-based processor emulation is also available. The TRW Dayton Avionics Engineering Center in Kettering, Ohio offers a line of software emulators called RePLACE (Reconfigurable Processor for Legacy Applications Code Execution) [7]. Both types of solutions execute the legacy code with no modification. Dual instruction set processors and software emulators provide a path for incremental migration to a second instruction set.

If changes to both the hardware and software are acceptable, another set of options becomes available. Several options exist that require changes to the legacy software, but maintain the capability of executing the bulk of the OFP unmodified. One such option, available in a few cases, is the addition of circuit cards containing legacy processors and memory using spare slots in an existing backplane. If the OFP is written to support a varying number of processors, then this may be a hardware-only change. If the OFP is written to operate with a fixed number of processors, then the operating system executive must be modified to take advantage of the added processors. This type of upgrade is possible with Navy aircraft that use the AYK-14 Standard Airborne Computer [8]. General Dynamics Information Systems has developed a legacy processor upgrade that retains the legacy processor and adds a separate PowerPC processor with a hardware bridge to coordinate between the two processors. The result is similar to a dual instruction set processor or software emulator except that this method uses two processors, both of which execute at the same time. The hardware bridge is required to connect the legacy processor and bus to the PowerPC processor, which uses a VME bus. Software and hardware must work together to handle accesses to common memory areas and I/O.

Current best practice calls for all software to be written in a high-level language using standard operating system interfaces, also known as application programming interfaces or APIs. Properly designed software that makes use of

standard language features, standard data formats, and standard APIs can be ported between host systems with little difficulty. In the best case, the OFP is simply recompiled for the new target hardware. Even in cases where the best practices have been followed, some hardware dependent code is required. Most changes to the hardware will require this portion of the code to be changed. If the hardware-dependent code is well encapsulated, then a combination of rewriting only the hardware-dependent code and recompiling the remaining software may be possible. Unfortunately much of the existing avionics software was written before such standards existed and is typically a mix of high-level and assembly language routines that are intimately tied to the underlying hardware, eliminating this option.

Another method of porting legacy software to new processors is to translate the software such that it can be reassembled or recompiled for the new target hardware. As indicated in Figure 1, translation can be performed at various levels. Binary translation has been used to port software from Intel processors to DEC Alphas [9]. Object code translation has been used to port legacy software from a CISC processor to a RISC processor and to port partially translated object code from a CISC processor to a hybrid processor [10]. In cases where the processor architectures are similar, such as Intel's 8080 and 8086 processors, translation from one assembly language to another has been performed. Translation from assembly language to a high order language such as C has been performed in several cases. Translation from obsolete high order programming languages, such as Pascal and JOVIAL, to current languages, such as C, has also been used [11, 12].

The advantage of translation is that it preserves the investment in legacy software while porting the code to a more supportable and capable hardware platform. However, there are several potential pitfalls. Many problems exist that prevent 100% of the legacy code from being translated automatically. Some manual translation and code development is required in virtually all avionics applications. The code generated by all of these methods tends to be longer than the legacy code. Translations at similar levels, binary-to-binary or HOL to HOL, can double the size of the translated

code due to the need to accommodate differences in architectures and languages. Translations from assembly to HOL require compilation back to a machine language. Preliminary efforts with an existing Navy OFP show a code growth of eight times the original length. Finally, translated code can be harder to maintain, as the resulting source code may be difficult for programmers to decode. Also, translation does nothing to modernize the structure of the software. In particular, translation cannot transform hierarchically organized, functional code to object-oriented code.

Some form of large scale software reengineering must be applied to restructure the software if such improvements are needed. Several approaches to restructuring avionics OFPs have been demonstrated. One approach encapsulates legacy code in software wrappers that can be treated as objects [11]. Code within the wrappers can be emulated, translated or rewritten as desired. Another approach automatically documents the existing structure in graphical form and translates legacy code to ease manual restructuring of the legacy OFP [12]. These approaches tend to be more expensive to implement than those that have been discussed previously, but less expensive than totally rewriting the OFP.

The final approach to updating legacy avionics subsystems is to replace the existing subsystem with completely new hardware and software. The cost of this approach is typically much higher than any other. It can be mitigated somewhat if an existing subsystem from another type of aircraft can be adapted to the aircraft in question. Some of the cost of adapting equipment from another aircraft may be offset by the reduction in support costs for both planes considered as a whole.

Addressing Legacy MIL-STD-1553 Communication System Limitations

Processing limits are only one choke point in legacy avionics systems. The communications data bus used to integrate multiple subsystems can also be a major limitation to upgrades. MIL-STD-1553 defines the serial data bus used to interconnect subsystems on most US military aircraft. MIL-STD-1553 is limited by a one megahertz bit rate and a small number of unique addresses for terminals and data words. As a result of these

limitations, existing aircraft typically incorporate multiple 1553 data buses to interconnect avionics subsystems. Six or more separate 1553 data buses are not uncommon. One processor, typically referred to as the mission computer, may control all the buses, or control may be arranged in a hierarchical manner. Despite being designed with multiple buses, requirements for new avionics capabilities often push the limits of the communications system.

Figure 2 lists options for addressing these limitations. It should be noted that, with one exception, changes to the communication system always impact two or more subsystems. The only exception is when the bus modification is compatible with the existing protocol and only impacts communications between redundant terminals in a single subsystem.

Many legacy avionics systems were designed so that all of the data transmitted on the 1553 buses is separately addressed and transmitted periodically at a rate between 60 Hz and 0.25 Hz regardless of the state of the mission. However, not all of this data is needed all of the time. For example, a tactical aircraft may have multiple operating modes for tasks such as takeoff, landing, air-to-air combat and air-to-ground attack. By reassigning the data addresses during each operating mode, more addresses may be obtained. By only transmitting data necessary for the current mode, communications bandwidth can be better utilized.

Data compression techniques can also be applied to reclaim communication resources. The type of compression technique that can be utilized depends on the type of data being transmitted. Video compression techniques can be used for image transmission, but are not suitable for data such as aircraft attitude and position. However delta-coding techniques can be applied. For instance, assume that aircraft attitude is transmitted as 16-bit quantities at a fixed rate. Using delta coding, the full 16-bit aircraft attitude data could be transmitted at a reduced rate. This data would be supplemented with 4-bit quantities transmitted at the higher rate that provide the change in the attitude data.

As the legacy system evolves, new functionality and additional processing capability may be added to new and existing subsystems.

These changes may provide opportunities to repartition functions. Such repartitioning may be used to minimize communications requirements between subsystems.

All of the above techniques are compatible with the requirements of the current MIL-STD-1553 standard and have the potential of being implemented without hardware changes. If an exception or modification to the specification is allowable, then several other alternatives may be possible. One of these involves the use of the redundant or standby bus. Most 1553 data buses used in avionics systems are implemented as dual standby redundant buses. Per the standard, only one of the buses is allowed to be active at any time. If an exception is granted to the standard, then these backup buses could be used for non-critical data transmission. Transmission of the non-critical data would only be allowed as long as both buses are fully operational. In case of a failure, non-critical data transmission would cease on the standby bus. Alternatively, a dual standby bus could be converted to two non-redundant buses. This would not require a change to the 1553 standard, but it would probably require changes to specifications for the avionics system that required the dual redundant bus in the original design. It should be noted that the hardware that is used in most 1553 terminals assumes standard operation including dual standby operation. It may not be possible to implement the techniques suggested in this paragraph without a change to the terminal hardware.

The SAE committee in charge of MIL-STD-1553 has considered an option called over-clocking for increasing the bandwidth [13]. Over-clocking allows the use of a higher clock rate on the bus. In order to maintain compatibility with existing systems, command and status words are transmitted at the normal clock rate. The clock rate used for data words transmitted between legacy terminals is also unchanged. However, data words transmitted between modified terminals occur at a clock rate of two to four megahertz. Legacy terminals will ignore the higher rate data. This technique allows terminals to be upgraded on an incremental basis. The main problem with such a scheme is that both the 1553 standard and currently available 1553 circuitry do not support such operation. This

situation is not likely to change until some program is willing to pay for the development of the required hardware. Additionally, over-clocking does nothing to increase the addressing limitations of the standard.

Another option that has been considered by the SAE is the use of a frequency-compatible modulation scheme. In particular, it has been shown that a scheme such as digital subscriber loop (DSL) used for high-speed Internet over existing phone lines can be adapted for use on 1553 buses. Transmission rates up to 30 MHz have been demonstrated on standard 1553 data buses in the lab. The frequency ranges of the two systems are far enough apart that they can coexist without modification to legacy systems that do not require the DSL capabilities. This system suffers from the same problem as the previous; no one has fully funded it for implementation.

Frequency-compatible modulation schemes are not limited to transmission over MIL-STD-1553 data buses. Similar schemes have been applied in the commercial world to transmit signals over power lines in homes. The same is potentially possible over power lines or other cables connecting avionics subsystems.

Both over-clocking and DSL options increase the frequency of the signals present on the legacy 1553 bus wiring. However, the wiring was not designed for use at these frequencies. Termination methods, loading, and bus lengths specific to each bus on each aircraft can impact either method, reducing the maximum frequency that can be transmitted without data loss. Radio frequency compatibility with other systems on the aircraft may further limit or preclude the use of these methods. The same holds true for schemes using other interconnections, such as power lines. Although limited testing has provided encouraging results, platform specific testing will be required to determine the applicability of such methods in legacy systems.

If none of the above methods provide sufficient bandwidth, then the only solution is to add a new communication system with new copper or fiber optic cables. This will involve significant expense. Commercial aircraft are increasingly adopting Fast Ethernet for high-speed data transmission. Military aircraft have favored various

implementations of the Fibre Channel standard. The main problem with Fibre Channel is that the standard allows many variations and there is little guarantee that implementations on different platforms are going to be compatible. Additionally, the Fibre Channel standard may not survive long in the commercial world, particularly for interconnecting processors in a switched-fabric system. This could prevent those aircraft that adopt it from enjoying the technology-insertion benefits that are supposed to accompany a COTS standard. Fibre Channel is competing for commercial market place with a host of similar standards including 3GIO, HyperTransport, InfiniBand, Packet-Switched Ethernet, RapidIO, StarFabric, and others [14]. The key lesson here may be to design any MIL-STD-1553 replacement system such that it can easily adapt to changing standards. Field programmable gate array technology can be used to great advantage for such purposes [15].

Choosing The Best Option

The best solution for a particular upgrade depends upon many factors. The result will be determined by the circumstances of the particular avionics system. The development roadmap predicting future additions for an aircraft will determine the requirements for any upgrade. The requirements of the future subsystems will directly determine the subset of options that provide acceptable performance. Within this subset, the cost required to integrate the future systems will influence decisions related to restructuring or replacing legacy software. Communication bandwidth requirements will determine requirements to upgrade or replace data buses. The computational throughput that is predicted to be required to support future subsystems will determine if new computers are needed. If a new computer is needed to support present current requirements, but existing computers are projected to be insufficient for future requirements, then the replacement computer should support future technology insertion.

The projected lifespan for the aircraft directly impacts the life cycle cost associated with every upgrade option. A system with a long projected lifespan will give more weight to insuring that the new system supports future technology insertion.

This favors solutions based on modular design and open systems interfaces. An aircraft that has an intermediate projected lifespan may weigh initial performance more heavily as upgrades may not be a consideration. Finally, a system with a short lifespan that needs only minor increases in performance during its remaining lifetime, might favor repartitioning functions across the subsystems or making minimal changes to one or more circuit cards in a few key computers.

Reengineering OFPs

Reengineering OFPs can be performed without changing the hardware. If properly performed, it has the advantage of improving the structure of the existing software making future maintenance easier. The cost savings provided by reengineering the software may justify the implementation cost, particularly if the roadmap for the aircraft predicts a long life with many future additions impacting the subsystem in question. However reengineering the software without purchasing new hardware does not address component obsolescence. This may be a problem depending on the existing hardware and the projected lifetime of the aircraft.

Reengineering existing OFPs requires consideration of the estimate of spare memory and processing resources that can be reclaimed. Spare memory can have a major influence on the cost of software maintenance. The effect of spare memory was calculated for two E-3 AWACS radars, the APY-1 and APY-2, which respectively were delivered with 9% and 34% spare memory. A three to one cost and schedule difference was documented when making the same change to both radars [3].

Hardware Emulation, Software Emulation, And Translation

The choice between software and hardware emulation is very dependent on both the legacy system and the target technology for the upgrade. The main discriminator between hardware and software emulation is speed. A hardware emulator can be optimized to run the existing instruction set at a rate comparable to that of the current generation of microprocessors. A software emulator must use a fixed instruction set to interpret

the instructions of the legacy processor. If the legacy processor has a complex instruction set that has been highly optimized for the particular application, then software emulation may be very slow. If the instruction sets of the two processors are similar, then software emulation may be very efficient.

Another discriminator between these two options is continuing support of the associated software development tools. Unless the hardware emulation includes a second instruction set (e.g. a dual instruction set implementation) the upgraded system will still be tied to the legacy instruction set. Although the emulator may come with modern upgraded software development tools at the time of the upgrade, these tools will begin to age as soon as they are delivered. If the legacy instruction set or the new development tools are supported in the commercial marketplace, then the tools will continue to be upgraded. Otherwise, the tools will continue to age. This can result in a requirement for additional upgrades in the future, depending upon the lifetime of the aircraft after the upgrade.

Upgrading Redundant Subsystems

When replacing computers in redundant systems, the choice can be made to replace all of the computers or just upgrade one computer. When only one of the computers is upgraded, it becomes the primary computer for the system. The remaining legacy computers become secondary computers. The full functionality of the system will only be available when the primary computer is operational. If the primary computer fails then the secondary computers must provide a critical subset of the functionality. Replacing only one computer in a redundant system reduces the hardware cost for the upgrade. However, it may increase the cost of software maintenance, as separate OFPs must be maintained for the primary and secondary computers. It may also serve as an incremental approach to upgrading redundant systems. One of the computers can be replaced originally. Later, when more money is available, the secondary computers can be replaced with the same upgrade.

Synergies

Whenever a subsystem upgrade is required for one reason, it may be possible to leverage the upgrade to improve system performance in other areas. Increased processing power may provide the opportunity to repartition functions from other overloaded subsystems. Alternatively, spare processing power in two subsystems may be used to implement data compression algorithms to reclaim data bus bandwidth. New I/O boards required due to obsolescence can be redesigned with FPGAs to improve flexibility. Data compression and smart I/O capabilities can be added to I/O boards to reduce both processor and data bus requirements.

Various combinations of the above techniques can be used to provide additional advantages. Combinations of translation and emulation can improve the resulting execution speed. An ASIC or FPGA can be used to implement a faster version of a legacy processor or a dual-instruction set processor. The legacy processor can be captured as a custom HDL design and a COTS intellectual property (IP) core used to provide the second instruction set. Buy capturing the IP for a replacement processor, the problem of future component obsolescence can be mitigated [16]. Upgrades to two subsystems could incorporate data compression, smart I/O, and DSL over the 1553 bus, speeding communications between the two subsystems and relieving the 1553 bus from their traffic. The result of such combinations may extend the capabilities and life of the remainder of the avionics system for little extra cost.

Conclusion

This paper has presented a comprehensive survey of options available for upgrading a legacy avionics system. The options were divided into two categories. The first category contains options for addressing processing speed limitations. Many of the options in this category can also be used to address component obsolescence and improve the maintainability of the associated software. The second category contained options for addressing limitations associated with the MIL-STD-1553 data bus. The best solution will vary with the particulars of each application. Guidelines were presented for consideration when choosing a solution.

References

- [1] August 1993, *Advanced Avionics Architecture & Technology Review*, Patuxent River, MD, Naval Air Systems Command.
- [2] 2001, *Aging Avionics in Military Aircraft*, Washington, DC, National Academy Press.
- [3] May 2000, *Guidelines for Successful Acquisition and Management of Software-Intensive Systems: Weapon Systems, Command and Control Systems, Management Information Systems, Version. 3.0*, Hill AFB, Utah, Software Technology Support Center, <http://www.stsc.hill.af.mil>.
- [4] CPU Tech's Multiple Instruction Set Processor Technology Renews Military and Commercial Electronic Systems, http://www.cputech.com/dual_use.htm (current July 2002).
- [5] Titan Systems Corporation – Rapid Retargeting, <http://www.titansystemscorp.com/media/docs/pdf/rapid.pdf> (current July 2002).
- [6] Madisetti, V. K. et al., Apr-Jun 1999, "Reengineering Legacy Embedded Systems," *IEEE Design and Test of Computers*, vol.16, no. 2, pp. 38-47.
- [7] Luke, J. A., D.G. Haldeman, and W.J. Cannon, Dec. 2001, "A COTS-Based Replacement Strategy for Aging Avionics Computers," *CrossTalk: The Journal of Defense Software Engineering*, vol. 14, no. 12, pp. 14-17.
- [8] June 1992, *AN/AYK-14(V) Navy Standard Airborne Computer Technical Description*, GSRCR-1314, Bloomington, MN, Computing Devices International.
- [9] Cifuentes, C. and M.V. Emmerik, Mar. 2000, "UQBT: Adaptable Binary Translation at Low Cost," *IEEE Computer*, vol. X, no. Y, pp. 60-66.
- [10] Silberman, G. M. and K. Ebcioglu, Jun. 1993, "An Architectural Framework for Supporting Heterogeneous Instruction-Set Architectures," *IEEE Computer*, vol. X, no. Y, pp. 39-56.
- [11] Corman, D., Dec. 2001, "The IULS Approach to Software Wrapper Technology for Upgrading Legacy Systems," *CrossTalk: The Journal of Defense Software Engineering*, vol. 14, no. 12, pp. 9-13.
- [12] Littlejohn, K., et al., Dec. 2001, "Reengineering: An Affordable Approach for Embedded Software Upgrade," *CrossTalk: The Journal of Defense Software Engineering*, vol. 14, no. 12, pp. 4-8.
- [13] <http://www.sae.org/technicalcommittees/aasd.htm> (current July 2002).
- [14] Andrews, W. and C. Ciufo, Nov. 2001, "Contenders Vie for Supremacy in Switched Fabric Market," *RTC Magazine*, vol. 10, no. 11, pp 2-12 of the "Interconnect Strategies" supplement.
- [15] Aaldering, M., Dec. 2001, "Hedging Your Bet: FPGAs Realize Myriad Serial Topologies," *COTS Journal*, vol. 3, no. 12, pp. 58-62.
- [16] "IP" Means Increasing Performance," *COTS Journal*, vol. 3, no. 10, Oct. 2001, pp. 62-63.