33

Advanced Distributed Architectures

Jim Moore
Smiths Industries Aerospace

33.1 Drivers and Trends
Technology Advance • Increasing Functional Complexity • Hardware/Software Cost Ratio Continually Falling • Integration • Modularity • Business Pressures
33.2 Integrated Modular Avionics (IMA)
Concept • Modular Architecture and Supplier Roles • Industry Standard Modules • Commercial Modules • Achieving the Wider Goals for IMA • Control of the Interfaces—Open Systems
33.3 Aircraft and Systems Architecture Issues
"Smart" Peripherals, High-Speed Serial Data Buses, Procurement Boundaries
33.4 Conclusions

In recent years a considerable momentum has built up on the assumption that future avionics will be based on centralized racks or integrated cabinets containing numbers of electronic modules of various types. Integrated Modular Avionics (IMA) in many sectors of the air transport industry has largely become the assumed way forward for the implementation of future avionics. Progress has already been demonstrated with first-generation civil IMA systems such as ELMS and AIMS on the B777. These are quite different implementations, having been optimized for their specific systems domains, and therefore appear to go only part of the way towards meeting the ultimate goals anticipated by the industry.

However, there are a number of reasons to suggest that relatively large groupings of tightly coupled and integrated computing resources may not provide the optimum architecture of future aircraft in the face of rapid technological change and ever-increasing business pressures.

This chapter will raise awareness of these trends and of alternative architectures in which the computing and interface resources are decentralized and operate autonomously within a networking communications environment.

33.1 Drivers and Trends

In addressing these issues it is important to identify and consider the changes that have taken place within the industry and the underlying trends to these changes and drivers

33.1.1 Technology Advance

Perhaps the greatest driver of change is the continuing rapid advance of electronics technology, affecting practically all areas of avionics systems — computers, sensors, displays, data buses, etc. Already there are examples of avionics computers introduced less than a few years ago that are now available at half the size, weight, and power consumption, but with considerably enhanced performance and functionality. For instance, a Flight Management Computer that 10 years ago required a size 8 MCU size box (10 in. wide) can now be implemented as a single-card module (<1 in. wide). Likewise there have been comparable advances in display technology, data buses, control panels, bulk memory media, sensor technology, etc.

These technology advances are being driven primarily by the demands of the large commercial markets, e.g., information technology, communications, PCs, consumer products, automotive, etc. The competitiveness and size of these markets are driving down costs and more and more are setting the performance standards. Aerospace components and communication standards are now being increasingly replaced by commercial ones, e.g.

- Microprocessors, microcontrollers
- Data buses (e.g., Ethernet, CANbus)
- Flat panel AMLCD displays and interfaces (e.g., OpenGL graphics language)

These trends are now firmly established in both the civil and military aviation markets.

The principal drawback is the relatively short lifetime (availability) of commercial components, which can become obsolete within just a few years, in much less time than the production cycle of the average civil aircraft programme. Therefore appropriate strategies are required to deal with part obsolescence.

The rapid and continuing advance of electronic technology is constantly driving down the cost of hardware, as the number and cost of the components used falls and much more functionality is achieved with less and less hardware. Figure 33.1 illustrates the growth in microprocessor performance over time and adherence to Moore's law and the expectation that performance will continue to double every 18 to 24 months, i.e., grow by a factor of over 1000 times in the next 20 years.

This illustrates the need for care in choosing hardware architectures with interface boundaries between the principal elements, which are the least likely to be affected by changes in the hardware technology, i.e., to prevent obsolescence/upgrade repercussions from propagating from the affected element (e.g., module) to other elements. This also applies to the interface to the functional software to protect it from the impact of changes in the processing platform on which it resides.

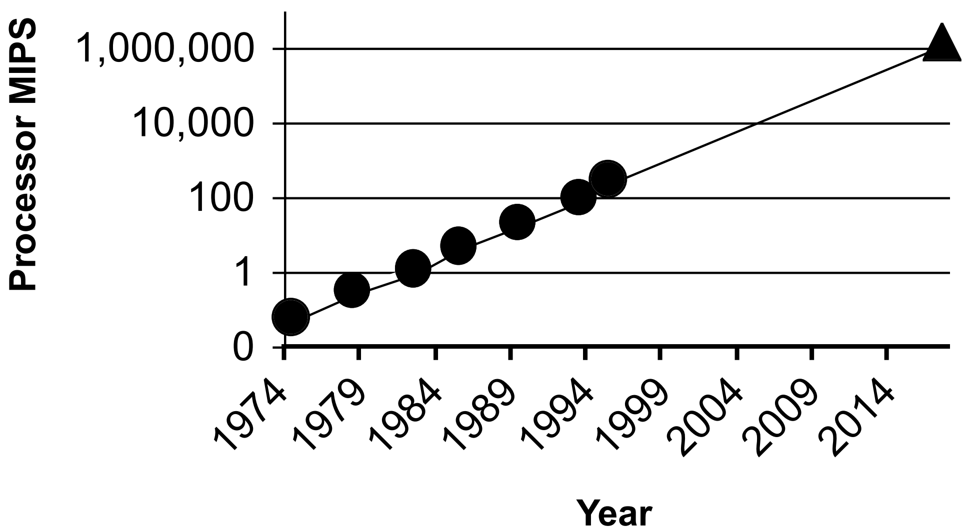

FIGURE 33.1 Increase in microprocessor performance.

33.1.2 Increasing Functional Complexity

Associated with the enhanced capability afforded by the technology, and as driven by the competitive pressures of the civil transport aircraft market, the functionality of avionics systems has continued to escalate. It is evident that avionics have long since become central to the ability of a manufacturer to manufacture and sell competitive aircraft and therefore are key to ongoing business viability. For example;

- Fly-by-wire flight controls
- FMS, full flight regime 4D flight management
- Full glass cockpits, large multifunction displays
- FANS capability to operate in the new air traffic management environment
- Passenger entertainment systems and commercial/business services
- On-board central maintenance computers and electronic documentation

Most of this added functionality is, of course, effected in software, and so the quantity of code embedded in avionics continues to rise. For the suppliers, more software generally leads to increased development and certification cost. High-level languages such as Ada and C are widely used to simplify the programming effort, and where possible, automatic coding is used to further simplify the process and reduce cost.

Much effort is being applied these days to ensure that application software can be reused on different hardware, i.e., be independent of the hardware. The primary driver for this is to avoid the high cost of new software for the application being hosted on different processors, for example, in the event that during the life cycle of the product the processor becomes obsolete or has insufficient capability to support growth in required functionality.

Hardware-independent application software requires an embedded software Operating System (OS) or executive which provides a generic interface to the application code and which translates it for the particular processor and hardware architecture being used. By the use of such an executive and a standardized application interface definition it is possible to achieve compatibility of applications developed by different suppliers capable of being used on a variety of different hardware platforms. Hence the concept of APEX (Application-Executive interface) as documentation ARINC 653.

33.1.3 Hardware/Software Cost Ratio Continually Falling

The combination of the above two trends suggests that, for the foreseeable future at least, the computing hardware cost is going to be less and less significant as a component of the overall cost of systems. The main cost drivers will be for software and integration. Already, today, software typically consumes 70 to 80% of the development budget and represents by far the highest schedule and cost risk to any new programme.

It is also important to recognize that there are recurring elements of software cost. Historically suppliers have included amortized software development costs in the equipment recurring price. But the actual value or market worth of the embedded software has been invisible to the purchaser. However, as the material (hardware) cost falls and as suppliers use or adapt outsourced COTS software elements (e.g., embedded operating systems, data-bus protocol software, networking software) the recurring cost and value of such software will become more significant. In the limit where the product is only software, for example an FMS to be integrated into the customers host system, the price will be driven by the market value per copy of that function, assuming the development costs have already been recovered. After all, it is the software that embodies the supplier's know-how, systems expertise, and would carry copyright.

The lesson this suggests as we consider the next generation of avionics systems architectures is to beware of overcomplicating the software and level of integration just in order to minimize the quantity of hardware employed.

33.1.4 Integration

As technology has advanced, so there has been a continuing trend of avionics integration. In the past most integration has been concerned with processes or functions which were already interdependent, and by so doing savings could be made in interface hardware, aircraft wiring, power supplies, etc. with little risk to the certification of the system. This is usually referred to as "vertical" integration. The integration of inertial sensors with computers to produce Inertial Reference Systems (IRS) is an example of vertical integration.

In other cases separate processors have been colocated into a single box because they can share the same I/O hardware on which they both depend, thereby eliminating one set of I/O hardware and simplifying the aircraft wiring.

More latterly there is a trend to go one step further by the integration of largely unrelated avionics functions onto shared processors. This may be referred to as "horizontal" integration. There are significantly higher risks because, while additional hardware resource may be saved, there are added complexities to provide "equivalent independence" or partitioning within the processing platform. Partitioning is used to ensure that malfunctions within one application (function) cannot affect the others, or that modifications made to one may be certified without the need to revalidate the others. A further complication may be that additional levels of hardware redundancy and associated monitoring and configuration management facilities may be needed to offset the risk of failures within the shared processor, causing simultaneous loss of all the otherwise unrelated application functions.

The trade between savings of hardware (e.g., processor modules) on the one hand, and the escalation in cost to achieve acceptable levels of segregation and integrity on the other, needs to be carefully weighed. This is a trade that would appear to be more difficult to support as the hardware proportion of overall cost continues to fall with time. For example a "cluster" of separate low-cost processor chips, each dedicated to a system function and each including it's own program memory, communicating via shared data memory areas and external bus interfaces, may become a cheaper and better alternative.

33.1.5 Modularity

The concept of modular design for both hardware and software has been with us for a long time. Most avionics Line Replaceable Units (LRUs) employ a modular approach to some extent or other. The object is to build up the design by repetitively using building blocks of a minimum number of different types. A supplier can substantially reduce costs if new designs reuse existing modules. There are benefits of scale in purchased parts and manufacturing. The modules used in greater numbers more rapidly reach maturity and modifications and repairs can be more readily implemented and tested. The number and cost of spare modules needed to support maintenance shops is much reduced.

However modularity also generally adds some complexity, e.g., for connectors and interface circuits that otherwise would not be needed. Generally, there will also be elements added to allow a generic design to meet specific requirements. Thus there is an overhead in modularity which tends to reduce the benefits of integration.

A primary goal of IMA is to establish the application of a standard set of hardware modules, directly line-replaceable, encompassing as much of the total avionics suite as possible. However, the adherence to prescribed module boundaries, e.g., packaging processors separately from I/O modules and from power converter modules, etc. but all within a centralized cabinet, sets up real barriers to the benefits that can be gained from the rapid advances in electronics technology.

This concept stems from previous generations of equipment in which these distinct electronic functions each required relatively large amounts of real estate and could not be economically integrated into the same circuit board or module. Also, it was believed that the industry standards would lead to a market for competitive sources of such modules and therefore yield freedom of choice. However the reality is that this approach delivers only a "closed" architecture, no effective industry standards, and no market for choice of modules.

33.1.6 Business Pressures

The aircraft manufacturers are constantly driving to reduce manufacturing lead times and cost at the same time as the product complexity increases and technology continues to change.

Manufacturers place contracts for large packages of systems to single suppliers or consortia, i.e., the "one-stop shop." This not only reduces the overhead and diversity but also reduces the OEMs integration task and engineering burden. For the smaller manufacturers who may be unable to support large engineering resources this is perhaps the only way to cope with the problem of increasing functional complexity and technological change.

For the larger manufacturers there is the added question of the extent to which they want to also be able to manage the internal "package" integration and modification task independently of the supplier. Coupled with this may be the need to ensure long-term freedom of choice of suppliers and ability to exercise control. This is a key issue bearing upon the degree to which the package of systems or functions has open internal interfaces and is supported with reliable tools.

Because of the high value and high degree of dependence upon such suppliers the manufacturers are increasingly requiring risk-sharing partnerships. This is also evident in the trend towards partner design and manufacture of major airframe parts, or modular aircraft manufacture. Costs, build time, and risk can be reduced when the interfaces between the parts support rapid assembly and the parts are delivered fully dressed and pretested for their systems content.

From a systems viewpoint this is made considerably easier when the distributed elements located in the various major airframe parts are interfaced via serial data buses, perhaps through Remote Data Concentrators (RDCs) or local controllers.

33.2 Integrated Modular Avionics (IMA)

33.2.1 The Concept

The concept of integrated modular avionics has come about in response to the above trends and pressures.

Various implementation architectures are described in ARINC 651, the industry's overall design guidance document for integrated modular avionics. They share the common theme of a number of "cabinets" that are connected together and with other peripheral equipment by means of a number of multiple-access serial data buses (refer to Figure 32.2)

Each cabinet is seen as a high-power computing center which, in essence, replaces a number of today's application-specific avionics computers (line replaceable units, LRU). Each cabinet contains a selected mix of line replaceable modules (LRM) e.g., core processor, input-output (I/O), power supply, and other

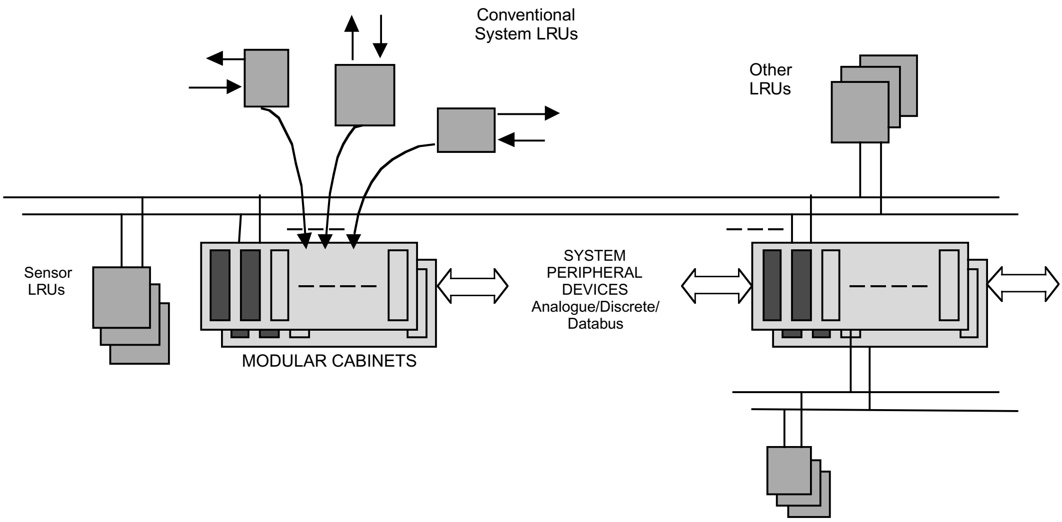

FIGURE 33.2 IMA overall architecture.

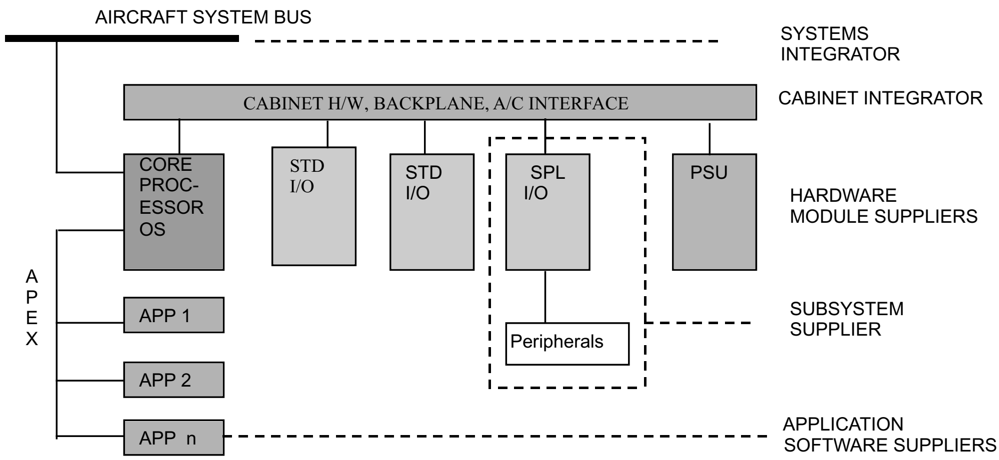

FIGURE 33.3 IMA concept.

special LRMs, all interconnected via an internal backplane bus. The cabinet enclosure is designed to provide a standard environment for the LRMs and effectively replaces conventional avionics equipment racks. Each cabinet is intended to have adequate processing and interface capacity to support the required integration of avionics functionality with spare capacity to meet life cycle growth requirements. The modules and backplane need to be fault tolerant to protect the group of integrated functions against single failures, or to provide high dispatch reliability, particularly for long-range or large aircraft. The system-specific application software is designed to meet the standard application-executive interface specification (APEX) to allow reusability on core processors of different hardware design.

33.2.2 Modular Architecture and Supplier Roles

The goal of standard, reusable, and interchangeable modules is central to the concept of IMA (refer to Figure 33.3). By rigorous definition and control of each module, both hardware and software, and of its interfaces, the aim is to permit a minimum number of different module types to support the broad range of current and future avionics system functions. The concept is expected to allow the systems integrator to procure each type of module from the most competitive source and to allow the user to support the aircraft with a greatly reduced number of spares. The concept, therefore, promises significant reductions in overall life-cycle cost.

It also will have a major impact on the way avionics business is conducted, since it opens up the possibility for LRMs to be supplied separately from the cabinet, software supplied separately from the hardware, and the whole cabinet to be integrated by an outside party.

To date, civil aircraft solutions like the B777 ELMS and AIMS have gone some way towards the envisaged IMA goals. Both approaches were optimized for the application domain and have brought benefits to the manufacturer and airline alike. It is important to recognize that neither AIMS nor ELMS are "open" architectures and all the parts and functions are manufactured, qualified, and integrated by the system supplier.

A key question now is whether and how a truly open common architecture and module suite can be developed and applied across the wider spectrum of avionics system domains. The following issues are considered relevant to this discussion.

33.2.3 Industry Standard Modules

Despite the original intentions, it is now fairly obvious that the industry will not be able to produce standard modules independently of the cabinet supplier/integrators. The specification process is too complex and the means are not available to independently validate and qualify modules for compatibility with the platform. The best that industry committees and agencies can do is to define form, fit, and interface requirements as a means of guidance and recommendations. It is also unlikely that the industry will be able to mandate a particular cabinet architecture. Apart from the complexities of these processes the time involved to achieve mature specifications is generally too long and will be overtaken by technology advance and real program-driven solutions.

33.2.4 Commercial Modules

The parallel is often drawn with the PC industry. The assumption being that because commercial interface standards have become widely established and the large market provides a wide choice of PC-compatible cards for the PC supplier and user, that the same "plug and play" approach will apply to the civil aircraft industry. In fact, for some avionics LRUs PC interfaces and modules are being used. However, the key difference is that the aircraft system integrator will generally have to prove and validate the use of any component used in the platform to a far higher level than used in the nonaviation world. This is likely to include a significant level of testing with the applications for a particular platform. The integrator carries the certification responsibility and will have to approve any alternative module sources.

33.2.5 Achieving the Wider Goals for IMA

The only realistic possibilities for achieving a wider application of common modules and integrated platforms appear to be the following:

- One platform supplier establishes his solution and module suite as the *de facto* standard, or
- A collaborative basis (probably aimed at a new aircraft program) in which the manufacturer and suppliers work together to define the standards and qualify a sufficient choice of interoperable parts.

Clearly the former is unlikely to generate freedom of choice and competitive component supply, and may simply establish a monopolistic market. However, some manufacturers may opt for this approach if they are not concerned about dependence on one supplier.

It is difficult to see how the larger vision of common modules across the wider systems domain along with freedom of choice and competition can be totally realized. Even the collaborative approach is likely to yield only a limited choice of supply. During the formative years the hardware and avionics functions will still be linked, if only to establish the certification principles. Then the technology will progress and the modules are likely to be obsolete for the next aircraft generation, requiring the cycle to be repeated. However, it is possible to see for a new aircraft programme that some module commonality could be achieved across platforms of different application domain, by collaboration between the platform suppliers.

33.2.6 Control of the Interfaces—Open Systems

A related issue is the question of the degree to which the manufacturer intends to be able to exercise control over the interfaces within the integrated system, e.g., within the platform between the hardware modules and between the application software and the operating system. For example he may wish to be in control of the integration of third-party software and variants arising from customer changes during the aircraft life cycle.

This is understood to be a clear objective for at least the manufacturers of large civil transport aircraft. In which case the same principles and features are needed as for integration management at the intersystems/aircraft level. It is this requirement, probably above all others, that will determine whether open architectures really do become established, and also the extent to which the platform internal architecture integration is maximized.

The key interfaces requiring control are the intermodule databuses (backplane) and the application software interface (APEX). It is vital for the aircraft buses to be truly open the communication over these interfaces (at the interoperability level) must be independent of the technology used in the system. Also, the interfaces must place minimum possible constraints on the functioning of the modules and equipment operation.

33.2.6.1 Software Interface

For the application software interface the manufacturer will require the same programming and integration tools as used by the platform supplier integrator. These tools would have been qualified by the platform supplier and would need to be maintained through the life cycle of the aircraft. Appropriate contractual arrangements will be necessary to define responsibility and liability associated with modifications or added functionality implemented by the manufacturer independently of the platform supplier.

33.2.6.2 Hardware Interface

For the module interfaces, it will be much more difficult for the manufacturer to add modules or introduce an upgraded performance module (for example) and then revalidate the platform if the platform architecture is tightly coupled, i.e., one in which the modules are controlled in lock-step fashion by the applications as an extension of the processor internal bus. In effect this is a wholly integrated platform, with the modules having only physical independence from each other. It may be open in the sense that the backplane bus meets a standard such as ARINC 659 or VME, but it does not readily permit the user to change or add hardware functionality independently of the integrator.

33.3 Aircraft and Systems Architecture Issues

33.3.1 "Smart" Peripherals

The technology advances that have enabled rack-mounted LRUs to be reduced in size or be integrated together over the years have also, in combination with serial data bus interfaces and other advances, enabled electronics to be used increasingly in various locations around the aircraft. Some examples of "smart peripherals" or "distributed systems" are

Full Authority Engine Control (FADEC) — Engine mounted
Air Data Modules (ADM) — Fuselage skin
Cabin entertainment systems — Cabin distributed
Smart autothrottle actuators — under floor
Smart display heads — Flight deck
Smart systems panels — Flight deck/cabin

With today's advanced microcircuits it is already possible to provide considerable functionality within many system peripheral devices without significantly affecting the space they occupy or significantly reducing their inherent reliability. As costs fall we are likely to see data-bus-interfaced smart peripherals being used increasingly, e.g., primary and secondary displays, flight control actuators, landing gear sensors/actuators, tank-mounted fuel quantity processing, remote data concentrators, radio sensors/transceivers, etc.

The main drivers are

- Reduction in installation complexity — wires, connectors.
- Robustness of peripheral interfaces to noise, improved HIRF/EME immunity.
- Reduced uncertainties in identifying fault location, i.e., whether fault is in the peripheral, the computer, or the wiring.
- Simplification in data interfaces, by incorporating special peripheral processing at the "point of action" and reducing communication data flow to higher-order parameters.

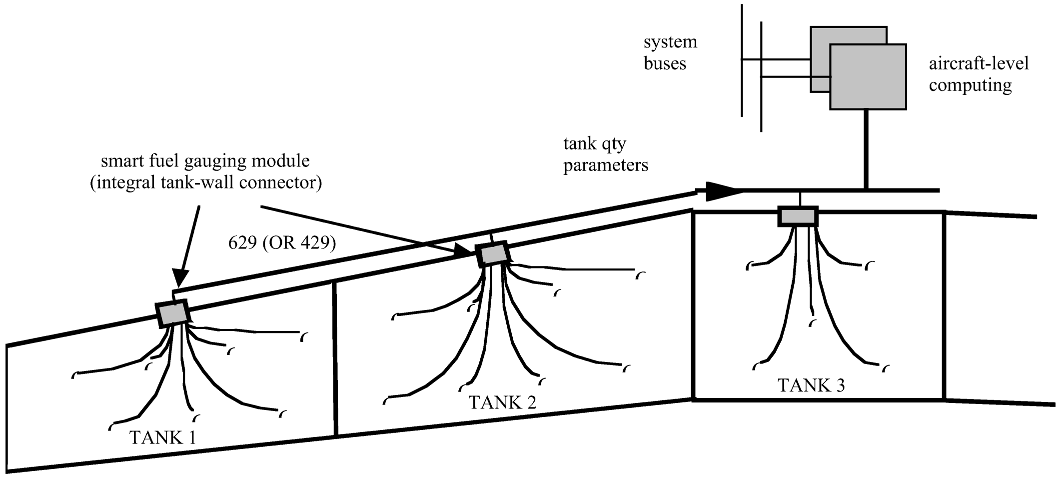

FIGURE 33.4 Distributed fuel gauging system.

- Architectural flexibility; the I/O interface is available on aircraft-level data buses for direct use elsewhere.
- Enabling the interface to be independent of the technology of the peripheral.
- Simplified procurement boundary promoting freedom of choice and control by the integrator.

Electronics is used at the peripherals and serial data buses provide the interface, so the systems become digital from end to end.

Figure 33.4 illustrates an example of a distributed Fuel Gauging System using the principal of smart peripherals. In a conventional high accuracy system the individual tank probes would be wired back to a gauging computer in the avionics bay. In a large aircraft there could be over a 100 probes and therefore at least 200 wires running over a considerable distance through the airframe, carrying sensitive analogue signals of very low strength. By locating highly reliable electronics close to the tanks this wiring is greatly simplified. All the probe signals for each tank are now combined onto a serial data bus. The signals are preprocessed into a form which is now independent of the probe technology and, being digital, are much less sensitive to wiring abnormalities and electrical interference. The processing task for the central computer is also simplified so that the gauging computation can now more easily be accomplished on a standard computing resource such as an Avionics Computer Resource (ACR).

The greater use of smart peripherals and remote data concentrators will also significantly diminish the need for analogue and discrete I/O modules within IMA cabinets and computer LRUs (see Figure 33.5).

33.3.2 High Speed Serial Data Buses

Serial digital data buses have been used on both civil and military aircraft of all types for many years. ARINC 429, which is a single-transmitter multiple-receiver topology bus operating at either 12 Kbps or 100 Kbps, is used on all modern civil transport aircraft. The disadvantage of this topology is that a single bus can carry data only from one source, and so each item of equipment needs as many inputs as the number of buses (sources) that it expects to receive data from. Accordingly, there are numerous bus spur connections and wires in the aircraft associated with most of the avionics boxes.

The Boeing 777 was the first civil transport aircraft to use a multiple access data bus in which all connected units can both transmit and receive via a single port connection (there are, of course, multiple redundant buses). This bus operates at 2 Mbps and uses active couplers to connect the unit port via a stub cable to the linear bi-directional bus itself. This is ARINC 629. This topology simplifies equipment

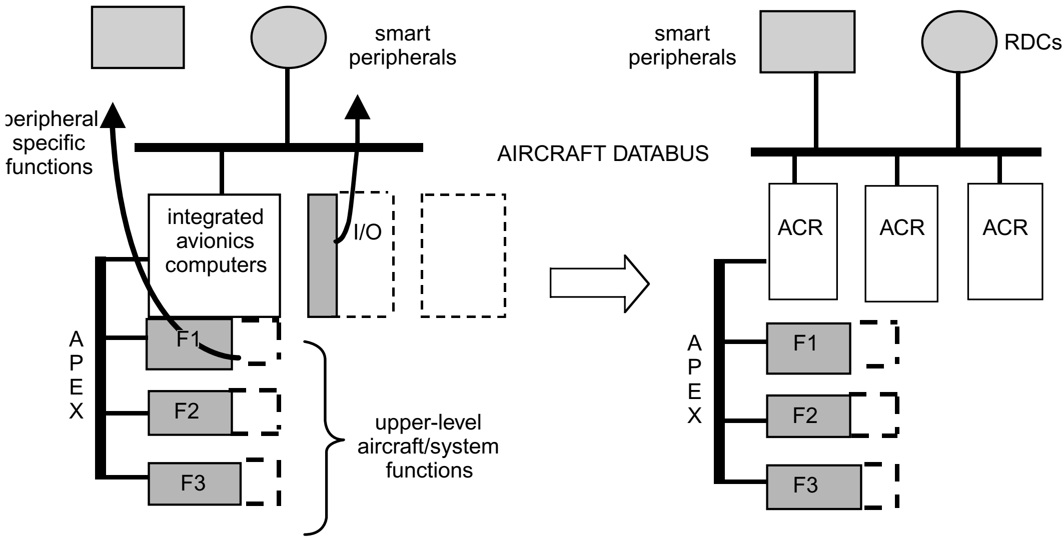

FIGURE 33.5 RDCs reduce need for analogue and discrete I/O modules.

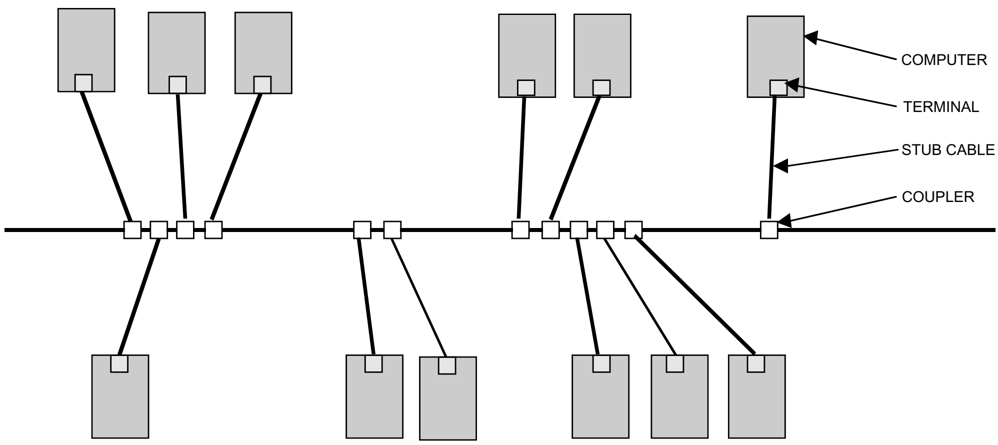

FIGURE 33.6 ARINC 629 topology.

and the aircraft wiring installation because one port and one bus stub connection only are needed for each unit to communicate to all the other units on the same bus. The bus is very deterministic in its operation and each terminal has high integrity control over the associated unit's access to the bus, the correct use of allocated bandwidth, and the selection of transmitted and received data. Its principal limitation, however, is the 2 Mbps rate. Also, the components are relatively expensive and there can be electrical performance factors which limit the physical spacing of coupler connections along the bus. Figure 33.6 illustrates the ARINC 629 topology.

More recently, fast networking data buses based on commercial standards are starting to appear in civil avionics systems. The industry has started standardization work on Ethernet and the first example of FDX (full duplex Ethernet) applied on a civil transport aircraft is for the interfaces between the various display units and computers in the cockpit primary display system on the B767-400 aircraft.

Airbus is also planning to use a variant of FDX (AFDX) as the main avionics and systems buses on the A3XX. These buses typically have a STAR topology in which the equipment ports are connected to

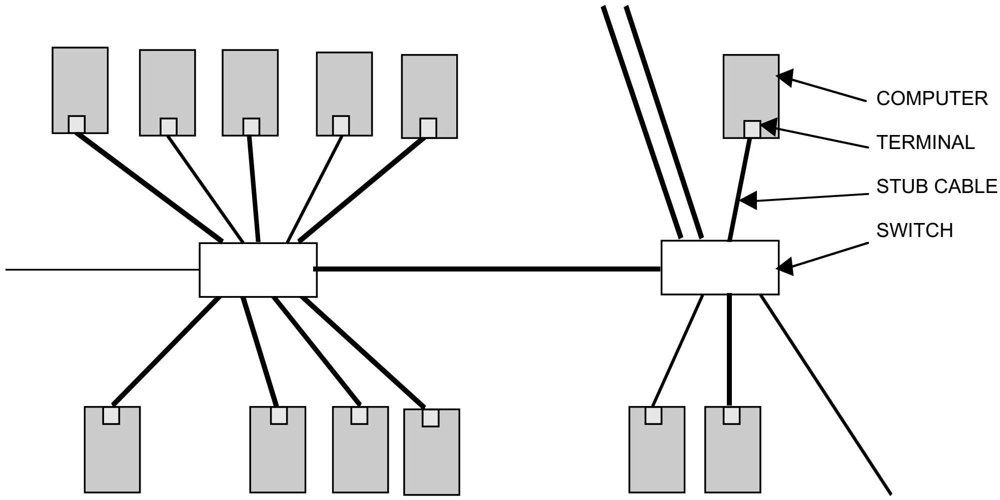

FIGURE 33.7 FDX (AFDX) topology.

each other via stub cables and the star point (Switch Unit). Again each item of equipment can communicate with all the others on the same bus through one port. In this case the data rate may be typically 100 Mbps, much faster than ARINC 629, although some of this bandwidth is used by higher formatting and protocol overheads. The concept is potentially much cheaper, since it can use the commercial standard directly (at least for the physical layer components), and it is likely to become the next civil aircraft industry bus standard. For most aircraft applications, however, modifications to the services and functions of standard Ethernet are needed to ensure a statically configured network, additional integrity, and controls on message routing and channel bandwidth usage. These changes impact upon the design of both the switch and the user terminal. The topology is illustrated in Figure 33.7.

An important advantage of buses like AFDX is that they can form a continuous network covering both the aircraft level (e.g., communication between LRUs, IMA centers, and other smart peripheral equipment) and the systems domain level (e.g., communication between modules associated with a particular group of systems). In an IMA context the switch provides the backplane connectivity for a domain and provides a direct interface to the aircraft-level part of the network, without the need for conversion gateways. In addition, only one set of bus management tools is needed to integrate the communication at all levels and for all domains across the aircraft.

33.3.3 Procurement Boundaries

An important architectural driver for the aircraft manufacturer is the procurement boundary, i.e., the interfaces through which a system or equipment to be procured operates in relation to its neighbours. Simplifying these interfaces is one of the goals inherent in determining where the boundaries should be. This is a necessary process to minimze the specification and integration task, reduce aircraft wiring and data flow, and achieve correct interoperability with minimum risk. Figure 33.8 illustrates these principles with a hypothetical example.

System functions A, B, and C (FA, FB, and FC) are highly related to each other and have complex interfaces between them, but have relatively simple or standard interfaces with sensors S1, S2, and S3. In addition FB and FC communicate with FD via a small number of readily specified signals. It is assumed that suppliers exist who are competent in all these system functions. The grouping together of functions A, B, and C makes sense because it significantly simplifies the manufacturer's engineering and procurement workload compared, say, to specifying and procuring FA combined with its associated sensors S1 and S2, separately from FB and FC. The (procurement) boundaries around the FA, FB, and FC group

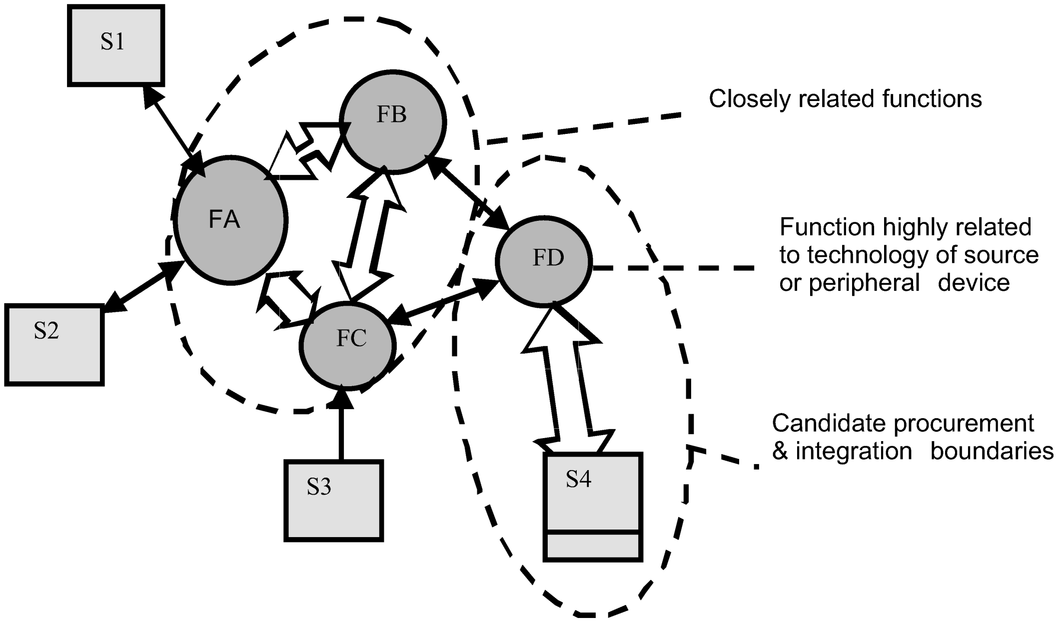

FIGURE 33.8 Integration and procurement boundaries.

are all relatively simple, easily controlled, and not dependent on the technology of the peripheral sensors. Therefore these functions are also likely to be good candidates for integration within a common computer, with least risk of requirements capture errors or overall integration problems within the aircraft.

On the other hand, FD is very closely related to sensor 4. The signal interfaces are complex or the function is highly dependent on the technology or specific design of the sensor. The simplest procurement boundary, therefore, encloses, both, and once again there is the potential for further simplification through integration of the function with the sensor.

Analyses of functional relationships and interdependencies across all the aircraft systems are required to determine the best options for functional grouping and integration. A top-down hierarchical approach is needed to establish the optimum groupings at the global systems level and the systems domains level. This approach assists the manufacturer to ensure that the functional groupings simultaneously meet technical and commercial requirements.

33.4 Conclusions

The foregoing acknowledges that benefits can be achieved by the careful application of cabinet-based integrated modular avionics for any given aircraft programme. The degree of integration, as distinct from modularity within an IMA platform or cabinet, and the degree of openness of the internal platform interfaces for any particular aircraft programme, will depend largely upon the purchaser's (e.g., aircraft manufacturer's) desire for control over these interfaces and the integration process.

If the aircraft manufacturer is to readily fulfil the role of platform integrator, a less tightly integrated backplane and module communication concept is needed. Here the modules would have autonomy of operation and communicate with identified data sets or messages. Each would be specified for its functionality and communication requirements. Each could be validated to a greater extent independently of the application environment. This concept is not dissimilar to the way LRUs interoperate. Indeed, it is to some considerable advantage to the manufacturer and the end user if the same data bus standard can be used for the backplane as is used for aircraft-level communications. Then exactly the same integration tools and maintenance tools can be used at both the aircraft and system domain level.

In the longer term, as the aircraft systems become more and more digital end-to-end through the use of smart peripherals and remote data concentrators, the need for analogue and discrete IO interfaces within centralized computing platforms will diminish. In addition, the processing within these platforms need not be burdened with functions associated with peripheral devices and can be simplified to that needed to perform upper-level system functions only.

It is also evident that the conventional segregation of processors, IO, and power converter hardware into distinct and separate modules is too restrictive and uneconomic compared with single modules containing all of this capability using the latest advances in technology. Hence, low-cost compact LRMs or LRUs providing a computing platform capability with substantial capacity and without internal architectural constraints are now achievable and are likely to become increasingly more competitive as time passes. Though these may differ internally from different suppliers, they can be designed to meet performance and interface standards (e.g., as for the ACR) including the APEX interface for abstracted application software and hardware interfaces for standard and special IO, both digital and analogue.

It is possible, therefore, to envisage the advantages hoped for from IMA to be obtained by the use of multiple ACRs, each hosting a number of system applications, RDCs, and smart peripherals, networked together using high-speed data buses such as AFDX.