Packaging, Interconnect and the Systems Integrator

by S. G. TYLER, B.Sc.
and T. P. YOUNG, B.Sc., Ph.D.
Marconi Technology Centres, Baddow

Steve Tyler graduated in Mechanical Engineering from North East London Polytechnic in 1982 and spent eighteen months as a tutor and research assistant. He joined the Marconi Research Centre in 1985 and, until 1992, he managed projects and developed processes in electronic interconnection and packaging, including laser processing and ultra-fine lithography. Since then, he has managed a team dedicated to R&D 3-5 years ahead of current production. Having submitted over 20 patents and published or presented frequently, Steve is a member of IMAPS and the PCIF and runs the Interconnect Systems Team within the Photonics Division at the Marconi Technology Centres.
(E-mail: steve.tyler@gecm.com)

Terry Young graduated in Electronic Engineering and Physics from the University of Birmingham in 1981, gaining a Ph.D. in laser heterodyne spectroscopy in 1986. He joined the Marconi Research Centre at Great Baddow in 1985 to simulate optoelectronic devices and was awarded the IEE Premium in 1989. Developing from devices to fibre systems, he has published or co-authored almost forty papers in total. In 1994/5 he moved into business development before a two-year stint running the Optoelectronic and Interconnection Division which embraced digital and RF opto-systems, displays and packaging. He is now a Business Development Manager for the Marconi Technology Centres.
(E-mail: terry.young@gecm.com)

Ever since Intel's Gordon Moore first predicted that transistor integration density was likely to double every 18 months, 'Moore's Law' has proved remarkably prescient, inexorably delivering bigger and faster processors and memory. Intel's 4004, produced in 1971 had 2300 transistors. The 1995 Pentium Pro contained almost 2400 times as many at 5.5 million, with a performance about 4500 times greater. Current predictions are that Moore's Law is good for another decade or two. An accessible retrospective on the electronics scene is given in reference (1).

Over a similar time frame, fibre optics has moved from undersea communications into local area networks and fibre systems and now cable TV. Optical back-planes are now emerging, creating a market for high-performance communications over a distance of a few centimetres.

Against this, developments in electronic packaging appear more prosaic. Whilst the metrics of silicon improvement appear as factors of thousands, the PCB industry would struggle to find a metric that has improved more than a hundred times over the same period and most improvement factors are well below that. Interconnection, it would seem, is destined to be the poor relation of silicon.

However, the relentless pace of change creates problems of its own. Shorter silicon generations have created a situation where chipsets used in the design phase of major programmes are no longer available at the start of the production run – to say nothing of the end of the product's life. A shorter silicon shelf-life, together with other pressures, such as the burden of qualification against the more extreme environments, makes the use of commercial chips for military, aerospace and safety-critical applications highly desirable. Managing the development and deployment of product over a period that represents many generations of its silicon 'innards' is now a major problem born out of silicon's own success.

For the system integrator, this problem is exacerbated: increasingly, the customer expects a long-life system to track the market (benchmarked generally by the PC sector), to benefit from frequent upgrades, face-lifts and even the introduction of modes of operation not conceived of when the purchase was first made. So the obsolescence issue is not merely a case of continuing to maintain performance over a moving foundation, but continuing to deliver state-of-the art performance over a decade or more.

Furthermore, silicon's ability to deliver more functionality within smaller volumes and for less overall power consumption has seen mainframe computing power creep under the desk, onto the desk, into the briefcase and into the pocket. The customer expects that system facilities which are currently fixed infrastructure will become portable and eventually, perhaps, wearable. As the limiting features in systems become the human interface, sensors, connectors or other interconnect, such expectations become harder to realize.

Finally (for the present), this shrinkage means that what was a system has become a subsystem – perhaps even a component. As this happens, the boundary lines change between system elements and amongst the specialist supplier domains from which they come. A shrinking system is one in which the partitions must continuously be redefined. What was a dedicated box, becomes a card in another box, a module on a card, or even a module integrated onto the flexible interconnect between two other parts of the system.

This paper sets out to show that some of these problems, created by the success of silicon, can be solved by advances in packaging technology. As

part of an overall strategy, packaging has a central role to play in the war against obsolescence. In the greater battle to manage through-life costs, packaging has a clear contribution to make in providing options to combine the best commercial and bespoke engineering products. Finally, through careful attention to packaging as part of a concurrently-engineered (or holistic) approach, miniaturized systems can continue to emerge and develop.

Having such potential in the present climate, we believe that packaging should occupy a more central role in systems engineering. It should be considered at the outset, rather than at the sunset of the design process. We believe that the time and resources dedicated to addressing packaging issues will more than pay for themselves over the product life cycle.

The advances in packaging and interconnect are well documented elsewhere[2] and this paper is not intended to describe packaging technology itself. For the unfamiliar, fig. 1 shows some of the options. This technology – involving laser drilling techniques, accurate alignment and high-resolution lithography – has produced a generation of boards with track width, separation and vias all around a few thousandths of an inch (50-100μm). At the same time, flip-chip, chip-on-board, and other assembly technologies[3] are increasingly enabling manufacturers to throw away the chip packaging and work straight onto the board. Microwave circuits and even optical fibre can be laminated into boards using processes that are already very close to mainstream manufacturing methods and are eminently amenable to productionization. Together with three-dimensional interconnect and a host of other developments, these advances are having a profound effect in the market for portable products and in the high street. Here we examine the implications for the large system integrator.

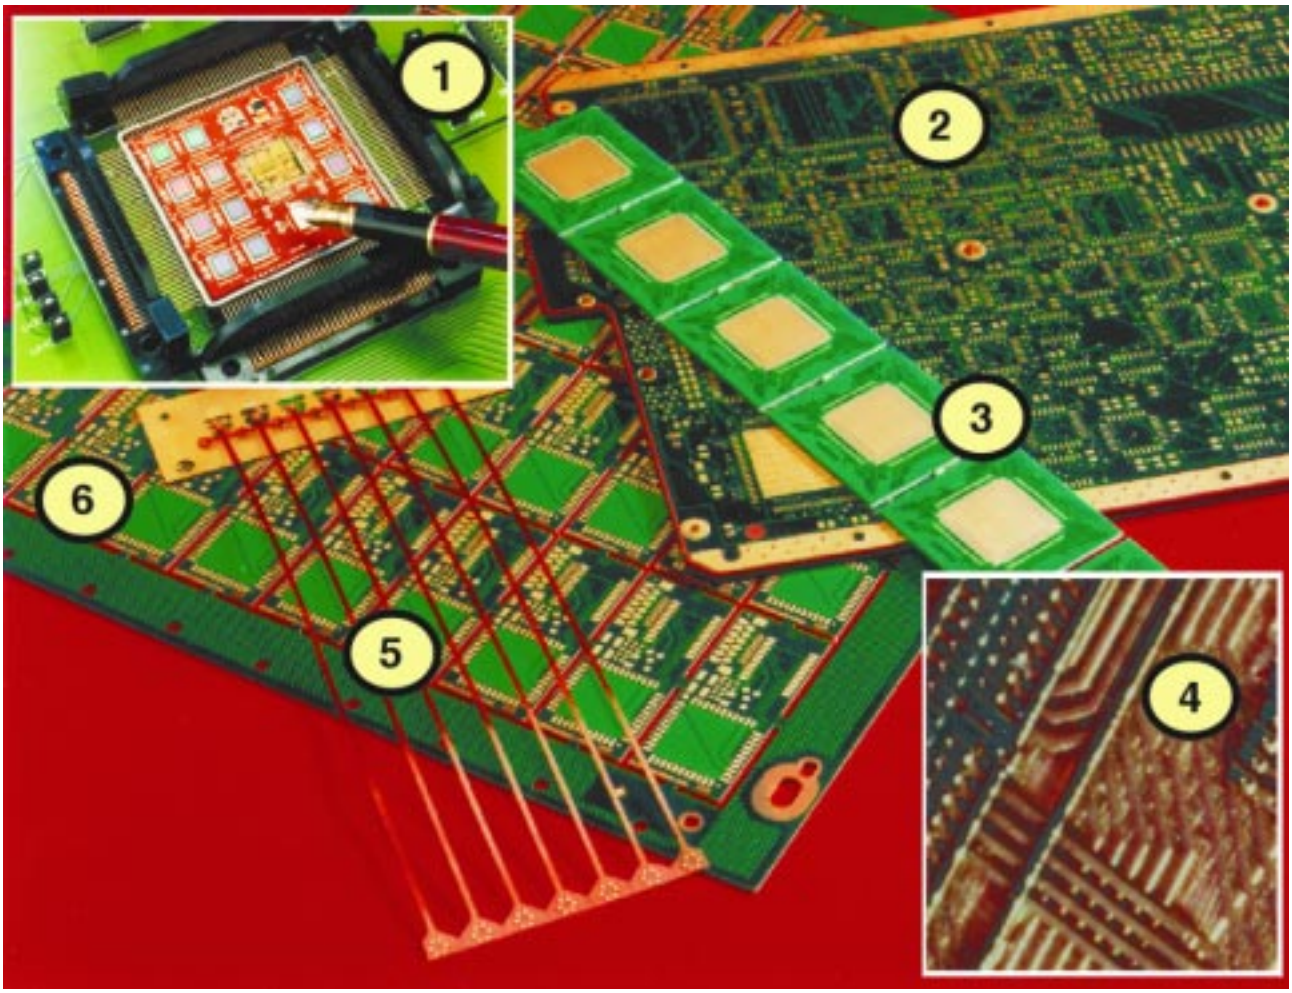

1 ***Examples of interconnection options: (1) a multi-chip module, MCM; (2) high-density laminate on heat-conducting core; (3) ball grid array single-chip carriers; (4) detail of high-density interconnect with dielectric etched away; (5) custom chip-on-flexi; and (6) microvia panel***

Obsolescence

Change Management in the High Street

The PC market is remarkable in so many ways that one must take care in looking to it for trends, lest a given trend should turn out to be simply an amalgam of special cases. However, as noted earlier, the PC market conditions our expectations in terms of the pace and management of change. The authors find it remarkable that home and office software developed in the late 1970s will still run on brand-new computers. However, to benefit from two decades of backward compatibility, one would have needed to have bought the right software at the time – no mean feat in view of the variety of computers on offer in that emerging market. It must also be allowed that, for the vast majority, a version of DOS is effectively a legacy management strategy, rather than an operating system of choice. Furthermore, even more recent systems have not maintained anything like strict backward compatibility.

A more subtle, yet useful, and certainly successful form of continuing compatibility lies in recognizing that people seldom need to run exactly the same software, provided that they can continue to access, edit and re-save the files generated using their old software. Thus, an effective change management strategy lies in maintaining (and even developing) a given 'look-and-feel' and providing software filters to legacy files. The idea of several layers to a legacy management strategy is most important.

Of course, it rarely matters if the legacy management is imperfect. Nobody's life is at stake, and the blame for failure generally attaches to the user, rather than the supplier. The situation is very different for the large system integrator, although the concept of multiple levels at which change and legacy may be managed is certain to be central to a successful obsolescence strategy.

Obsolescence Management and Packaging at the Physical Layer

At its simplest, a higher density interconnection medium offers the ability to add function within the original size and weight constraints. Fig. 2 shows a scrambler circuit realized on fineline PCB and designed to fit between the battery and the transmitter of a two-way radio.

This provided enhanced functionality within the original design, eliminating the need to redesign the original radio unit into which it was integrated. Fig. 3 shows a filter bank, designed to fit onto board that had been in service for some time. Again, the performance of the system was enhanced within the constraints of an existing module. Both examples use high-density PCB technology[2,3], both involve a degree of ingenuity to pack the components onto the board. Both deliver new functionality to an existing product and do so without the need to redesign the original system or, indeed, to change the physical layout.

Whilst these examples are intrinsically remedial, involving an R&D exercise after the basic system was built, they indicate a basic way in which advanced interconnect may be used to counter obsolescence.

This concept can be developed at the chip level using single-chip carriers (SCCs), as shown in fig. 4. Here a small, high-density, laminated PCB accepts a single chip (such as a processor) on top and provides a ball grid array (BGA) for interconnection onto the main board. In fact, The Marconi Research Centre, Great Baddow (now part of he Marconi Technology Centres), was the first organization in Europe to develop 1mm pitch BGAs. This was as part of a collaborative programme in which it was demonstrated that laminate BGAs are inexpensive enough to displace ceramic as the preferred technology for this type of application.

The SCC or BGA immediately provides a simple barrier against the incoming tide of obsolescence:

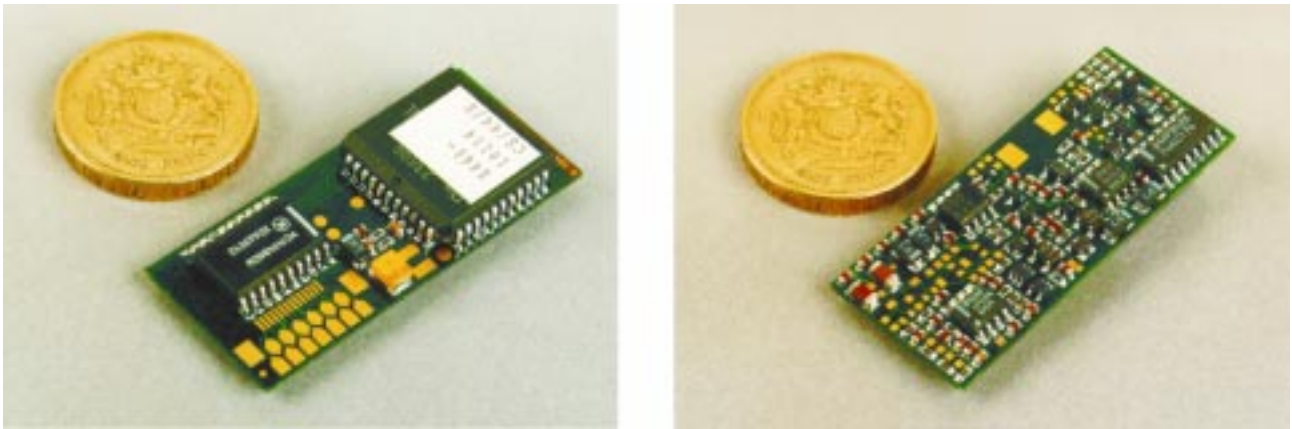

2 *Scrambler circuit realized on microvia PCB*

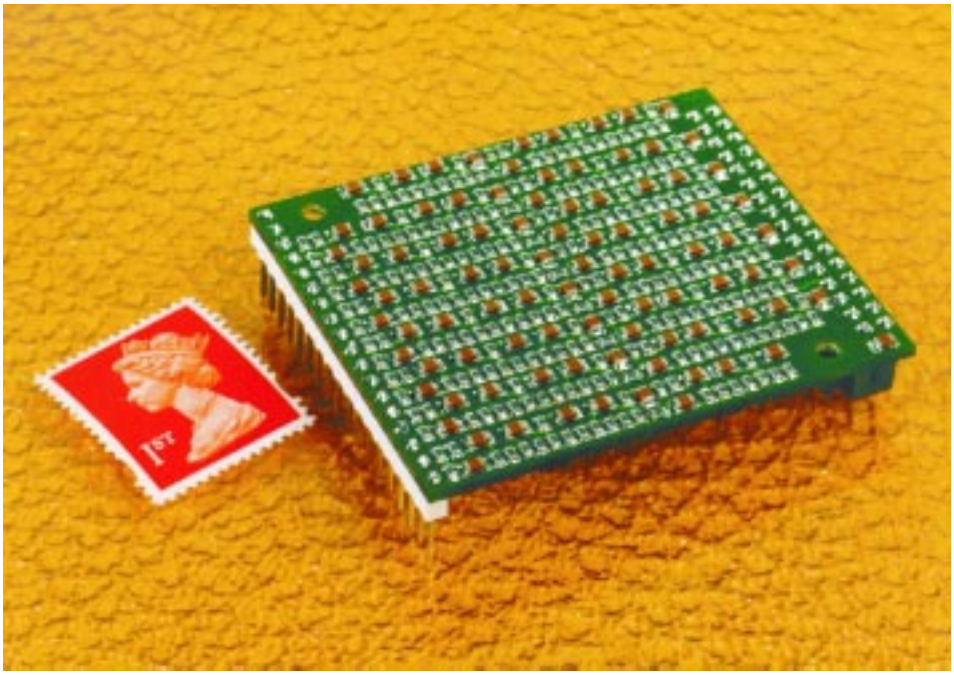

3 *Filter bank upgrade*

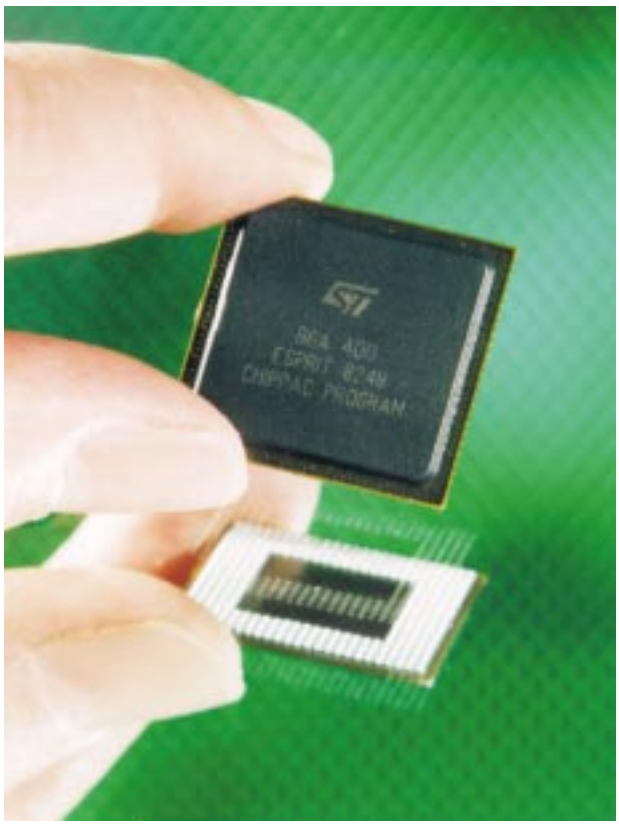

4 *SCC with pentium processor assembled onto it*

if the higher power successor to the current chip has different pin-outs, the interconnection patterned into the BGA can be modified to restore pin-to-pin compatibility to the board beneath. In fact, given the trend towards increasing processing power and shrinking chips, one could predict, with confidence, that it would always be possible to assemble a circuit above the board to deliver a performance envelope, to the board beneath, that fully embraced the original specification.

This concept can be extended to the multi-chip module, or MCM, in which a few chips, or chips plus passive components, sensors, clocks, etc., can be integrated onto a small substrate that interfaces to the main board through, for instance, a ball grid array. The MCM is a very powerful form of packaging in this context, providing a substrate onto which disparate package types may be integrated within tight size constraints. The small size makes it ideal for high-speed circuits and it offers low tooling costs and short implementation cycles.

Advanced Packaging Concepts for Obsolescence Management

Fig. 5 shows a packaging concept known as solder-free interconnect (SFI). In this concept, a series of multi-chip modules (MCMs) sit on a motherboard. With miniaturization, each module would contain a circuit that might conventionally fill a board. These modules are not soldered into place, however, but are pressure-connected using a variety of micro-connection techniques between the module and the board. Here the interconnect facilitates in-service upgrades, provided typically by a technician with a screwdriver.

Through SFI and other related approaches, one may provide islands of function that can be upgraded. Given the trends in technology, these islands will typically contain a great deal more

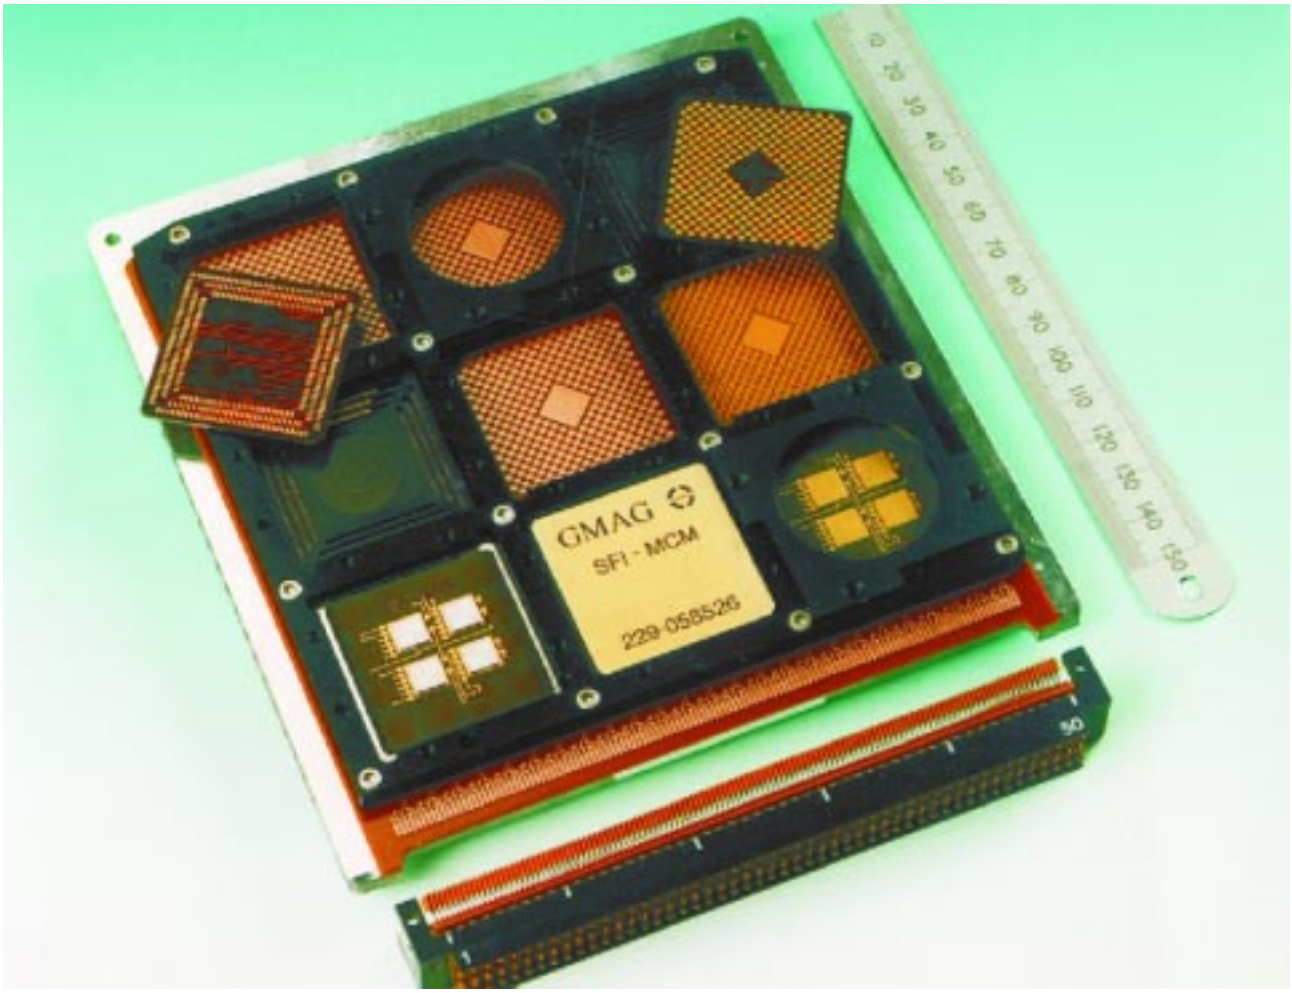

5 ***Solder-free interconnect (SFI) demonstration module***

processing power towards the end of the product life than at the start. It also provides the means to manage these upgrades whilst maintaining full functionality throughout.

Clearly, SFI on its own will not win the war against obsolescence, but in combination with field-programmable gate arrays (FPGAs) and the system software, one has options for managing an evolutionary route forward over many years. It is important to note, however, that the decision on partitioning and packaging must be made as part of the first design on a concurrently-engineered basis, or the whole strategy disappears. Algorithms for partitioning such systems are being developed in collaboration with UK universities.

One need not go for a fully modular design such as that shown on the SFI board: it may be enough to have one or two BGA sites on a board for current or future use in upgrading the board. A pressure connector would have to be developed for isolated SFI modules. Fig. 6 shows what such a board might look like.

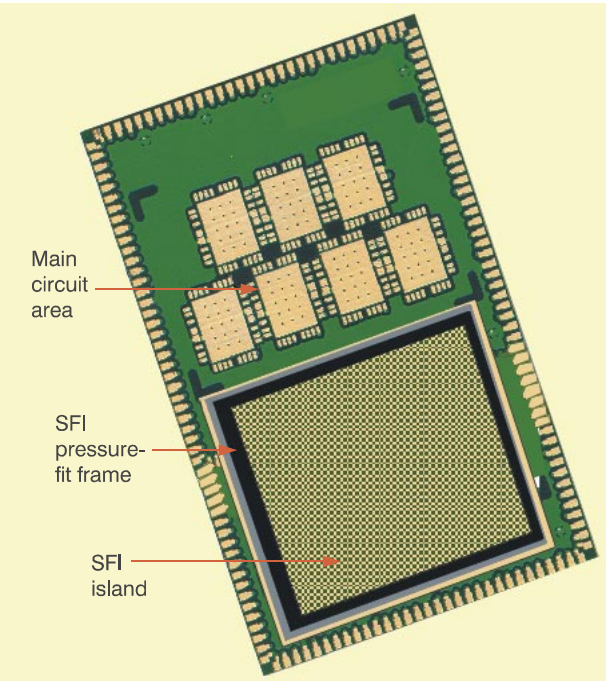

6 ***Advanced board with area reserved for SFI***

Cost and COTS

Two attractions of procuring commercial, off-the-shelf (COTS) technology are the cost benefits and the level of choice afforded by commercial markets. In practice, there are several barriers to a straightforward appropriation of COTS technology for military, aerospace and safety-critical applications. These include:

- obsolescence (see earlier section);
- safety, security and confidentiality;
- reliability; and
- ruggedization for harsher environments.

Packaging may have an effect on safety and security issues, but it is rather oblique. It might, for instance, include PCMCIA modules for encryption.

The introduction of COTS systems into harsh environments has been achieved simply by repackaging the outer casing to fit it for the application. However, in other situations, the cost of this hardening undermines the cost savings effected by going COTS. The standard response to this type of problem is to redesign from scratch. Familiarity with the packaging options and control of the technology offers another way forward: repackaged COTS systems. As such, this option is distinct from mechanical repackaging of the outer casing, recasting the system as an ASIC, or redesigning the system from scratch, but using COTS chips.

Repackaged COTS Systems

In the new world of packaging, it may pay specialist suppliers to acquire complete system designs along with proven chipsets and add value by repackaging them for their applications domains. A recurring theme across the hardware and software industry is a redefinition of traditional boundaries: software libraries are routinely leased and used where suppliers would once have produced in-house solutions. ASIC designers use libraries of functions, including complete processors, in order to reach their objectives swiftly and to focus on the aspects of the design and application at which they have chosen to excel.

This concept of repackaged COTS has yet to define the conditions under which it is economically viable – clearly the integration exercise must add more value than the original designer receives in IPR payments, royalties or licence fees. However, it provides a route towards volume and weight savings without prejudicing the integrity of proven functionality.

Power savings also accrue, as one can see from the way in which the mobile equipment market uses new packaging technology so extensively. The smaller circuits are, the less interconnect there is to drive and this translates directly into power savings. However, a second effect accelerates the power saving with miniaturization: as more chips reside on a card, the interface chips (line drivers and receivers) simply disappear along with their attendant power drain. This in turn creates more space for more chips. Given the fact that bus interfaces typically occupy 10% of the real estate of a line-replaceable module and may take as much as 25%, this virtuous circle can have a profound effect on the overall power and weight budgets. Some data on the environmental performance of this type of technology is given in reference (4).

Fig. 7 shows the trends in packing efficiency (silicon area/board area) for different types of interconnect. Traditional DIL packages on double-sided PCB are around 2%, whilst MCMs – which can pack silicon in several layers – can exceed 100%. The timeline along the bottom indicates the relevance of these advances to the present discussion.

Clearly, there is a new discipline to emerge in determining how best to take a working product and repackage it so as to maintain the integrity of the original design whilst shedding unwanted power dissipation, volume, mass and susceptibility to mechanical stress and vibration, coupled with ability to manage the thermal problem of high-density electronics. The EMC dimension of this exercise is not a trivial one either, although reductions in size and power consumption generally favour the EMC designer. The system integrator of the future will require a thorough and far-reaching understanding of the packaging technology available.

This links very directly to the earlier discussion on partitioning. For instance, as cards shrink to multi-chip modules (MCMs), interconnected using, say, SFI, the card line-drivers disappear and the bus interface has effectively become the card: the plug-in interface is now a planar interface. In the simple case where this advantage is traded for more performance (that is, more cards of MCMs going into the same box) the partitioning issue is trivial. However, for systems where the card of MCMs is a self-contained subsystem, the ability to package it in planar fashion as part of an interconnect fabric might have a profound impact.

Before leaving this section it is worth noting, then, that the ability to shrink the interconnect and to increase the density of functions redefines the old system partitions. A 'card's worth' of circuitry may typically occupy only 10-20% of the original

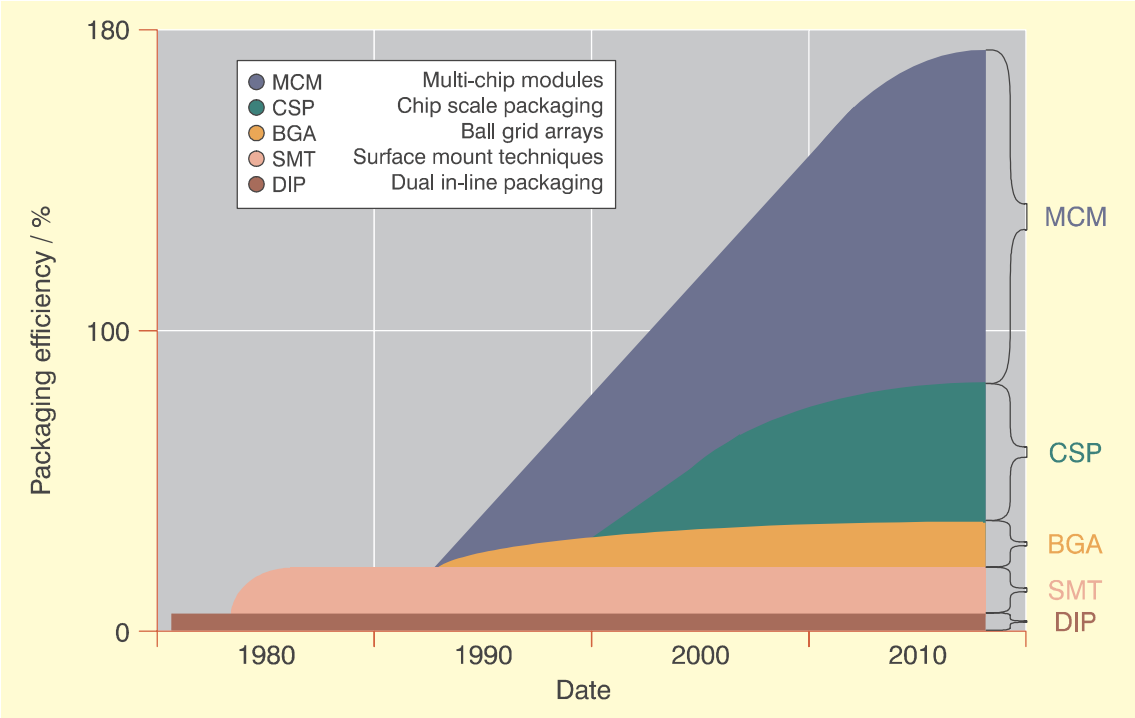

7 ***Board fill-factor as a function of technology (key to right-hand column from bottom: dual-in-line packages (DIP), surface mount techniques(SMT), ball grid arrays (BGA), chip-scale packaging (CSP), and multi-chip modules (MCM))***

real estate and so old concepts of how the system is best divided up need drastic revision. However, this repartitioning of systems is a recurring theme in the packaging revolution now underway – not just in terms of what goes on which card, but in terms of what goes on fixed cards and what goes on flexible interconnect; what goes on the substrate supporting a sensor and what is part of the 'real circuit'. These new options create the vision of interconnect as a seamless entity supporting a diversity of functions through different fabrics (laminates, flexis, embedded MCMs, and so forth).

System Partitioning and Integration

Behind this lies a larger debate in terms of the suppliers of functionality. A supplier with an isolated system to sell, complete with its own power supplies, interfaces, built-in test, etc., now discovers that much of that system has been subsumed, through the integration exercise just alluded to, into the bigger system fabric.

The residue of deliverable hardware may no longer represent viable business. Sensor suppliers who have traditionally relied on providing the back-end processing to boost turnover and margins, may discover that the sensor alone is all the system integrator needs, because the integration exercise provides access to ample processing power elsewhere. This trend is not purely a product of advances in packaging and interconnect, but is fuelled by the emergence of open networking, protocols and operating systems. A by-product of the packaging revolution could, in theory, be to move sensor vendors up-system.

The packaging part of the argument, however, may also be made to work in favour of sensor or 'widget' suppliers, because they can now integrate more functionality around their speciality at little extra cost (in terms of currency, consumption or weight). The strategy for survival, however, depends critically on understanding these trade-offs and then pitching in a product with the right amount of processing, the right type of interfaces, the right size, weight and cost, to make it easiest for the customer to integrate into a system-of-systems.

The evolution of TV cameras is an excellent illustration of the way in which system partitioning has changed over the years. People-sized systems with identifiable boxes of electronics have eventually emerged as elegant accessories that lie easily in an open hand. The key to success has not simply involved making everything smaller, but in using the interconnection fabric to integrate fully all the disparate sensors, modules, and peripheral drives, as well as to make best use of the available space.

Assembly of some 'high street' portable products is almost an exercise in origami, as the flexible interconnect is folded up and springs into place on installation. The concept of the 'electronic driver card' or the 'power supply module' is disappearing. The system is engineered as an integrated whole – with all that that means for design, fabrication and assembly.

At the Marconi Technology Centres, Great Baddow we have focused on the development of the fabric to meet a range of requirements. This technology has been available for some time and progress in the commercial market shows how far-sighted some of it has been. Two further developments have also been made in order to support the type of integration required by companies such as Marconi Electronic Systems: microwave circuitry and fibre-in-board.

Microwaves on PCBs

Because the cost of substrates in microwave systems is generally insignificant when compared with the cost of the components, it does not make sense to develop microwave PCBs purely on cost grounds. However, as noted above, there is a great deal to be gained in having a common medium that will embrace all aspects of the system. Fig. 8 shows a 10GHz signal distribution manifold, built for a radar application, which uses laminate, or softboard, PCB technology. This demonstrator has tackled such areas as wideband right-angle transitions and 1000-way microwave splitting, whilst maintaining the dimensional tolerances needed for an active antenna array. The structure contains multi-layer, screened, digital signal distribution in polyimide glass with stripline active networks and buried passive components in multi-layer PTFE on an aluminium carrier.

Fig. 9 shows a fibre-fed communications modules designed to perform the final drop (for example, lamp-post to the home) for interactive ATM services on RF carriers between 27GHz and 30GHz. The microwave filters and patch antenna arrays have been fabricated on softboard. Fig. 10 gives a typical response of the filters. The tolerance required here is 5% on dimension of <50μm. As a stand-alone substrate, this technology will have competitors; as a microwave circuit on a seamless fabric embracing an integrated system, it is unique.

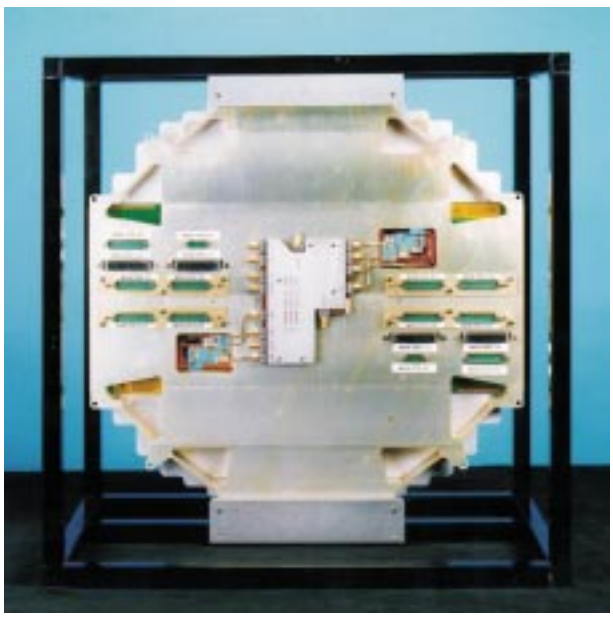

8 ***10GHz distribution system on softboard***

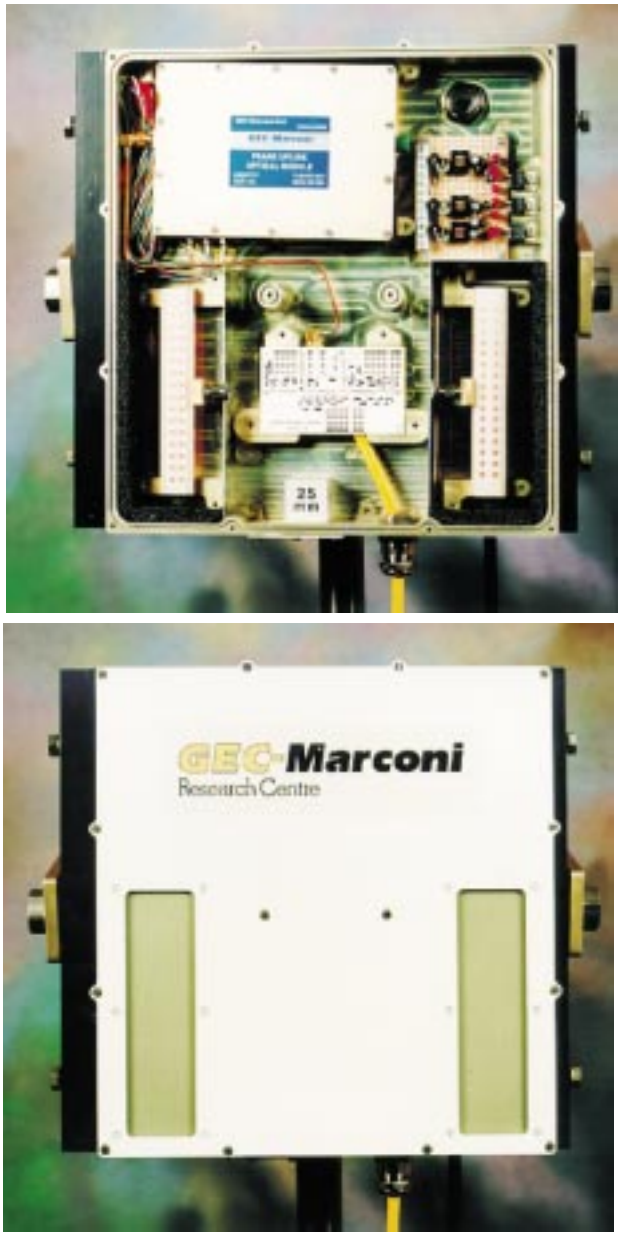

9 ***Communications unit providing a radio interface to a fibre network***

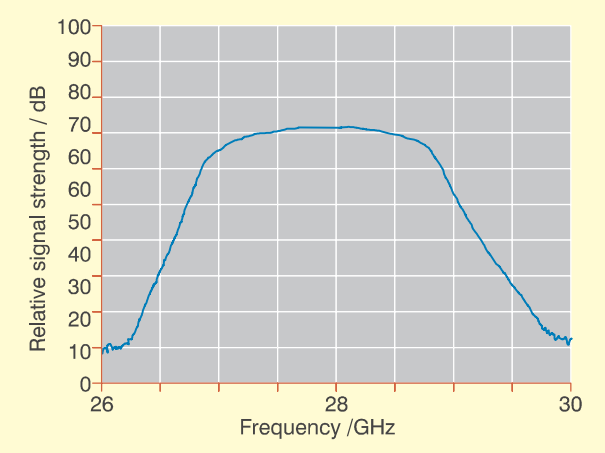

10 ***Response of a softboard filter***

Fibre-in-board

Another development has been in support of closer integration of electronic and optical integration. Fig. 11 shows an experimental structure in which fibre is embedded in one of the laminate layers that would typically serve as the dielectric separators between metallization layers. The technology has already been developed as an avionics backplane[5] and is currently under development for telecoms applications.

The benefits in terms of fibre management are important. The technology introduces fibre backplanes into systems in an evolutionary, rather than a revolutionary manner, which is also important. However, the key advantage from a systems viewpoint is that both optics and electronics can be integrated onto a common substrate. Whether the unit be a high-functionality module (such as a fully-integrated true-time-delay subsystem for a multifunctional radar) or part of the system infrastructure (for example, a high-speed, EMC-immune, backplane for a supercomputer), the system designer has access to both optical and electrical interconnect in a common medium.

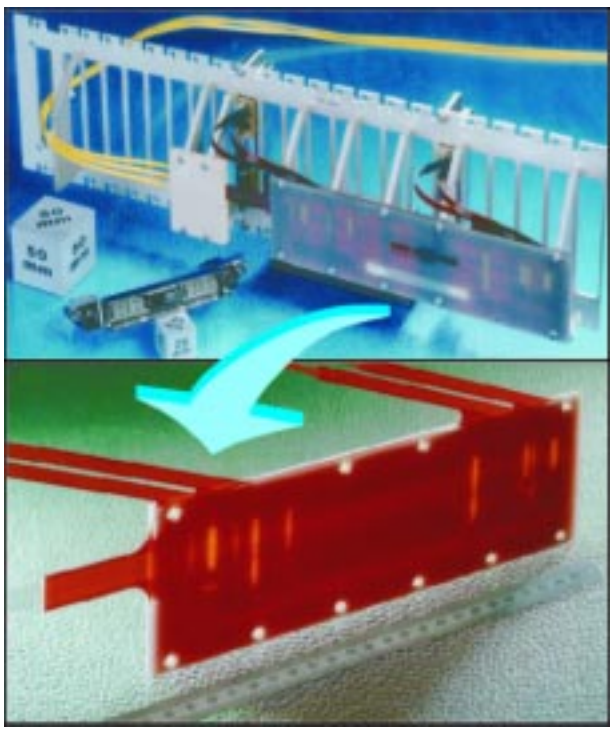

11 ***Laminated fibre-in-board (bottom) and backplane embodiment (top)***

Who Pays?

The key question here is, 'who benefits?'. In this paper, we have argued that the most significant benefits are achieved by addressing packaging earlier rather than later – in design, rather than in production. Furthermore, we have argued that the biggest benefits accrue in terms of life cycle costs, rather than purchase price; in system integration, rather than product manufacture. For these reasons, we believe that packaging is primarily a systems issue. The team at the Marconi Technology Centres, Great Baddow has pioneered advances in packaging and interconnect for the past fifteen years for precisely this reason.

Many would argue, however, that the board manufacturers are making the money from the technology and should therefore fund the R&D investment. Sage[6] argues cogently in this respect that the source of innovation funding depends on the state of the interconnect cycle. Where evolutionary change is underway, it lies with the board manufacturers – and particularly their suppliers – to squeeze new performance out of the existing product and to offer better value to the system customer.

However, in the process of revolutionary change, the board manufacturers, with their tight margins are unlikely to invest in new types of equipment without a clear view of the applications and volumes. The system integrator has exactly this insight and therefore the most to gain. At times of revolutionary change, system integrators must, then, support the changes. The current state of flux in packaging clearly lies on the revolutionary rather than the evolutionary side of that divide.

However, other factors must also be considered, especially because many large systems houses have dispensed with their traditional in-house manufacturing. They are thus dependent on subcontracted support and yet are being asked for financial support towards the sea-change in packaging needed to maintain competitiveness. Recent analysis of the supplier market[7] indicates that two types of board supplier will survive in future: the small contractor and the large conglomerate. The former will survive because of its fast turnaround and ability to provide a high-quality,

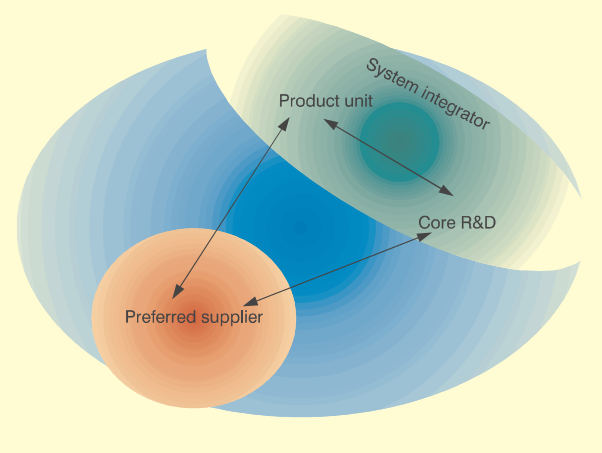

12 ***Interaction between a systems company and an interconnect supplier***

versatile service. The premium this service attracts compensates for the less competitive aspects of low-volume production. The mass producer (£100 million+ p.a.) will continue to grow because it can afford the capital investment and benefits from massive economies of scale. According to this analysis, everyone in between will be squeezed out and the big will continue to get bigger.

We believe the role of a system integrator in such a scenario is to maintain a leading-edge R&D team and to form preferred supplier relationships with one or more candidates from the two categories outlined above. This type of arrangement is shown in fig. 12 and frees the system manufacturer from the capital outlay involved in maintaining its own production, provides it with a vital degree of foresight in the field, and ensures, through technology transfer, that the right technologies are there in production when required. The flow of cash and intellectual property rights (IPR) involved in reaching such an agreement is beyond the brief of a paper such as this.

Conclusions

Conventional system design has placed packaging (that is, the technology that will be used to produce the final product) low on the list of priorities and late in the design process. We believe that packaging is more beneficially considered as part of the initial design. In particular, as part of an integrated design strategy, packaging delivers vital benefits in three areas: obsolescence, access to COTS, and in integrating the next generation of systems.

Appropriate packaging provides a response to obsolescence at the physical level where BGAs and MCMs can be used to ensure that a consistent function is delivered to the motherboard, despite changes in the silicon above the interface. At the more advanced level and in conjunction with FPGAs and the software design, it should be possible to ensure that the functional envelope of an evolving system embraces fully the original requirement and that upgrades are introduced in a controlled and manageable manner. That this is possible is evident, but R&D is needed to understand completely the impact of using such a strategy.

As far as COTS is concerned, advanced packaging can provide a route between wholesale use of COTS and bespoke systems. By repackaging COTS, one has opportunity to reap the benefits of proven systems whilst providing the right environmental capability for a given application.

Finally, integrated miniature systems are forced to rely on an increasingly complex array of integration technologies in order to bring silicon chips, sensors, interfaces and peripherals together. The Marconi Technology Centres, Great Baddow, has supported this drive by integrating technology for microwaves and optics into the basic PCB laminate processes.

Overall, packaging is a weapon in the system integrator's armoury. Various levels of systems edge can be obtained by appropriate deployment of the technology. In particular, teaming arrangements with the most competitive suppliers should enable the benefits of leading-edge R&D to be fed through in support of advanced systems that, in turn, will be better future-proofed and able to deliver improved performance over their life cycle.

Acknowledgements

Many of the demonstration units referred to have been partially supported by the EC through various Framework programmes, including CHIPPAC, FRANS and FLINT. Some UK Ministry of Defence (MoD) support is recognized under various initiatives, but specifically under AMSAR and A^{3}P.

References

1 New Electronics, 30th Anniversary Issue, August 1998.
2 CHANDLER, N. and TYLER, S.G., 'Ultra-fine feature printed circuits and multi-chip modules', *Microelectronics J.*, **26**, p. 393-404, 1995.
3 JOHNSON, R.W., TENG, R.K.F. and BALDE, J.W. (eds), 'Multi-chip modules – System Advantages, Major Constructions and Materials Technologies', IEEE, ISBN 0-87942-267-X, 1991.
4 DOYLE, R., O'FLYNN, B., BARRETT, J., SLATTERLY, O. and LAWTON, W., 'Chip-on-board for large area die – an assessment of low stress glob-top materials', *Journal of Microcircuit and Electronic Packaging*, **20**, 2, 1997.
5 JIBB, D.J. and GIDDINGS, B.J., 'Single-mode fibre-optic data networks for integrated modular avionics', ERA Report 95-0364 (Conference Proceedings), 1995.
6 BPA Report, System 2000, BPA, Dorking, Surrey RH4 1QT, 1996.
7 TYLER, R.L., 'Modelling and predicting the future shape of the PCB industry', Proc. Conference on Technology & Business Strategies for the Next Five Years: The Successful PWB Company, Florence, June 16-18, 1997.