By Glynn Carter,
Technical Manager,
Force Computers.

Enabling technologies for internet appliances

The technological leap-frog game, always providing us with solutions to problems we didn't realise we had, has now arguably provided us with our biggest ever solution - the internet appliance. Assuming that the network infrastructure keeps pace, engineering innovation can now provide us with a whole wealth of tools that we never dreamed we needed, but now cannot live without. The birth of the internet appliance has not brought with it any very new engineering challenges. If anything anengineer's life has become more mundane, with technology becoming increasingly modular. The innovative challenge is now choosing the correct and often unique combination of technologies.

TIME TO MARKET

Once the unique technology profile has been chosen for a particular internet appliance, time to market is the critical success factor in the market place today. The pressure on engineering is huge, with the innovation to product cycle ever shortening and successful execution non-negotiable. With engineering skills becoming more focused and design teams generally getting smaller, more companies are turning to eTOTS (enabling technologies off-the-shelf) to complete their design.

Historically, new products have been managed in-house from technology development, product design through to manufacture and service. Over the years it has become more cost-effective for companies to outsource elements of the process to specialist companies. Initially this design outsourcing would have centred around specialist power supply or enclosure requirements.

These enabling technologies can be bought in as silicon, for expanding the functionality of an existing networked appliance, such as enabling local level interaction on a mobile telephone. Alternatively, for a point application such as a home medical attendant or office based network attached storage, it can be a complete solution built up from a number of modular components.

With the advent of the 'Internet Appliance', the embedded computing market is broader, more dynamic and faster growing than ever. The choices and interdependencies at the operating system, computing platform, bus and i/o level are at first pass vast. However this choice, combined with the fact that outsourcing is now rapidly extending to include virtually all elements within a product design, can mean that companies face the risk of adopting potentially dead-end technology.

GLOBAL APPROVALS

Packaging has taken on a wholly new perspective. Not only must a product be globally safe, it must also be globally approved and has to look right. At Force, we design the hardware to fit the packing, not the packing to wrap the hardware. Often the initial concept is created by an artist with the ergonomics being rigidly set before the device even gets to the drawing board. This creates a new set of design challenges, around areas such as power dissipation, product safety and maintainability.

A recent proposal for an add-on modem for a popular handheld device left designers with an interesting challenge. The packaging forced the need for a horseshoe shaped circuit. The challenge was not necessarily in the electronics design, but in the design for manufacturing and test.

While industry-standard embedded software drives hardware over a wide variety of chip architectures, apparently simple issues such as big-endian/little-endian have always provided an artificial restriction to platform choice. Today, a well accepted operating system such as Linux can be viewed as being technology agnostic, running on platforms ranging from Pentium, SPARC, PowerPC and Alpha through to MIPS, Crusoe, ServerSet, 68K and StrongARM.

With the spreading acceptance of these technology-agnostic operating systems, a subtle hardware architectural difference between two platforms can mean, for example, that an apparently standard Linux PCI driver works on one, but is virtually unusable on the other. These subtleties are often not even encountered until the design phase has been completed and a team is well into product development.

Option interconnects have broadly standardised on the PCI technology. Assuming that this simplifies software development (it doesn't always), designers are then faced with a mechanical problem. PCI, while electrically stable, provides the designer with a choice of implementations of both bus width and speed. Having started with 32-bit at 25/33MHz, we now see a small but increasing number of options requiring 64-bit at 66MHz. Mechanically we can choose to implement PCI options such as video, sound, Ethernet, modem and RAID in several ways.

PCI, in either its commonly seen desktop and server options card or in the CompactPCI form factor, can give us potentially higher performance and densities, along with onboard processing capabilities. At the other end of the spectrum we have the device on a chip, which can save money and footprint at the cost of flexibility and stability. Between we have the mini-PCI and PMC (PCI Mezzanine Card) both of which are

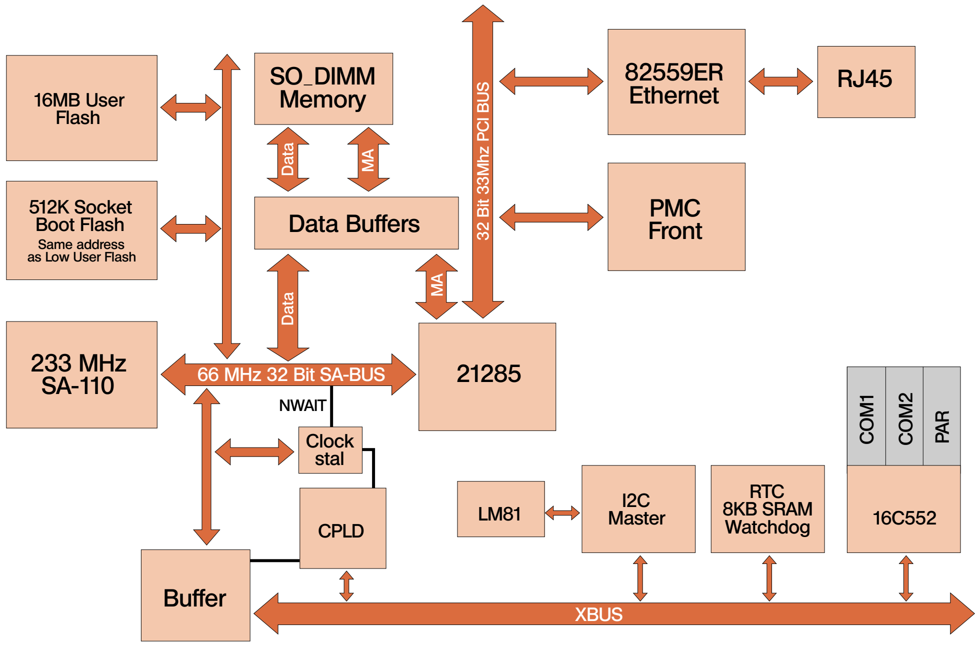

Figure 1.

aimed at distinctly different markets and have very little commonality of options.

Communications, while broadly standardised and internationalised, are far from straightforward. The simple modem, now well into its third decade, still provides us with the only common, globally accessible transmission medium. The customer expects speed of access whether it be by cable, wireless or local-loop. they have no sympathy for the designers looking at the global flexibility of a new design.

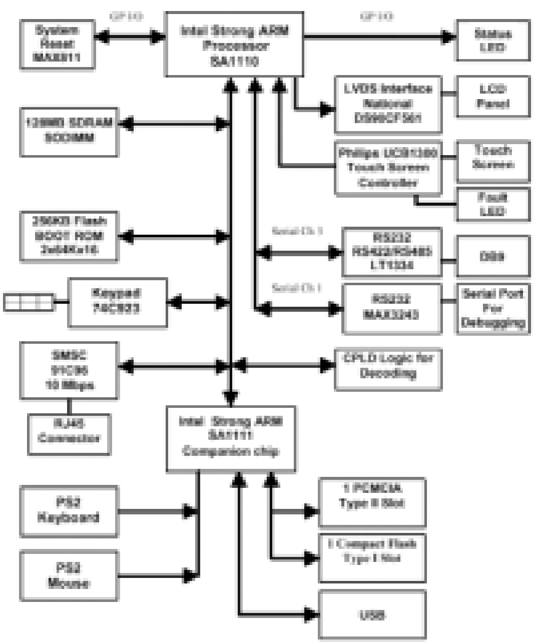

Figure 2.

VENDING MACHINES

As an example, if we take a look at what may initially appear as the simple requirements for the increasing number of internet enabled vending machines. One of the driving factors for enabling such a device is to limit the human interaction, cutting down on the number of unnecessary visits for cleaning, recharging of consumables, change refilling and reactive maintenance.

The design of figure 1 was aimed at 'farms' of ticketing machines, such as those found at airports, stations and exhibition centres. It gives a generic platform to implement the basic functions to limit human interaction, whilst still allowing for a small amount of flexibility through the PMC carrier.

Now compare this with an updated variant (figure 2), aimed specifically at 'improving' the usability of the ticketing machine by allowing implementation of multilanguage and multipayment support. A hidden benefit of this approach was to allow local advertising.

While the application is the same in both instances, the implementation is wholly different. Both solutions can be modified to encompass other similar applications, but how would you alter the platform so that it would allow for wireless consumer communication? Additionally, how could the design be improved to make the device capable of handling technologies such as RF, Ethernet, ISDN, ADSL and GSM for central systems communication when used as standalone devices in areas where the 'farmed' approach isn't applicable?

These basic requirements are being combined with specific interfaces with such as GPS to give time and position information (on-board ticketing for trams and buses), biometrics devices (high value payments) and speech recognition, as well as interfaces required for electronic payment and the last ditch 'push here for help' human interface.

Perhaps more importantly for this ticketing example, is the investment in software lost between the first and second generation products. Had the designers backed themselves into a technology cul-de-sac with their initial platform choice or was it an informed decision? Is the design likely to change radically for the next generation product?

All of these enabling technologies are available off the shelf; the experience required to glue the technologies together isn't. Whilst the number of point technology companies expands, companies who have the integration skills, the broad spectrum technology experience and proven track record to successfully implement a design together are limited.

Even when designed within the rigorous constraints imposed by the perceived or actual window of opportunity, successful and sustained execution is essential. You can be first in the market to demonstrate a prototype that is impossible to manufacture, test or maintain. Getting your innovation from concept to market is an added dimension to the internet appliance market ■

Glynn Carter has a degree in Engineering. He started his career in 1985 as a software programmer, developing software for the process control industries. In 1987 he joined Digital Equipment Company Ltd as a Technical Consultant working with various networking technologies. In 1995 and as a Principal Technical Consultant, Glynn moved into the Technical OEM Division of Digital working primarily with OEMs who were logically or physically embedding Digital's OEM products.
Glynn joined Force Computers in 1999 as the Regional Technical Manager and focuses on addressing the diverse technology needs of companies within the emerging networked appliance markets.