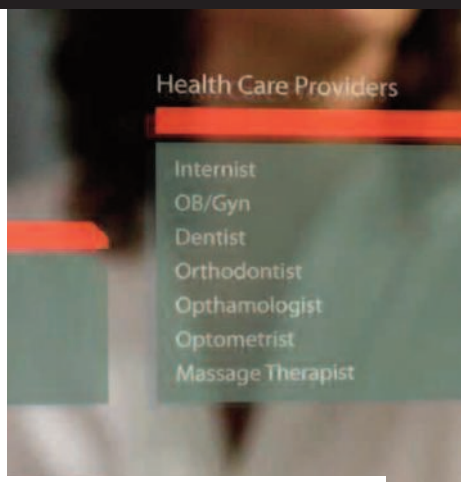

Things worthy of consideration

The Internet of Things is pushing microcontroller developers to move in unexpected directions. By **Graham Pitcher**.

It's not too long ago since the 'go to' microcontroller for electronics designers was a humble 8bit device, often featuring the 8051 core. But times are changing, as 32bit MCUs become the device of choice and 8bit parts begin to be phased out.

Those were changes already in progress, but the Internet of Things (IoT) appears to be having a major influence on how microcontroller technology develops.

In broad terms, the IoT comprises three elements: edge devices, which often perform one dedicated function; hubs or fusion devices, which integrate data from edge devices; and larger processing elements. It's the first two categories which are currently focusing the minds of MCU developers. The reason? The IoT demands two things above all others – minimal power consumption and the lowest possible cost. Thomas Barber, director of marketing for wireless products with Silicon Laboratories, explained why. "Estimates suggest there may be 30billion edge devices. They can't all consume 1W because we don't have enough energy, so they need to consume microWatts and they need to be inexpensive."

Andreas Eieland, Atmel's senior product marketing manager for flash MCUs, agreed. "Being able to have the right features at the right power consumption will be critical. Edge devices will need to run from harvested energy or for their full lifetime from a single battery."

Geoff Lees, general manager of Freescale's microcontroller business, added: "We are being directed towards more and more integration. What were low cost devices need to have more RAM, more resources, but we can't increase the power budget or the price. That is pushing us to move geometries such as 40nm and 28nm more quickly than we would have expected."

While the three companies agree on the general shape of the IoT market, where they differ is on the technology that will be applied at the various layers. But all agree that MCUs targeted at the IoT will need to offer wireless connectivity.

Lees: "We are being directed towards more and more integration."

Silicon Labs, for example, has just released the Si106x/108x range of wireless MCUs, designed with the IoT in mind (see fig 1). And sitting at the heart of these parts is an 8051 core; not the 32bit device you might expect. There's also a sub GHz RF transceiver and peripherals such as a 10bit A/D converter, comparators, 16bit timers and serial interfaces.

Barber said: "Every time there's an 'obit' for the 8051, it keeps going and we're not betting on it going away. There is a large base of developers who like it and it will continue to sell." But Silicon Labs can also offer devices based on the ARM architecture; not only its own designs, but also from recently acquired Energy Micro. "ARM sees the Cortex-M0 as an 8bit killer," Barber continued, "but it will come down to customer preference."

Eieland noted that power requirements will drive this element. "MCUs will have to consume less than 1µA and less than 200nA in deep sleep. Whatever 'wins' will have to offer 8bit type performance.

"Atmel makes 8051 based MCUs, but if you want to connect to ZigBee, for example, our AVR cores are a better choice." Atmel's offerings here include devices from its ATMega range. "But we also have a Cortex-M0+ part that competes in that sector."

Lees noted: "You can make the 8051 work, but what happens when you have a lot of sensor data which is more than 8bit? We see the Cortex-M0+ as a major platform for edge node devices, but we're already seeing more products with what would have been thought of as an excess of

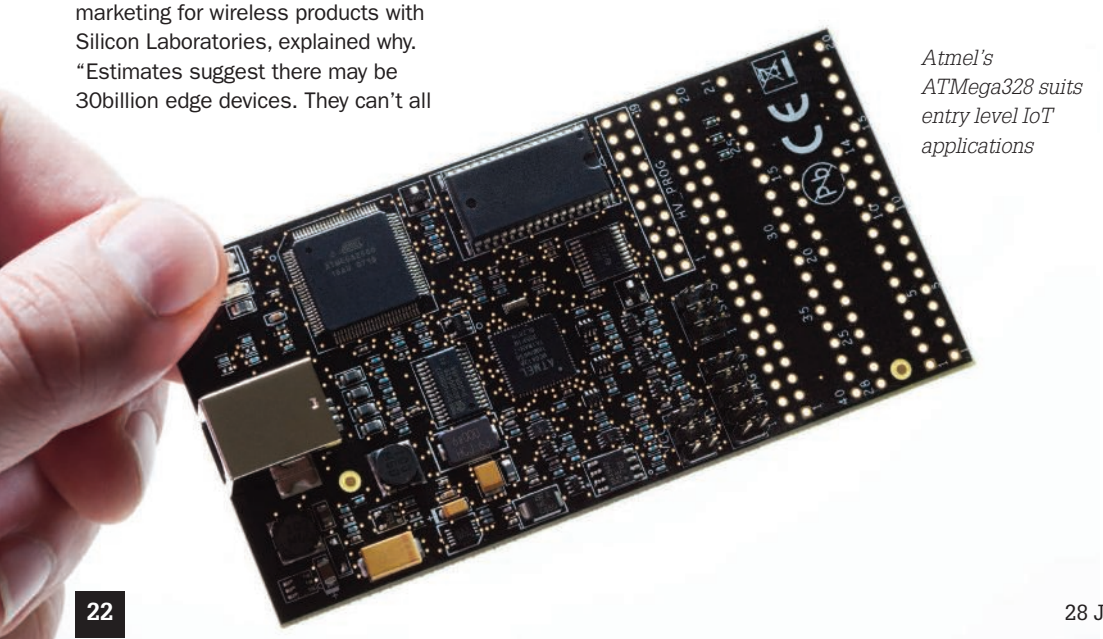

Atmel's ATMega328 suits entry level IoT applications

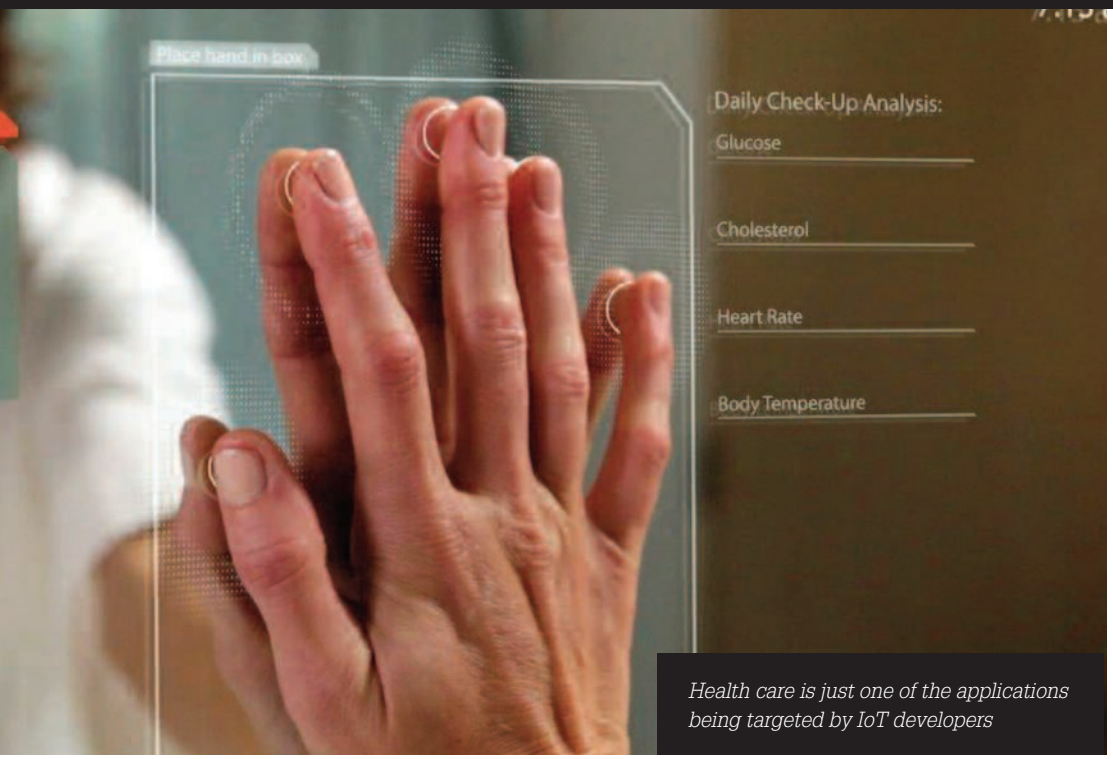

Health care is just one of the applications being targeted by IoT developers

processing." For the moment, Freescale is targeting its Kinetis-L range of M0+ based devices at edge devices, but Lees perceives that, even at the edge, there will be the need for more performance. "Where there would have had one processor, we're now seeing an M0+ as sensor hub and a Cortex-M4 as the apps processor handling communications."

What is also generally agreed is that the standalone microcontroller will be, to some extent, superseded by an SoC implementation. Barber said: "We'll see devices with lots of peripherals and memory and we won't have one chip that meets all needs." And he believes that, in the long term, integration on die will be preferable to in package integration. "It's always better to integrate on die; it's cheaper and a better solution."

Lees pointed out: "As fast as you can integrate, there are more sensors – and some we've not heard of before. So these parts will have sensors, low power connectivity and sufficient processing."

In terms of in package integration, Atmel has worked with Bosch Sensortec on the development of the BNO055, said to be one of the first application specific sensor nodes to be announced. "There are six dice in the package," Eieland noted, "including a SAM D20 MCU. It's a good example of how we are working with sensor manufacturers to get the most size efficient solutions possible."

Barber: "Every time there's an 'obit' for the 8051, it keeps going."

While Lees is confident that MCUs for IoT applications will need to be made on low geometry processes, Barber and Eieland are not so sure. Eieland thinks refinement of existing technology may be sufficient. "We don't think we want to go to 55nm because leakage at that node will be significant. Processes in the range from 100nm to 150nm may well be suitable, with one more product generation needed to evolve the low power aspects."

Barber implied that smaller geometries may not be at the top of Silicon Labs' agenda: "We will win the process technology battle assuming we've done the right things to control leakage. Architectures will become more important, with more efficient computation. The longer you can keep the MCU powered down, the better off you will be, so the device will need to come in and out of sleep modes and to feature autonomous peripherals which don't wake the core up. We're still trying to maximise battery life - energy consumption is critical. How you have low consumption for a device with a lot of functionality? But I think another two device generations will get us to the price and power points we need to hit."

Lees noted: "Freescale has already moved Kinetis to 90nm and there will be a more rapid move to smaller geometries. We are fighting leakage in deep sub micron, but we are putting a lot of work into developing libraries, power gating and so on. All of this work is now becoming a fundamental part of MCU design.

"The power budget is a continuing discussion point in all my design teams; it limits what we can do and directs our road map towards next generation energy efficient cores."

And that road map could be interesting. "Although Freescale is a lead partner with ARM for its Cortex-M cores, we are beginning to see more IoT applications suited to low end Cortex-A processors, rather than high end Cortex-M cores," Lees concluded. "Application processors are coming to the domestic market, and that's something that was unexpected."

Fig 1: Simplified block diagram of Silicon Labs' Si106x/8x wireless microcontroller