

# Slow light, fast computers

As everything else seems to get faster, why are scientists trying to slow light down?

**David Gevaux** spoke to Fengnian Xia and Yurii Vlasov at the IBM Thomas J. Watson Research Center to find out.

## Why all the interest in optical information processing?


Our major interest is optical interconnects in future high-performance computers. As the speed of computers increases, connecting different chips within a computer or even different parts within a chip becomes a big issue owing mainly to bandwidth limitations of electrical wiring. Photonics provides us with a solution as it can intrinsically handle a huge data rate. For example, it has been shown that a single optical fibre can handle data rates larger than one terabit per second. Although such high-speed optically interconnected networks on a chip are a long way off in practice, the allure of such bandwidths stimulates a great amount of interest within IBM and in the microelectronics industry in general. This interest is also seen from the support our work was given from the Defense Advanced Research Projects Agency within the programmes on Slow Light and on Intrachip Interconnects.

## What role do optical buffers have to play?

In any network, one of the problems is signal congestion at the switching nodes. The use of optical buffers can effectively avoid congestion in optically interconnected networks by temporarily storing light signals for the time needed to resolve the congestion — a sort of optical memory. There are interesting parallels here with the early computers — in the 1950s, electrical delay lines based on acoustic waves moving through rods of mercury were used as memory in the first programmable machines. The development of optical buffers removes the need to convert optical signals into electronic signals to do the buffering. Instead, the signal can be kept in the optical domain, which ultimately means that higher networking speeds and better overall performance may be attainable.

## What makes for a good optical buffer?

The first thing is the buffering capacity (or the number of buffered bits), although the exact capacity required is dependent on the application. For real buffering applications



Yurii Vlasov, Lidija Sekaric and Fengnian Xia in their slow-light lab.

in optical networks, we need hundreds of bits at least, which corresponds to delays of tens of nanoseconds at data rates of several gigabits per second. The maximum capacity achieved in our work is 10 bits, so we are still a long way off, but we believe that our results (p65) represent a major step in the right direction. In addition, low loss, tunability, and a small footprint [device size] are also very important. For on-chip applications the footprint is obviously of utmost importance as it allows the integration of optical buffers with other optical elements on the same chip.

## Why take the optical resonator path?

Our aim here at IBM is to build a toolbox of devices, not only optical buffers but switches, modulators and routers, with a view to building optical networks on a chip. The use of optical resonators has allowed us to shrink the device footprint, leading to 10 bits of buffered information within an area less than 0.09 mm<sup>2</sup> (and this can be further reduced to 0.03 mm<sup>2</sup> after optimization). Compatibility with current complementary metal-oxide semiconductor (CMOS) technology is

also vital, hence the devices are fabricated in a standard CMOS fabrication line and might eventually be integrated with microelectronic chips. Our approach also has the advantage that it uses a completely passive structure; it does not require external lasers or an external electrical power to drive the system to its slow-light state — it is built in silicon and nothing else is required.

## In what ways can these optical buffers be improved?

A major limitation at present is the insertion loss. Although the current structures are the result of a lot of work on design and fabrication optimization, for real buffering applications the losses are still too high (about 20 dB of loss for a total buffering of 500 ps). To make optical buffers practical we must either find a way of decreasing the intrinsic losses or come up with a design that will be less sensitive to fabrication errors. Another approach is the implementation of amplifiers to offset the losses.

A second improvement would be to increase the attainable operational bit rate. At present, our delay line can work error-free up to 5 Gbits s<sup>-1</sup>. Recent designs of optical resonators can lessen the negative impact of dispersion and enable higher operational bit rates to be achieved. A final requirement is the ability to change the delay — tunable buffering. In 2005 we demonstrated that an integrated micro-heater can be used to achieve this with a response time as small as 50–100 ns.

## Do you think optical networking will one day completely replace electronic networking?

For on-chip networking it will not replace electrical signals completely, rather, photonics will be used in synchronization with electronics to boost performance and perhaps eventually result in a supercomputer-on-a-chip.

*Vlasov and his co-workers have an article on optical buffers on p65 of this issue.*