Data-acquisition and triggering with transputers

J. C. Vermeulen
NIKHEF-H, P.O.Box 41882, 1009DB Amsterdam

Properties of the transputer

A transputer of the current generation (1) is an integrated circuit with a 16 or 32-bit CPU, a fast on-chip memory (2 or 4 kByte) and interfaces for 4 bi-directional serial 20 Mbit / s point-to-point links with DMA support. Each link consists physically of two wires. All links can be used simultaneously. Each link can transport data in both directions simultaneously. The throughput per link per direction is of the order of 1 MByte / s (1.77 MByte / s at maximum). The micro-code of the CPU contains a two priority level multi-tasking kernel. It provides support for process scheduling, inter-process communication and descheduling of processes while waiting for timers or for an external signal. The communication mechanism is synchronous: a process is descheduled when attempting to communicate with another process and is made executable again when data transmission (by the CPU when both processes run on the same transputer, or under DMA control via a link) is completed. The 32-bit CPU's computing power is comparable to that of a Motorola 68020 CPU for integer operations. Some transputer types have a floating point unit with a performance of about 1 million floating point operations per second. Task switching takes at maximum 4 μs. A transputer has a 16-bit address / 16-bit data or 32-bit address / 32-bit data external bus interface for connecting external memory and interfaces. Transputers can be programmed in occam (2), a language with support for multi-tasking and inter-process communication, or a.o. in a parallel dialect of C or FORTRAN.

Systems can be built from single transputers - which may be equipped with external memory and / or dedicated interfaces - by connecting them together via the links. The links provide not only inter-processor communication, but also a boot path, as a transputer can be booted after a reset via one of its links. The links are fast enough for many present high-energy physics on-line applications. It has been shown to be possible to make them 100 m long, when using balanced drivers and receivers and high-quality twisted pair cable.

The functionality of systems built from transputers can be augmented with 32 by 32 crossbar switches for links, which are available from INMOS as single integrated circuits. Also available are integrated circuits for interfacing a transputer link to a parallel microprocessor bus. These devices are functionally comparable to simple conventional interfaces for asynchronous serial links ("ACIA's"). Finally, at NIKHEF-H an uni-directional fan-out circuit for links has been developed, that can be used for broadcasting messages to up to 16 destination transputers.

Present applications of transputers in high-energy physics

At present three high-energy physics experiments are or will be using transputers for data-acquisition and / or triggering:

- The JETSET experiment at LEAR, starting data taking, applies three transputers for event building (3). The transputers communicate with VALET-PLUS systems via the parallel bus to link interfaces mentioned earlier,
- The UA6 experiment, also starting data taking, uses about 60 transputers for data read-out, event building and third-level triggering (4). Third-level triggering is not done using the conventional farming approach, i.e. assigning the data of one event to a single transputer, the data of the next event to the next transputer, etc. . In stead different transputers execute different parts of the algorithm, while the data is distributed accordingly,
- The ZEUS experiment at HERA, scheduled to start data taking in August 1990, is the largest scale application with an estimated total of about 550 transputers, which will be used for data read-out, second-level triggering and event building (4, 5). About 400 of the transputers will be residing in VME-crates in so-called 2TP-VME modules. This type of module has been developed at NIKHEF-H and contains a.o. 2 transputers, which both can access the VME-bus. The systems of ZEUS will be partially dynamically reconfigurable. For data-acquisition the crossbar switches will be used for multiplexing data from many links onto a few links. In the event-builder crossbar switches are used to route the data to the correct third level crates. In all cases the crossbar switches are controlled by 16-bit transputers, with parallel bus to link interfaces attached to the external bus. Every transputer that needs to output data via the crossbar switch has a direct link to one of the interfaces for sending switch requests to the 16-bit transputer and for receiving messages acknowledging that the switching has been done. The NIKHEF-H broadcast circuit will be used a.o. for distribution of trigger decisions.

Programming of these transputer systems is done mainly in occam.

For a newly proposed high rate experiment at the CERN SPS it is intended to use transputers for event-building and third-level triggering and for interfacing (via SCSI) to digital audio cartridge tape drives (6). Other projects, not directly connected to an experiment, consist of the development of a DSP + transputer combination (7), the development of a combination of an Intel 860 processor, an Intel 960 communication engine and a transputer (8), and work on embedding transputers in the Fastbus environment (9).

These applications prove that transputer technology makes it possible to construct high-energy physics data-acquisition and trigger systems in a modular and natural way on account of the transputer links and the synchronous communication mechanism, supported by the on-chip multi-tasking kernel. The synchronous communication is very helpful in the data-driven environment of a data-acquisition system. There is no need for complicated interrupt or handshake schemes, as a process sending data to another process, that could be located in another processor, will only continue when the data is received (time-outs can be provided, if necessary). The on-chip multi-tasking kernel furthermore facilitates to run monitoring processes as background processes.

The new generation of transputers

In 1991 Inmos intends to start marketing a new type of transputer with code name H1 (10). This device is expected to offer 5 - 10 times the computing power of the present 32-bit transputers and to have a peak performance of 20 million floating point operations per second. The instruction set is a superset of the instruction set of the present T805. There are again 4 bidirectional links, but now physically consisting of 4 wires. The throughput per link and per direction should be 10 MByte / s. It will be possible to couple present transputer links via protocol converters to H1 links. The internal memory is increased to 16 kByte and can also be configured as cache memory. The on-chip real-time kernel will be somewhat extended, so that the implementation of a multi-level scheduler on top of the existing facilities will be eased. There will be some form of memory management. An important innovation, illustrated in fig. 1, is the introduction of hardware support for "virtual channels". Consequently it will be possible to define communication channels between any arbitrary pair of processes, regardless on which transputers they are executing. The hardware will take care of multiplexing several "virtual channels" on a single physical link, if necessary, and will also take care of routing the messages. For routing a method is used that is known as "wormhole routing". A path between source and destination is opened by the first packet of a message. This path is also followed by all subsequent packets of that message. The last packet closes the path again. With this technique no buffering of data at intermediate nodes is needed. A packet contains 32 bytes. However, the last packet of a message can be smaller. A handshake mechanism (one acknowledge per packet in stead of one acknowledge per byte as is used in the present transputer generation) guarantees that the synchronous communication model is obeyed. It is possible to randomize the route taken to a certain extent in order to avoid the occurrence of "hot spots".

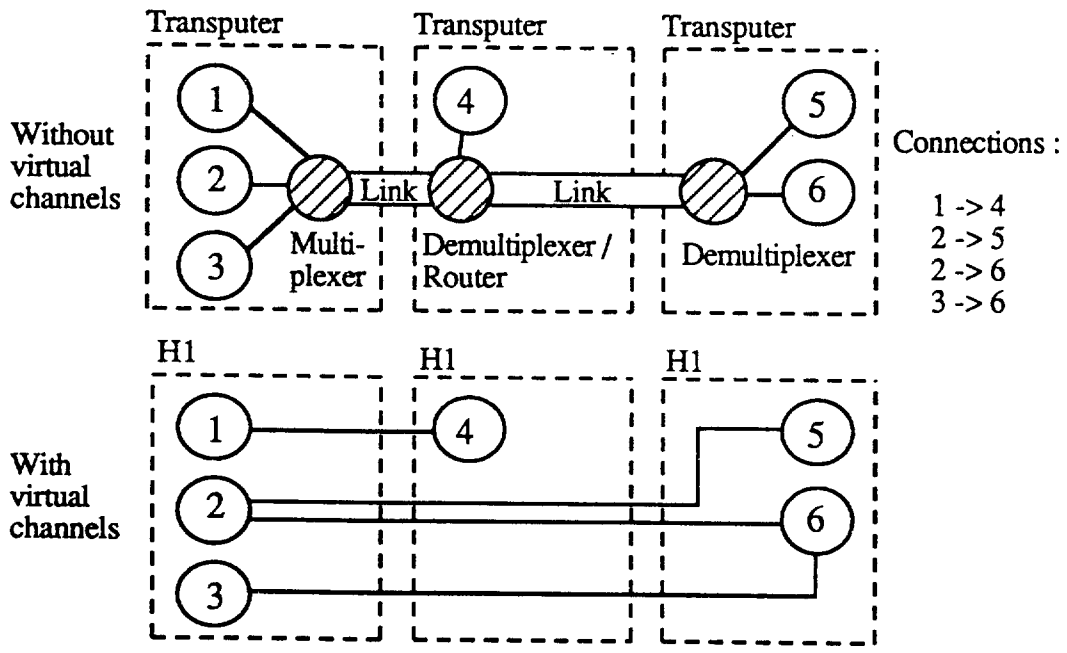

Fig. 1: For the present generation of transputers multiplexer, demultiplexer and router processes have to be applied to make communication possible between processes executing on arbitrary processors. The programmer of a system of H1's can make use of virtual channels between any pair of processes, i.e. the hardware takes care of any muliplexing, demultiplexing and / or routing needed.

After the introduction of the H1, it is expected that a new 32 by 32 crossbar switch will become available, which also will provide hardware support for routing. Only when starting up the information necessary for the routing of messages has to be loaded in the device, i.e. there is no need for a transputer controlling the switching actions of it. By combining the new crossbar switches with H1's it becomes possible to construct large systems with relatively high bandwidth and low latency inter-processor communication between arbitrary processors in the system.

The new generation of transputer hardware thus supports at the level of the application software a transparent mechanism for inter-processor communication. Point-to-point links in combination with the new routing crossbar switches guarantee an efficient use of the available bandwidth and also a low latency. These properties are very desirable for data-acquisition and trigger systems for LHC experiments.

Concluding remarks and outlook

The transputer is characterized by the integration of communication engines and a CPU - with a multi-tasking kernel in micro-code, which supports synchronous inter-process(or) communication - on a single chip. Its use for data-acquisition and triggering in three high-energy physics experiments demonstrates that its properties satisfy the requirements of these experiments in a natural and scalable way. Next year an order of magnitude faster transputers with extended functionality can be expected to be available. Especially the support for "virtual channels" is of interest. On basis of their functionality, the H1 and also the new crossbar switch are attractive components for applying in LHC experiments. As the H1 will be 1991 technology, even more advanced devices featuring faster operation, higher data transfer rates and perhaps also a further extension of functionality may be expected to become available during the design phase of these experiments. A recipe for implementing a functionally equivalent alternative is not simple. It would consist of building a transputer-like device, probably from one or more powerful CPU / CPU's, capable of fast task switching, and coupling it / them to suitable intelligent interfaces, capable of autonomous data transfers, to fast point-to-point connections. Also autonomous transfers between interfaces in one device, necessary for the implementation of virtual channels, should be possible. Furthermore a suitable multi-tasking kernel should be applied. This of course can all be done, but does not seem to be an attractive approach when sufficiently powerful real transputers are commercially available.

References

(1) INMOS, The transputer databook, second edition, 1989.
(2) INMOS Limited, occam® 2 Reference Manual, Prentice Hall, 1988.
(3) A. Bogaerts et al. , "The JETSET event builder", paper presented at the workshop on on-line applications of transputers in nuclear and high-energy physics, 2 March 1990, CERN.
(4) R. W. Dobinson, J. L. Pagès and J. C. Vermeulen, Particle World, to be published.
(5) L. W. Wiggers and J. C. Vermeulen, Comp. Phys. Comm. 57(1989) 316 - 320 .
(6) G. D. Barr et al. , CERN/SPSC/90-22.
(7) D. Crosetto, CERN-DD/89-33.
(8) R. Belusevic et al. , these proceedings.
(9) B. Martin, "Transputers in Fastbus", paper presented at the workshop on on-line applications of transputers in nuclear and high-energy physics, 2 March 1990, CERN.
(10) Dick Pountain, Byte, April 1990, "Europe and World" supplement, p. 3 - 12.