QuickRing Delivers High Bandwidth

Low-Cost Controller Connects Chips, Boards, Systems at 200 Mbytes/s

by Curtis P. Feigel

By cleverly combining technologies, National Semiconductor has introduced QuickRing, a refined interconnect scheme with a raw data-transfer speed of 200 Mbytes/s. A QuickRing of only moderate cost can support system throughput greater than 1 Gbyte/s. Even before the company announced production quantities of its QR0001 controller, QuickRing had garnered design wins in high-speed LAN routers and hubs. Now the technology is moving into servers, workstations, and personal computers.

QuickRing is at once an architecture, a hardware specification, and several layers of communication protocols. Developers can incorporate the controller chip into existing card designs, allowing QuickRing connections on cards to be used in conjunction with standard buses. This auxiliary, or "airplane" interconnect can boost the data-handling capacity of systems by an order of magnitude or more.

Progress has been ongoing in the realm of processors: clock speeds and data-handling capacities are greater than ever before. Less progress has been made, however, in interconnecting those processors to peripherals, and to each other. With the emergence of multimedia, virtual reality, and various kinds of image processing, transferring data throughout a system—especially off-card—creates a bottleneck.

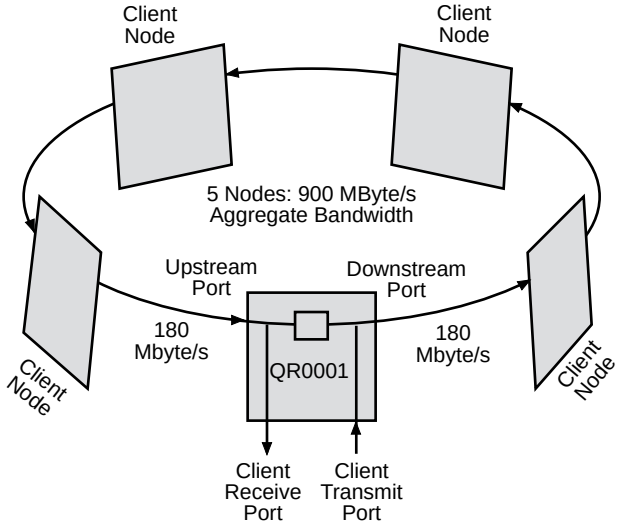

Figure 1. The QuickRing controller provides a high-speed interface to other processors and peripherals. Up to 16 nodes can reside on one ring, with each issuing and receiving multiple data streams.

One solution is to use wider buses; another is to use high-speed serial links. QuickRing provides a third alternative: use the high-speed signalling typical of a fast serial link, but use it impressed onto what is essentially a narrow bus connecting two devices. As shown in Figure 1, QuickRing arranges multiples of this link in a unidirectional loop.

Overall, the scheme achieves transfer rates similar to a wide bus or a fast serial link but at lower cost, and provides error detection and traffic management as well. National claims QuickRing is aimed not at replacing bus-type architectures but as a compliment to them, and that QuickRing can provide a standard, seamless interconnect for use between ICs, between boards, and between systems.

The concept for QuickRing originated with Paul Sweazey while he was with Apple's Advanced Technology Group. National helped refine the concept and developed the QR0001 controller chip, while BetaPhase (Menlo Park, Calif.) designed the special connectors and flex-circuit "cable" to handle the high-speed signals. As the advantages of QuickRing are recognized, more players have joined the group, including Molex, Amp, and Woven Electronics. Sweazey himself has now joined National.

Design Beats Bus Limitations

QuickRing overcomes one obvious limitation of the common bus: it is not limited to carrying a single transaction at a time. A QuickRing ring can contain up to 16 nodes, each capable of transmitting and receiving at the same time, so it's conceivable that one ring could carry 16 transactions simultaneously. The trade-off is that latency increases for each node between the target and the destination, although this can be mitigated if the application favors a certain arrangement of nodes on the ring.

Other not-so-obvious limitations of the common bus appear as the clock speed increases. At high clock rates, long signal paths with multiple stubs become more susceptible to transmission-line effects such as ringing, reflections, and crosstalk. Also, noise margins erode, and ground bounce becomes more difficult to control in the devices driving the bus. QuickRing's designers attacked this problem by using a point-to-point connection, carrying differential signals with a small voltage swing and controlled rise and fall times, all of which minimize EMI and maximize noise rejection.

The result is an interconnect with a raw transfer rate of 200 Mbytes/s and a sustained data rate of 180 Mbytes/s. Taking into account protocol and management

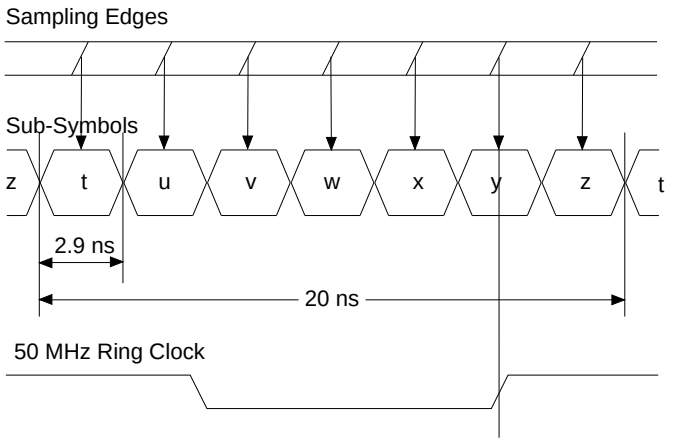

Figure 2. The ring segment carries seven sub-symbols of six parallel bits each during every cycle of the 50-MHz clock.

overhead, National's simulations place the aggregate maximum throughput of a 16-node ring at over 1.7 Gbytes/s. The practical throughput of a four-node ring ranges from 180–640 Mbytes/s. In all, QuickRing may have higher latency and more protocol overhead than a simple bus, but because of its pipeline nature, it is ideal for high-bandwidth streams of data, such as digitized video or audio.

Ring Has Advantages for Data Streams

QuickRing's ring topology brings with it a mixed bag of advantages and limitations. Because a ring allows multiple transfers to take place concurrently, the throughput of the system can be quite high. This characteristic points to streaming-data applications where, as with any pipeline, the latency is conquered once the stream is started.

Perhaps the best scenario for QuickRing occurs when the physical arrangement of nodes on the ring reflects the order of processing. For example, a certain image-processing application may require data decompression, then some floating-point operations, followed by some integer operations, before finally being displayed. It would be a straightforward task to arrange special-purpose processors on the ring in that order, allowing image pixels to stream from one to the next. On the downside, latency can exceed a microsecond if the nodes are arranged particularly poorly, or if the application has no inherent processing order.

Small Voltage Swing Is Essential

QuickRing derives its signalling technology from lessons National learned while helping define SCI. QuickRing uses the LVDS (low-voltage differential swing) scheme embodied in the IEEE 1596.3 standard. This specifies ±150 mV thresholds, with a maximum swing of ±400 mV and a smooth trapezoidal waveshape. This small voltage swing reduces the high-frequency energy contained in the signals, minimizing crosstalk and RFI. It also means that the signals can be driven quickly, even if capacitive loading on the line is high. Because they are differential, QuickRing signals have good common-mode noise rejection.

Data is represented by the wavefront of a voltage transition propagating down a wire. Multiple wavefronts sent down the same wire, separated by a few nanoseconds, can successfully be decoded at the other end.

Rather than trying to route a single master-clock signal all the way around the ring, one node is designated to originate it, then it is passed from node to node and regenerated at each step. This important because it allows each node to guarantee that the clock and data it transmits are synchronous. The clock is a 50 MHz differential signal called DnCLK at the transmitting node, and UpCLK at the receiving node. Even though the data signals occur much faster than DnCLK, all are guaranteed to have a specific phase relative to it.

To move these high-speed signals from one node to the next, QuickRing uses a cable with etched copper circuit traces on a polydimide plastic substrate—a flexible printed circuit. There are seven pairs of signal lines: six

Price and Availability

The QR0001 QuickRing controller chip is currently shipping in production quantities. It is priced at $45 per unit in lots of 1000. For literature and data sheets, contact National's Customer Response Center at 800.272.9959. Membership in National's QuickRing Developer program is free, while the complete QuickRing Developer Kit will be available in 1Q94 for $1000. For more information, contact Sean Long at 408.721.3046.

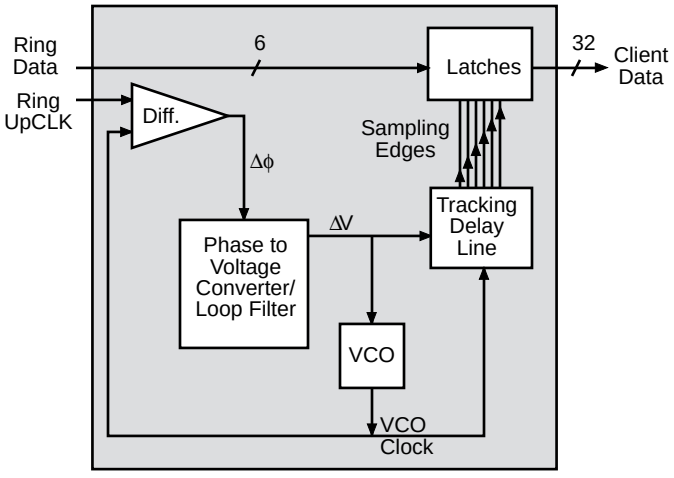

Figure 3. in this detail of the QR0001, a multi-tap delay line creates seven sampling edges from the ring's single 50-MHz clock.

differential pairs for data, and one differential pair for clock. This construction creates a very regular, controlled-impedance signal path that can carry QuickRing data with little loss in quality. Special edge connectors clamp securely onto the client circuit card while maintaining accurate spacing with a 66-pin-per-inch density.

Figure 2 shows how, during each cycle of DnCLK, a QuickRing segment carries a 42-bit symbol (32 bits of Data, seven bits of EDC, one Frame bit, and two Control bits). This symbol is transmitted as seven sub-symbols, each six bits wide. While this may seem an unusual way to break down the data, National's analysis showed it to achieve the best trade-off between data width and data rate. Aligned with the next-to-last sub-symbol, a DnCLK transition prepares the receiving node for the start of the next 42-bit symbol.

A 50-MHz clock is not nearly fast enough to directly support QuickRing's data-transfer rate, which requires a sub-symbol every 2.9 ns. Instead, UpCLK is passed to a multi-tap delay line, as shown in Figure 3, which produces accurately-timed clock edges to operate seven separate sets of six-bit latches. Each latch operates at only a 50-MHz rate, and signals on the ring are sampled at a 350-MHz rate. The complete 42-bit symbol is easily assembled into parallel format. Basically, the 50-MHz clock serves as a way to keep the sampling edges at the receiving node in phase with the transmitting node. This clock also feeds the QR0001's core, and is eventually retransmitted as the DnCLK serving the data sent to the next (downstream) node.

Seven bits of every 42-bit symbol are devoted to error handling. The QR0001 uses Hamming Code to provide single-bit error correction and double-bit error detection, although initial samples of the chip will not actually perform correction.

Controller Handles Data Streams

Connecting two cards, or "clients," via QuickRing is logically equivalent to placing a large FIFO buffer between them. To the logic on each card, the QuickRing interface appears as a separate transmit and receive port, each 32-bits wide (see Figure 4). A simple handshaking scheme indicates when the transmit buffer is getting full, or when data is available at the receive buffer. Each client port may be clocked separately at up to 50 MHz.

When the client begins to transmit, it first writes a 32-bit wide "head" word followed by the stream's 32-bit "payload" words into the QR0001's transmit port. The head holds the node IDs of the data stream's source and destination, as well as other information that uniquely identifies the stream. This data stream may be any length; data is considered to be part of the original stream until the client writes a new head. The client can multiplex several data streams onto the ring by writing a unique head for each stream.

Data coming into the transmit port is resynchronized to account for any frequency or phase differences between the client clock and the ring clock. Then, based on the contents of the head, the data is routed to one of three independent FIFO buffers. The X and Y buffers are each meant to handle one independent high-speed data stream. The LB buffer is intended for low-bandwidth transmissions, but can handle multiple streams, each with a unique head. The controller will direct a given data stream to one buffer only. If the data stream fills up the buffer, the controller will signal the client to stop transmitting until more space is available.

The controller also watches the amount of free space in the ring-transmit pipeline. When 20 open entries remain, it negates the TxOK signal to indicate that the client must stop transmitting data within 20 transactions. As the transmit pipeline clears, the controller asserts TxOK and the client may resume sending.

No Limits to Stream Length

Although the design sets no bounds on the length of the data stream, data is put onto the ring in packets with a maximum length of 21 symbols (one head symbol, and one to 20 payload symbols). The end of a packet is indicated with a tail symbol. This is simply a payload symbol (containing user data) that has a special flag set.

Packeting the data allows intervening nodes the chance to transmit in-between packets already on the ring. The controller attempts to concatenate payload symbols from the same data stream to fill a packet, but does not wait for any specific amount of data to be loaded into the buffers—it puts data onto the ring as quickly as possible.

As a packet comes from the ring into a node, the

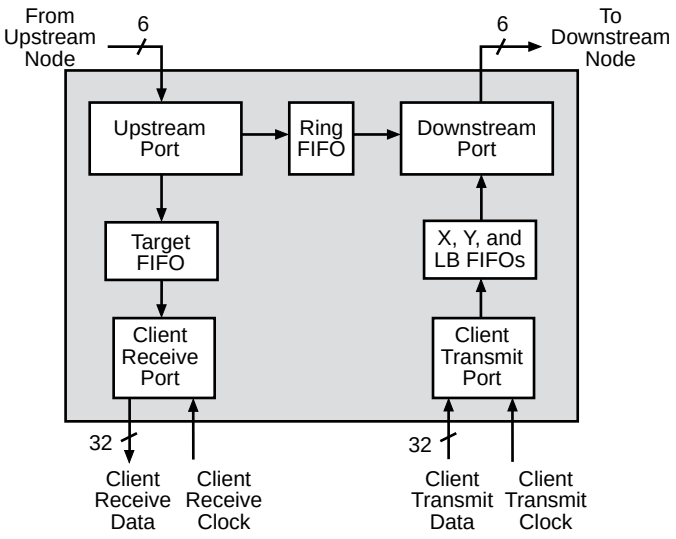

Figure 4. Data coming off the ring can either be routed to the node's client circuitry or forwarded to the next node.

	QuickRing	VL Bus	PCI	NuBus
Typical Bandwidth	>350 M/s	120 M/s	100 M/s	20 M/s
Slots	16	2	5	>8
Client Width Approx. Gates	32 40,000	32 1,000	32 4,000	32 1,000
Motherboard Cost	$36	$47	$63	$44
Daughterboard Cost	$36	$2	$2	$22
Bandwidth per Dollar	4.86 M/s	2.45 M/s	1.54 M/s	0.30 M/s

Table 1. QuickRing can out-perform common buses for bandwidth-per-dollar. (Source: National)

head is examined to determine if its destination is this node, or some other node around the ring. Data for this node is routed to the client through the controller's receive FIFO; data for other nodes is shunted into the controller's ring FIFO so it may continue its journey around the ring.

In the receive block, logic removes redundant packet heads, concatenating payload symbols into as long a stream as possible. This data is held in the Target FIFO, which is 75 words deep. From here, the client logic can retrieve the data words in sequential order. A resynchronizer block takes care of frequency or phase differences between the ring clock and the client clock. The data stream appears at the client receive port with a single head, with the payload words in their original order. Again, no bounds are set on the length of the data stream. If more than one stream is sent, either from single or multiple nodes, each stream will emerge preceded by its own head.

Vouchers and Tickets Help Manage Traffic

The QR0001 controller chip implements an intriguing protocol to ensure that a destination node has room in its FIFOs before a source node transmits data. This is also the source of much of QuickRing's latency.

In using this scheme, the controller generates a "voucher," a special control symbol put onto the ring, as soon as it receives a head word from its client. The voucher is targeted at the destination node specified in the head; it circulates the ring to the destination node, telling that node that the source has at least one payload symbol to send. If this node has room in its receive buffer, it issues a "ticket," another special control symbol put onto the ring, which tells the source node to transmit the data. Vouchers and tickets may be nested within packets. The key is that the transmitting node sends no packets unless a corresponding ticket is first received.

Because of the size of the receive FIFO, a node can have only three normal and six low-bandwidth tickets outstanding at any one time. If all tickets have been issued, the node will queue incoming vouchers to be handled later (up to a limit of 30 normal and 10 low-bandwidth vouchers). In the unlikely event that the destination node can neither issue a ticket nor queue the voucher, it will reject the voucher, which will be returned to the source node. The source node will wait for 100 clocks before attempting the same transaction again.

The source node may launch more vouchers before the first ticket arrives; up to three vouchers may be outstanding from the same target node. When a node receives vouchers or tickets not intended for it, it gives them priority over whatever data packets it has waiting to be transmitted. These vouchers and tickets "hop over" data packets in the forwarding block so they can be put onto the ring and sent to the proper node as quickly as possible.

According to figures from National, best-case latency for a five-node ring with light traffic is approximately 1 µs, from the time a client writes the head word to the transmit port at the source node, to the time the head appears at the receive port of the destination node. Part of this delay is due to the fact that tickets and vouchers encounter a typical delay of 40 ns at each node, and that between them, the voucher/ticket pair must circulate through all nodes in a ring. Calculating latency for a ring can be complex because it involves a host of variables, including the number of nodes and amount of traffic on the ring.

Favorable Price/Performance

Perhaps the most compelling reason to use QuickRing is the performance it brings for moderate cost. Although its high-volume price of $72 for two nodes is more than some popular bus schemes, QuickRing is still one or two orders of magnitude less expensive than fast serial links, such as FDDI or SCI, which cost $500–$1,000 or more per node. Table 1 shows that QuickRing can deliver 2 to 20 times the bandwidth-per-dollar of a typical bus interconnect, assuming the application is favorable to pipelined data streams. If an application cannot take advantage of concurrent data streams, QuickRing's latency issues may put it out of the running.

National Pushes Interoperability

Although QuickRing seems to have disappeared from Apple's product plans, National has carried through with the project, completing the design, expanding upon it, and evangelizing it to third parties. Working with developers, National has put together protocols and specifications to employ QuickRing as a mezzanine intereconnect and as a way to connect systems over short and long ranges. At Comdex, National demonstrated a full-speed QuickRing connection which passes stored, uncompressed video between two PCs over a three-meter cable. Other developments include a method for connecting multiple rings.

National has fostered a growing group of developers, many of which are close to releasing products based on QuickRing technology. For example, EnerAnalytics (San Ramon, Calif.) uses QuickRing to connect its JPEG compression board to its Advanced Graphics Engine, a 24-bit color display card with an embedded RISC processor for graphics acceleration. Xedia (Wilmington, Mass.) uses QuickRing in its high-performance networking hub, and TeleGlobe Communications (North Andover, Mass.) employs QuickRing in its CX-3000 broadband multimedia hub.

National sponsors a developer's program that not only helps developers with designs, but acts as a clearinghouse for user-generated comments to refine QuickRing as a product and as a standard. The group regularly distributes mailings with information about applications and software, runs developer forums and events, and coordinates media and marketing issues.

Designers can also purchase a developer's kit that includes reference hardware designs, reference software, driver interfaces, prototyping cards with connectors, and access to National's applications engineers and QuickRing bulletin-board system. Two versions of the kit are available; the one for PCs supports the VL-Bus and ISA buses, while the one for Macintoshes supports the NuBus and Quadra PDS buses. National plans to support QuickRing across multiple platforms, including a variety of workstations. Additionally, it will help define multivendor interoperability specifications to make QuickRing a "plug-and-play" standard.

An Uphill Battle

QuickRing has made inroads into the market, but considering the number of new buses and serial links being developed, National faces an uphill battle just to be noticed. PCI looks like it may be the next standard bus for personal computers and workstations, with system-logic support from the likes of DEC, Intel, Motorola, and IBM. Apple has announced that future PowerPC machines will use PCI, and appears to have dropped QuickRing as an internal project, although numerous Macintosh developers continue to show interest. PCI and VL-Bus both have extended 64-bit standards that could be more sensible in some applications.

Apple has teamed with IBM to promote the P1394 standard, which, among other things, is proposed as a serial link for connecting SCSI devices. IBM is also pushing another fast serial link called SSA, or serial storage architecture, for connecting systems to disk arrays.

Considered in a system, none of these give datastream performance on par with QuickRing. Perhaps more importantly, National can move more quickly than a standards committee, so QuickRing may suffer fewer delays getting into the market. The QuickRing concept is a good one and the timing is right, so expect to see more rings in your future. ♦

