

G-Link: A Chipset for Gigabit-Rate Data Communication

Two easy-to-use IC chips convert parallel data for transmission over high-speed serial links. A special encoding algorithm ensures dc balance in the transmitted data stream. A binary-quantized phase-locked loop is used for clock recovery. An on-chip state machine manages link startup automatically.

by **Chu-Sun Yen, Richard C. Walker, Patrick T. Petruno, Cheryl Stout, Benny W.H. Lai, and William J. McFarland**

The last decade has seen a tremendous increase in computing power with only modest advances in the bandwidth of the data links used to interconnect these computers. Between 1982 and 1992, the speed of a high-performance engineering workstation has increased from 0.5 MIPS (million instructions per second) to 100 MIPS, an increase of over two orders of magnitude. In that same period of time, computer network bandwidths have gone from Ethernet at 10 Mbits/s to FDDI at 100 Mbits/s, an increase of only one order of magnitude. In addition to faster computers, other factors, such as the widespread use of multimedia applications, will put pressure on network bandwidths, threatening to create an I/O bottleneck for modern computing systems.

Unlike computer systems, serial links cannot exploit parallelism and must run at proportionally higher rates for each increment in performance. At clock rates below about 100 MHz, traditional printed circuit board design techniques can be used to implement link circuitry with collections of packaged parts. But as link speeds approach the gigabit-per-second range, interchip timing skews make it impractical to build low-cost gigabit links in this way. Although long-haul telephone networks have used gigabit-rate data links for many years, these links use nonintegrable components and require adjustment and maintenance. Such systems are easily justified when the cost is amortized over millions of users but are too costly and complex for computer use.

To support the needs of computer and other generic data transport applications, the HP HDMP-1000 gigabit link (G-link) chipset has been developed. It is the first commercially available 1.4-Gbaud link interface in two chips, a transmitter chip and a receiver chip, requiring no external parts or adjustments.

The architecture of the G-link chipset greatly eases the job of the system designer. Communication between the chipset and the user's system takes place through a low-speed parallel interface. All gigabit-rate signals, with the exception of the serial electrical data stream, remain internal to the chips and are never routed on the printed circuit board. Thus the designer is able to use standard printed circuit board design techniques to deliver gigabit-rate performance. For fiber

optic applications, the high-speed serial signals are easily connected to lightwave transmitter and receiver modules. To simplify the designer's job further, a link-management state machine controller implemented on the receiver chip insulates the user from many of the details associated with link startup and error monitoring.

The chipset was designed in HP's 25-GHz f_T silicon bipolar process and incorporates patented circuit techniques developed at HP Laboratories, namely the encoding scheme and the phase-locked loop circuit. These new techniques, described later in this paper, represent departures from traditional telecommunication practice and have made practical the integration of an inexpensive and easy-to-use gigabit-rate chipset.

Overview

Fig. 1 shows a typical G-link application supporting a full-duplex interconnection between two hosts. One transmitter and one receiver chip are used for each end of the link.

From the user's viewpoint, the chipset behaves as a "virtual ribbon cable" for the transmission of parallel data over serial links. Parallel data is serialized by the transmitter chip and deserialized by the receiver chip into the original parallel form. The chipset hides from the user all the complexity of

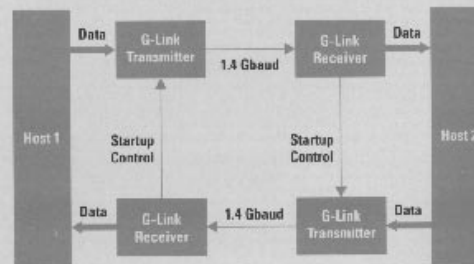


Fig. 1. A duplex link built with the HP HDMP-1000 gigabit link (G-link) chipset.




Fig. 2. Simplified transmitter chip block diagram.

encoding, multiplexing, clock extraction, demultiplexing, and decoding needed for high-speed serial data transmission.

The transmitter chip (Figs. 2 and 3) accepts the user's parallel data word and clock. The word-rate clock is internally multiplied up to the serial rate in the transmitter chip phase-locked loop. This high-speed serial clock is used to multiplex the encoded data. The encoding algorithm, called *conditional inversion with master transition*, or CIMT,¹ creates a frame* for data transmission by appending four coding bits to each input data word. The resulting frame is then transmitted in either normal or inverted form,² as necessary, to maintain dc balance of the serial bit stream for transmission over optical links or coaxial cables. This CIMT line code distinguishes itself by being efficient and simple to implement compared to other line codes such as 8B/10B.

To support modern network protocols, the chipset allows the transmission of three different types of frames. Generic user data is transmitted with *data frames*. *Control frames* are the second type of frame, and are used for the transmission of information that should be treated separately from data, such as packet headers. *Fill frames* are the third type of frame, and are sent automatically by the link during startup and to maintain synchronization when the user has neither data nor control information to send.

In the receiver chip (Figs. 4 and 5), the clock and frame alignment are extracted from the incoming data stream with a phase-locked loop. The data is then demultiplexed and decoded back to its original parallel form. In addition to these basic functions, the receiver chip also includes a state machine controller, which performs an end-to-end handshake and provides both bit and frame synchronization. This handshake avoids the false lock problems that are typical with clock extraction circuits that accommodate a wide range of clock frequencies.

An unconventional "bang-bang" phase-locked loop³ is used in the transmitter and receiver to provide adjustment-free bit retiming at very high data rates. Using the special master transition built into the line code, the phase-locked loop provides frame synchronization without the periodic insertion of special frame synchronization words.

A very compact chip layout was achieved by using three layers of metal and a quasi-gate-array ECL design methodology.

* In this paper, a frame is defined as an encoded input word.




Fig. 3. Photomicrograph of the transmitter chip.

The 68-pin surface-mount package (Fig. 6) is designed to maintain good performance for 1.4-GHz signals.

The key features of the chipset are:

- Parallel ECL bus interface
- 16 or 20 bits wide, pin selectable
- Flag bit usable as extra data bit (17th or 21st)
- CIMT encoding and decoding
- Ac/dc coupled
- 110 to 1400 Mbaud serial line rate
- On-chip phase-locked loops for transmitter clock generation and receiver clock extraction
- Local loopback mode for troubleshooting
- Single -5V ±10% supply voltage
- 2W power dissipation per chip (typical)
- Can be used with fiber optic links
- On-chip equalizer for use with coaxial cable
- Standard 68-pin CQFP (ceramic quad flat package).

Because of the simplicity and flexibility of the G-link chipset, it can be used for a wide variety of applications, including computer backplanes, video distribution, peripheral channels, and networks.




Fig. 4. Simplified receiver chip block diagram.




Fig. 5. Photomicrograph of the receiver chip.

G-Link Line Code

Many coding schemes have been developed to allow communication of information over various types of channels. In synchronous communication links, clock and framing information must be transmitted along with data in such a way that the clock and data can be recovered at the receiving end of the link. Therefore, it is necessary for the transmitted encoded serial bit stream to have enough embedded clock information for the receiver to recover the serial clock. There must also be some method of frame alignment so that the boundaries of a frame can be located at the receiver.

In optical links, it is desirable to ac couple the data signals to simplify laser bias circuitry and optical receiver design. This is also true in repeater design, since the components are commonly ac coupled. A problem with ac coupled systems is that the baseline will shift when the transmitted digital data is not dc balanced. This shift makes detection difficult and degrades the system noise margin. To overcome this problem, arbitrary data is typically encoded before transmission to achieve dc balance. The receiver restores the data to its original form by decoding.

In the G-link chipset, the CIMT coding scheme performs the following tasks:

- The transmitter chip supplies a master transition in every frame for clock recovery and frame alignment at the receiver.
- Frames are conditionally inverted as necessary to maintain dc balance.
- Information is provided in the transmitted frame about the type of frame transmitted and whether or not the frame was inverted.
- At the receiver, decoding is done to determine what type of frame was received and whether or not the frame was inverted.

- If the frame was inverted at the transmitter, it is inverted again at the receiver to restore the information to its original form.
- The receiver performs error checking on portions of the frames to detect loss of lock.

This method of encoding and decoding has several advantages:

- Clock information is available in each frame, indicating both phase and frequency alignment.
- There is no need for the user to send any special characters to indicate the start of a new frame. The G-link chips perform frame alignment transparently.
- There are no restrictions on the user's input bit patterns. Dc balance is maintained by frame inversion and a maximum run length is guaranteed by the master transition.
- By checking for framing errors, the receiver can detect loss of lock and reinitiate the link startup process. (A discussion of link startup can be found under "Startup State Machine Controller" on page 109.)

Data is encoded by appending four extra coding bits (C-field) to the input data (D-field). The serial combination of the D-field and the C-field makes a frame. The user can choose to transmit either data frames or control frames. In addition, two types of fill frames are internally generated for transmission when there is no input supplied by the user or during startup. To maintain dc balance, data and control frames are either inverted or not inverted. Information about inversion and the type of frame is contained in the C-field. Unlike typical codes with fixed data width, the CIMT code can accommodate multiple data widths.

The G-link chipset is designed to transmit either 16-bit-wide or 20-bit-wide data words. Both the transmitter chip and the receiver chip have an input pin that allows the user to select




Fig. 6. Transmitter chip in 68-pin ceramic quad flat package (CQFP).

the parallel word width. There is also a flag bit, which can be used as an extra data bit. A frame consisting of the D-field plus the appended C-field is then either 20 or 24 bits long. In the case of control frames, two bits in the D-field are used for encoding, resulting in 14 or 18 bits available for transmitting information. The flag bit is obtained by selecting between different sets of coding bits in the C-field.

Table I shows the contents of different frames generated at the G-link transmitter for the case of 20-bit data. DAV (data input available) and CAV (control input available) are supplied by the user to indicate what type of user input is to be transmitted. If neither data nor control inputs are available, a fill frame is sent. FLAG is the additional flag bit input. D0 to D19 are the parallel inputs. INV is a logic signal internally generated on the transmitter chip that indicates whether the frame is to be inverted.

Table I
Contents of Different Frame Types
for CIMT Encoding of 20-Bit Data

FLAG DAV CAV INV	D-Field	C-Field	Frame Type
		MT ↓	
X 0 0 X	11111111 11 00000000	00 11	Fill (FF0)
X 0 0 X	11111111 11 00000000	00 11	Fill (FF1L)
X 0 0 X	11111111 11 00000000	00 11	Fill (FF1H)
X X 1 0	D0-D8 01 D9-D17	00 11	Control
X X 1 1	D0-D8 10 D9-D17	11 00	Inverted Control
0 1 0 0	D0-D19	11 01	Data, FLAG Low
0 1 0 1	D0-D19	00 10	Inverted Data, FLAG Low
1 1 0 0	D0-D19	10 11	Data, FLAG High
1 1 0 1	D0-D19	01 00	Inverted Data, FLAG High

X = Don't Care

MT = Master Transition

The C-field bits were chosen so that a master transition always occurs between the second and third bits of the C-field. For data and control frames, this transition can be in either direction. The C-field bits were also chosen so that the codes

for data and inverted data frames are complements of each other. The same is true for control frames. This allows the entire frame to be inverted with the correct C-field bits for a particular type of frame.

There are two types of fill frames, referred to in Table I as FF0 and FF1. FF0, a training sequence used during startup, has a single rising edge at the master transition and is a square wave with 50% duty cycle. The receiver's clock recovery circuit is able to lock onto this signal, extract the serial clock, and provide frame alignment. FF1, another training sequence used during startup, is also sent after startup whenever the user does not supply inputs for data or control frames. FF1 is similar to FF0 except that the position of the falling edge moves by one bit forward or backward, creating a square wave that is two bits heavy (FF1H) or two bits light (FF1L). The decision to send either FF1H or FF1L is made depending on the disparity* of previously transmitted bits, in an attempt to reduce the disparity to zero. Since FF0 is dc balanced and the two types of FF1 frames are sent to reduce disparity, fill frames are not inverted.

Noninverted control frames have the same C-field as fill frames, but are distinguished from fill frames by the center two bits of the D-field, which are 01. Control frames are inverted when appropriate, but then have a different, unique C-field.

All other possible C-field codes that are not listed in Table I are not allowed and are considered to be errors if received. The receiver detects the loss of a master transition or a forbidden C-field code as a frame error. This information is used by the receiver's state machine to derive the link status. In addition, if the flag bit is not used by the user, it is used for additional frame error checking. The flag bit is alternated internally by the transmitter and this alternation is checked at the receiver.

Coding Implementation

Fig. 7 shows a block diagram of the transmitter chip. The user supplies the parallel inputs D0-D19, a frame rate clock, the DAV and CAV inputs, and the FLAG input (optional). The high-speed and subrate clocks are derived from the frame rate clock by a phase-locked loop circuit. "System I/O"

* Disparity is the number of 1s minus the number of 0s.

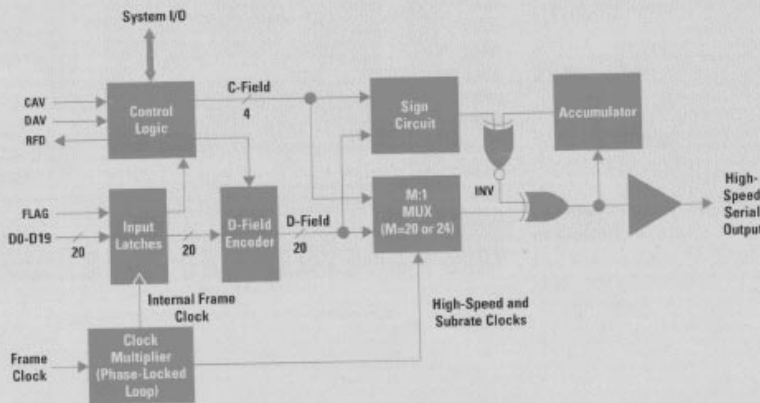


Fig. 7. Transmitter encoding circuitry.

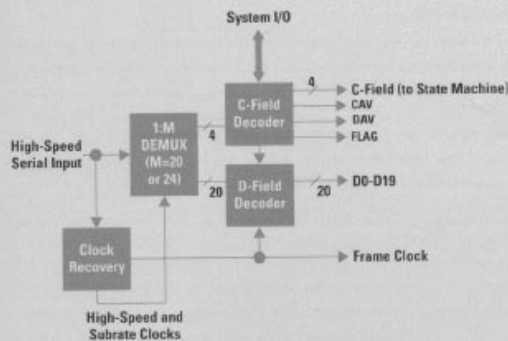


Fig. 8. Receiver decoding circuitry.

refers to other signals that are involved in the link's configuration and status. RFD (ready for data) is an output indicating to the user that the link is ready to transmit data. The D16-D19 inputs are ignored when the user selects 16-bit parallel word width.

Depending on the DAV, CAV, and FLAG inputs, the C-field coding bits are generated and any necessary encoding of the D-field is performed. Then the C-field and D-field bits are evaluated in a sign circuit whose output is the sign of the disparity of the frame. A separate accumulator keeps track of the disparity of previously transmitted bits. The decision to invert or not to invert a frame is made based on the outputs of these two circuits and is indicated by the signal INV. If the signs of the disparities of the current frame and the previously transmitted bits are the same, INV is high and the current frame is inverted. If they are not the same, INV is low and the frame is not inverted. Only data and control frames are inverted; the invert function is disabled for fill frames. The frame is serialized with a circuit that multiplexes the parallel inputs into a serial bit stream and performs any necessary frame inversion. The output of this circuit is then transmitted across the serial link.

A block diagram of the decoding portion of the receiver chip is shown in Fig. 8. After startup, the serial clock and the framing information are produced by the receiver's clock recovery circuitry, allowing the receiver to recover the serial data and demultiplex it back to parallel form. The frame clock is provided as an output for use in the user's system.

By examining the C-field bits, the C-field decoder determines what kind of frame has been received and whether or not it has been inverted. With this information, the D-field decoder restores the parallel data back to its original form. In addition, the C-field decoder provides DAV, CAV, and FLAG information back to the user. These signals have the same definitions as the corresponding transmitter inputs. The C-field bits are also used by the receiver's state machine to check for frame errors.

Encoding Circuitry

Encoding on the transmitter chip is performed mainly by logic cells and two on-chip programmable logic arrays (PLAs). However, there are two special parts of the frame inversion function. The first is an analog sign circuit which

determines whether a frame has more high or low bits. The second is an accumulator which keeps track of the disparity of the previously transmitted data.

The sign circuit on the transmitter consists of one differential pair per bit, a summing circuit, and a comparator. To prevent errors in determining a frame's sign, it is important for the differential pairs to have matched current sources. Therefore, each differential pair is supplied by two current sources from an array of current sources laid out in common centroid fashion. This reduces the effects of process and temperature gradients on the value of each pair's combined current source. In addition, large-geometry resistors are used to improve matching of the current sources.

The currents are summed at shared collectors through resistors, creating a differential voltage proportional to the difference between the numbers of 1s and 0s in the frame. When there are more 1s than 0s, this voltage is positive; when there are more 0s than 1s, it is negative. This voltage then drives a comparator, which produces a high or low logic signal depending on the sign of the input voltage. This method of determining the sign of a frame is simpler and faster than a digital solution.

The accumulator circuit keeps track of the disparity of previously transmitted bits. It is implemented with a 6-bit up/down counter. To relieve timing constraints, the counter operates on two bits at a time. This allows it to operate at a clock rate that is half the serial output rate.

The counter can count from all 0s to all 1s and is reset at startup to the midpoint, which is considered a balanced state. The range of this 6-bit counter is then -32 to $+31$ bits, where 0 is the balanced state. With two input bits, there are four possible combinations: 11 which has a disparity of $+2$ bits, 00 which has a disparity of -2 bits and 01 or 10 which are balanced with zero disparity. Since we only need to count up or down by multiples of 2, we can allow one bit of the counter range to correspond to a disparity of 2 bits. Thus the effective counter range, in bits of disparity, becomes -64 bits to $+62$ bits. The worst-case disparity that can occur with this coding scheme is ± 31 bits, which is well within the range of the counter. The most-significant bit of the counter is compared with the output of the sign circuit to decide whether to invert the frame.

Accumulating two bits at a time is the most convenient approach. If the counter were to operate on one bit at a time, it would still have to count either up or down and one bit of the counter range would correspond to one bit of disparity. Thus, the range of a 6-bit counter would be -32 to $+31$ bits of disparity, which would not have enough margin beyond the worst-case disparity of ± 31 bits. A higher-order counter would be required, and it would also have to run at the full serial output rate, resulting in increased power consumption.

If the counter were to operate on four bits at a time, it would have the benefit of running at one fourth of the serial rate, but it would have to count up and down by 4, up and down by 2, or remain unchanged. One bit of the counter range could correspond to two bits of disparity as in the case implemented, but the counter design would be more complex.




Fig. 9. Typical clock extraction and data retiming circuit requires phase adjustment and wide bandwidth.

Phase-Locked Loop

In a serial data link, the clock signal is not explicitly transmitted, but is instead implied by the transitions of the data stream. By examining the transitions in the data stream with a clock extraction circuit it is possible to create a replica of the original clock that was used to transmit the data. This recovered clock can then be used to sample and restore the potentially degraded analog input.

Many high-speed clock extraction techniques exist, but most have been developed for long-haul telephone applications. Telecom systems are designed to maximize the distance-bandwidth product of the link. This criterion minimizes both the number of physical repeater sites and the number of fibers that have to be installed in a given run. As a result, a much higher premium is placed on clock-extraction performance than on cost-effectiveness. These objectives have made this class of clock extraction techniques unsuitable for datacomm applications.

Traditional Telecom Clock Extraction Circuits

Fig. 9 shows a representative clock extraction and data retiming circuit that is used for high-bit-rate telecom systems. The incoming analog data stream is split into two parallel paths: the clock extraction chain and the data retiming path.

Because an NRZ (nonreturn to zero) data stream does not have a spectral component at the clock frequency, some nonlinear process must be used to derive a clock signal from the data stream. In the typical circuit of Fig. 9, a time derivative is applied, followed by an absolute value function. This combination of elements creates a narrow unidirectional pulse for every transition of the data. This new waveform contains a spectral component at the clock frequency. Once the clock component has been created, it can be isolated either by a filter, typically implemented with a SAW (surface acoustic wave) device, or by a phase-locked loop.

There are two problems with this configuration. The first is that, although the circuit extracts the correct clock frequency, it does not extract the correct phase. There is a large phase shift between the input data and the recovered clock. The phase relationship between the clock and the data must then be adjusted somehow to compensate for process and temperature variations. The second problem is that the creation of narrow pulses requires high circuit bandwidth. This is often the speed-limiting factor for gigabit-rate clock recovery circuits.

G-Link Solution

A design goal of the G-link chipset was to eliminate all external parts and user adjustments and effectively hide the system complexity from the user through monolithic integration. The clock extraction circuit was most impacted by these requirements. To achieve these aggressive goals, a new phase-locked loop circuit was developed based on a binary-quantized ("bang-bang") phase detector.

The phase-locked loop circuit used in the G-link chipset (see Fig. 10) works hand in hand with the CIMT line code to avoid both the phase adjustment problem and the bandwidth requirement of the traditional techniques. In this circuit, the incoming data splits into two paths (just as in the traditional telecom approach). Instead of a complex phase detector, which is potentially mismatched in delay to the retiming latch, two matched latches are used at the front end of the circuit. One latch is used for retiming and the other for phase detection. Because both latches are laid out identically on the chip, their delays are well-matched.

The two latches are driven by the VCO through a complementary buffer. If the VCO is properly aligned, the top latch samples the center of the data cell on rising edges of the clock while the lower latch samples the data transitions on the falling edge of the clock.

Because the G-link line code provides a guaranteed transition at a fixed, defined location in every frame, the sample of this transition can be used as an indication of the loop phase error. The VCO output is divided by either 20 or 24, depending on the selected word width, to produce one sampling pulse per frame. That clock pulse is used to take a sample in the vicinity of the master transition so that a phase update is generated, once per frame, indicating whether the VCO is early or late with respect to the master transition. Assuming a rising master transition, as shown in Fig. 11, if the VCO is too high in frequency, the sampling point drifts to the left of the master transition and a low value is sampled. If the VCO is too low, the sampling point moves to the right and a high value is sampled. This circuit then produces a one or zero indication from the phase detector that tells whether the VCO is early or late with respect to the incoming data.

Since the fastest operating element in this circuit is a latch operating at the serial rate, this circuit is usable up to the




Fig. 10. Simplified diagram of the G-link binary-quantized (bang-bang) phase-locked loop and data retiming circuit.

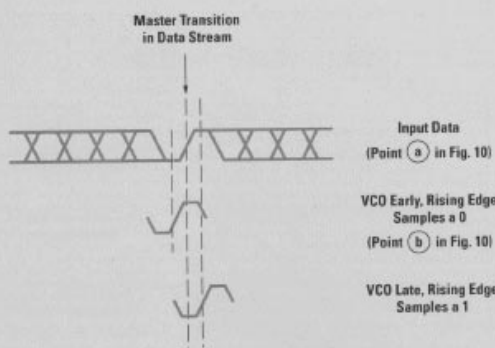


Fig. 11. Once per frame, the phase-locked loop detects whether the VCO is early or late with respect to the master transition encoded in each frame.

highest frequency at which a given process is capable of making a functioning latch. In addition, the circuit inherently provides excellent phase alignment between the VCO and the data. Note that the output of the phase detector latch is not linearly proportional to the loop phase error, but is instead a binary-quantized representation of the error. This characteristic renders the loop equations nonlinear and requires unconventional design methods (see "Bang-Bang Loop Analysis," page 110).

False Locking and Frame Synchronization

During initial link startup, it is necessary to ensure that the phase-locked loop correctly determines the frequency of the incoming data and finds the location of the master transition.

In many clock extraction circuits, the clock frequency is extracted from a coded, random data stream. A common difficulty with this approach is the problem of the phase-locked loop locking onto wrong frequencies that are harmonically related to the data rate. To avoid this problem, most systems limit the VCO range so that it can never be more than a few percent away from the correct frequency.

A narrow-band VCO using external components was not consistent with the goal of building a completely monolithic chipset. Integrated oscillators rely on low-tolerance IC components and are typically limited to $\pm 30\%$ tolerance on the center frequency. For customer flexibility, it was desired to extend the oscillator range to cover at least an octave. This range, in conjunction with digital dividers, allows the G-link chipset to operate over a range of 110 to 1400 Mbaud in four bands.

A second design problem is frame synchronization. At the receiver, some method must be employed to determine the boundaries between frames so that they can be properly deserialized back into the original parallel words. The G-link chipset establishes and monitors frame synchronization by using the embedded master transition. Unlike other links, the G-link chipset allows the continuous transmission of unbroken streams of data, without the insertion of special frame synchronization words.

Startup State Machine Controller

To eliminate the problems of false locking and frame synchronization, the G-link chipset uses a startup state machine and the special training fill frames.

Because the internal VCO is capable of operating over nearly a 3:1 range of frequencies, a frequency detector is necessary to avoid false locking problems. The frequency detector operates only when simple square-wave fill frames are being sent. A conventional sequential frequency detector, built of two resettable flip-flops, determines the sign of the frequency error. When the phase error is less than ± 22.5 degrees, the output of the phase detector is used. Otherwise, the loop filter is driven by the frequency detector output. Because the frequency detection circuit cannot operate on data frames, the state machine controller must disable the frequency detection circuit before allowing data to be sent.

Neither node of a duplex link can achieve lock unless the opposite side is sending special fill frames. Neither side of the link can stop sending fill frames and start sending data unless the other side has successfully achieved lock. The state machine uses the two distinct fill frames FF0 and FF1 to allow one side of the link to notify the other side of its current locking status. This guarantees that fill frames will be sent whenever needed to restore lock, and only as long as necessary to achieve lock.

As described previously, FF0 is a 50% balanced square wave with equal numbers of 0 and 1 bits. FF1 consists of two modified square-wave patterns. These two patterns are used as needed to maintain dc balance on the link. Both FF0 and FF1 have a single, rising transition, which is in the same position in the frame as the master transition of data and control frames. The rising edge of the fill frames is used initially to establish an unambiguous frame reference. After initial lock, the master transition of the data frames is used to maintain frame lock.

Fig. 12 shows the state machine handshake procedure for a full-duplex link in greater detail. Both the near and far ends of the link independently follow the state diagram of Fig. 12. The three states are defined by the state variables STAT0 and STAT1. At power-up, each end of the link enters the sequence at the arc marked "Start."




Fig. 12. State machine handshake procedure for a full-duplex link, showing the values of the state variables STAT0 and STAT1 (0,0, etc.).

Bang-Bang Loop Analysis

A simplified version of the clock recovery phase-locked loop of the G-link chipset is shown in Fig 1. Only the transition sampling latch is shown, and the input is assumed to be a square wave at the same frequency as the VCO.

The VCO is controlled through a loop filter that consists of the sum of an integral signal and a proportional signal. Because the phase detector is quantized, the VCO frequency switches between two discrete frequencies, causing the VCO to ramp up and down in phase, thereby tracking the incoming signal phase.

If the loop is properly designed, the system can be considered to be composed of two noninteracting loops. These are the paths labeled proportional branch and integral branch in Fig. 1. The first loop includes the connection of the phase detector to the VCO input through a proportional attenuator, while the second loop drives the VCO through an integrator.

The proportional signal tunes the VCO, causing the output of the phase detector to switch rapidly between 1s and 0s at a fairly high frequency. Other than the dc component, the bulk of the phase detector output signal spectrum falls outside the effective passband of the integrator branch of the loop. Thus the integrator branch operates on just the dc component of the phase detector output. Its job is to servo the center frequency of the VCO so that the two discrete VCO frequencies programmed by the proportional input will always bracket the frequency of the incoming data signal. This frequency adjustment occurs so slowly that it does not materially affect the operation of the high-frequency bang-bang portion of the loop.

Proportional Branch

To simplify the analysis of the first branch of the loop in Fig. 1, the integrator output can be replaced with a constant reference voltage so the proportional tuning input will cause the VCO to bracket the incoming frequency. The VCO will then run at two discrete frequencies: at a frequency slightly higher than the incoming data, thereby advancing the phase, or at a lower frequency, thereby retarding the phase.

If the incoming frequency is midway between these two discrete frequencies, the loop will switch between the two frequencies with approximately a 50% duty cycle. If the incoming frequency is slightly higher than the nominal VCO center frequency, the duty cycle will shift such that the loop will spend a higher percentage of time at the high frequency than at the low frequency. In general, it can be shown that the duty cycle present at the output of the phase detector is proportional to the difference in frequency between the incoming signal and the nominal VCO center frequency.

Integral Branch

The second branch of the loop contains the integrator. Because the integrator effectively filters out the oscillatory portion of the phase detector output and only reacts to the average value of the phase detector output stream, the proportional branch of the loop can be ignored here by replacing the phase detector with a virtual frequency detector. The integrator extracts the dc component and thereby




Fig. 1. Simplified version of the phase-locked loop. For analysis, the loop can be considered a combination of two noninteracting loops: a proportional branch and an integral branch.

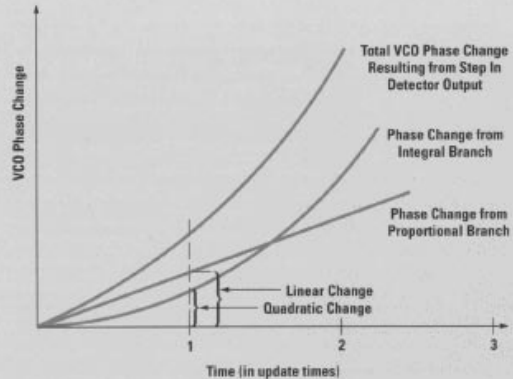


Fig. 2. Contributions to VCO phase changes. Stability factor is the linear phase change divided by the quadratic phase change in the same time.

tunes the center frequency of the VCO so that it is always equal to the incoming data rate.

In a conventional linear phase-locked loop, the loop error signal is proportional to phase error but is used to control the VCO frequency. This introduces an integration in the loop transfer function. This integration, in conjunction with the loop filter, creates a second-order feedback loop. Such loops can exhibit an underdamped response to changes in input phase, leading to an undesirable exponential buildup of jitter in systems with long cascades of repeaters.

In the G-link phase-locked loop, the phase-detector dc component is proportional to frequency rather than phase. Because the the frequency of the VCO is controlled by a frequency error signal rather than a phase error signal, no extra integration appears in the loop transfer function. This means that no jitter buildup results from the action of the integral branch of the loop. The jitter statistics are simply dominated by the hunting behavior of the high-frequency proportional branch of the loop.

Loop Stability

To reach a qualitative understanding of the loop behavior, the two branches of the loop were assumed to be noninteracting. For this assumption to be valid, certain conditions must be met.

It is important that the loop be set up so that, between phase samples, the action of the proportional branch of the loop dominates over the action of the integral branch. This can be verified by creating a step change from the phase detector and tracking its effect on both halves of the loop. Fig. 2 shows the contributions to the VCO phase change. In the proportional path, the VCO is programmed to make a small step change in frequency, which causes a linear ramp in the phase error. In the integral path, the integrator programs a linear ramp in VCO frequency, which causes a quadratic walk-off in the VCO phase.

The ratio of these effects at the end of one frame update time gives a figure of merit for the loop design. The phase change from the proportional branch of the loop must be greater than or equal to the phase change from the integral branch of the loop for the system to be stable. In the G-link design, this stability ratio is designed to be always greater than 10.

Richard C. Walker
Principal Project Engineer
Hewlett-Packard Laboratories




Fig. 13. The VCO consists of three variable-delay cells configured as a ring oscillator.

Each state in the state machine has three notations. The top notation is either "FDet" or "Phase." FDet stands for frequency detect mode, and implies that the frequency detector has been enabled in the receiver chip phase-locked loop. When the chip is in this mode, it is important that no data be sent, because the frequency detector is only able to lock onto one of the special training fill frames FF0 or FF1. The Phase notation means that the receiver phase-locked loop has been switched to phase-detect mode and is ready to allow data transmission. The middle notation in each state is the fill word that is currently being sent by the node's transmitter chip. The last notation is the ready-for-data (RFD) status of the transmitter chip. When RFD is low, the transmitter chip signals the user to hold off any incoming data while it is sending fill frames. When RFD is high, data is sent if available, and if not, fill frames are sent to maintain link synchronization.

The two bits bracketing the master transition are monitored by the receiver chip to detect a locked condition. If these two bits are not complementary for two or more consecutive frames, it is considered a frame error. The receiver chips at both ends of the link are able to detect data, control, FF0, and FF1 frames and frame errors. Transitions are made from each of the states based on the current status condition received by the receiver chip. Each of the arcs in Fig. 12 is labeled with the state that would cause a transition along that arc.

If either side of the full-duplex link detects a frame error, it notifies the other side by sending FF0. When either side receives FF0, it follows the state machine arcs and reinitiates the handshake process. The user is notified of this action by the deasserting of RFD.

This startup protocol ensures that no user data is sent until the link connectivity is fully established. The use of a handshake training sequence avoids the false lock problem inherent in phase-locked loop systems that attempt to lock onto random data with wide-range VCOs.

Loop Implementation

The VCO is built from three variable-delay cells configured as a ring oscillator (Fig. 13). The ring provides a wide-range tuning input and a small "bang-bang" tuning input. The wide-range input adjusts the delays of each stage from one gate delay to three gate delays, thus giving a 3:1 VCO frequency range. This wide range allows the final system to be specified with a 2:1 range over both process and temperature variations. The bang-bang tuning input programs a small change in the VCO frequency and is driven by the proportional branch of the loop filter.

The loop filter is implemented with a charge pump integrator and a 0.1- μ F external capacitor, which is housed within the package. The integrator is based on a unity-gain positive feedback technique (Fig. 14) which cancels out the droop in the integrator filter capacitor. The effective dc gain of this circuit approaches infinity as the feedback gain approaches unity. The unity-gain technique achieves high dc gain while avoiding the stability and noise sensitivity problems of on-chip high-gain operational amplifier designs.

G-Link Chipset Implementation

To achieve the best speed and power performance, the G-link chips were designed using the HP B25000 25-GHz f_T silicon bipolar process. This process allows mixed-mode designs ranging from dense low-power logic structures to high-performance analog cells. A three-layer metal system allows compact layouts, minimizing chip area and cost. This process features transistors with minimum pitch of 2.6 μ m. Only simple npn transistors and p+ and p- resistors were used in the design.

Building Block Design

The G-link chipset is a fully custom circuit using specially designed cells as building blocks. These include (1) logic cells consisting of gates, latches, and flip-flops, (2) PLAs for low-speed logic, and (3) I/O cells, which include all of the low-speed ECL and high-speed input and output drivers. A band-gap reference was also designed to stabilize chip performance with variations in temperature and power-supply voltage.

Logic Cells and Arrays. Since logic elements are used most widely in the G-link chipset, considerable effort went into optimizing their performance, power, and active area. A three-level tree structure was chosen to implement the logic functions. All signals are differential to improve noise margins and to reduce ground currents, which could disrupt the analog circuitry. The inputs and outputs of these gates and latches are fully level-compatible for ease of routing. Each functional cell has resistor options by which the speed can be traded off with power. In all, there are four power classes for each logic cell. An example of a master-slave flip-flop with a 2:1 input multiplexer is shown in Fig. 15. This circuit is designed to operate up to 2 Gbits/s at a junction temperature of 125°C with a fanout of 10.

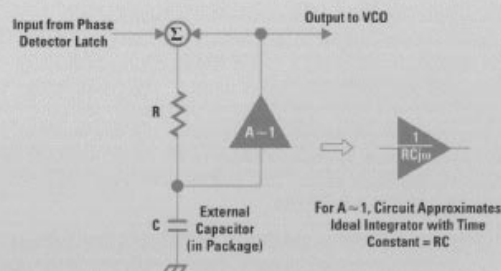


Fig. 14. The loop filter is implemented with a charge pump integrator based on a unity-gain feedback technique.

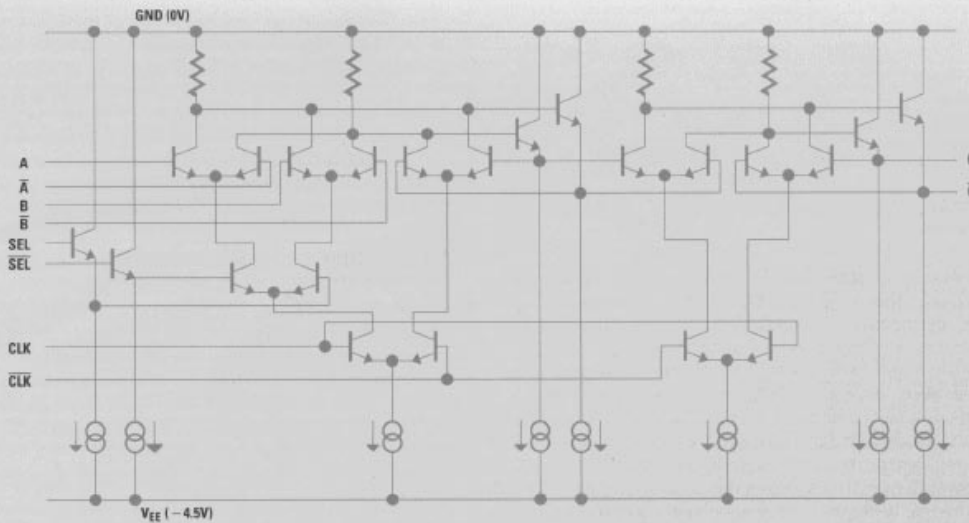


Fig. 15. Schematic diagram of the master-slave flip-flop with 2:1 input multiplexer.

Low-speed logic is implemented, where possible, with array structures for compactness and reduced power. The single-ended logic PLAs with AND-OR planes are designed to be programmed using only metal layers. Altogether, two PLA cells are used in the transmitter and one in the receiver.

Input/Output Cells. An effort was made in the I/O design to make the chips easy to use. Except for the high-speed serial signals, all of the chip I/O is 100K ECL-level-compatible. To minimize the power dissipation of the chip, ECL outputs are limited to driving 10 cm of transmission line with a minimum characteristic impedance of 50 ohms terminated into 300 ohms. For added convenience, unconnected inputs are internally biased to low ECL logic levels, and are sensed as high levels if the inputs are grounded.

A special input cell was designed for all gigabit-rate input signals. Both differential inputs of the cell are biased to ground with 50-ohm terminating resistors. This configuration allows singled-ended or differential input signals to be conveniently ac or dc coupled. This cell is used for the strobe and high-speed clock inputs of the transmitter and for the data and high-speed clock inputs of the receiver.

The G-link chipset is designed to work with either optical fiber or copper coaxial cable media. For cable applications, the data input cell of the receiver has an optional equalizer to extend the usable distance of the link. The equalizer circuit is designed with 3 dB of gain peaking at 600 MHz to compensate for signal roll-offs caused by the skin loss effect in coaxial copper cables. Operating at 1.2 Gbaud with RG-58 coax, the equalizer extends the usable cable length by over 50% for a given bit error rate.

All high-speed outputs are driven by buffered-line-logic cells. Buffered-line-logic drivers¹ provide differential outputs capable of delivering 0.7V into 50 ohms, ac or dc coupled to ground. If dc is coupled into -1.3V, the levels are ECL-compatible. In addition, the source impedance of the driver is matched to 50 ohms with a VSWR of less than 2:1. This

makes the high-speed connections of the G-link chips very convenient and easy to use. The only requirement is that unused outputs be terminated into 50 ohms.

Band-Gap Reference. To minimize circuit drifts caused by environmental changes, a band-gap reference with power supply compensation was designed. This circuit provides a reference voltage that powers up all cells in both chips. A power-down feature in this circuit enables portions of the chips to be turned off to conserve power.

Layouts

To minimize the design and layout effort, a generic design structure was used as the basis for all cell layouts. Each of the various logic cells was built from the generic array of transistors and resistors by customizing the metal interconnections. The ratio of devices used to total devices available reached over 95% in this design. This layout technique has the advantage of easy reconfiguration for design revisions. The I/O port locations are uniformly defined for all cells to simplify cell interconnection.

An example of a master-slave flip-flop with a 2:1 multiplexer input is shown in Fig. 16. This circuit array, measuring just 104 by 135 μm , is customized with two layers of metal.

All cells and power buses are designed to be placed using a coarse grid. This simplifies the placement of cells in the system design level. Another feature is that all cells have test probe points accessible at the top metal such that all connection signals can be test probed for diagnostic purposes.

The transmitter and receiver chips each measure 3.5 mm on a side. The high-speed and low-speed pads for each chip are arranged so that a single package design accommodates both chips.

The design of the chips relied heavily on simulation and verification tools such as the Spice simulation program and HP's proprietary Bipolar-Chipbuster IC layout system. The Spice




Fig. 16. Chip layout of the master-slave flip-flop with 2:1 input multiplexer.

circuit description files were extracted from the artwork including parasitic capacitors for final simulation before fabrication.

Packaging

A custom 68-pin ceramic quad flat package (CQFP) was designed specifically for the G-link chipset. It features 50-ohm transmission lines for the high-speed I/O pins and internal 0.1- μ F capacitors for power supply bypassing and for the integrator of the phase-locked loop. It also has internal ground vias to minimize inductance, thereby reducing noise. Its outline conforms to standard 68-pin packages. The typical chip-to-case thermal resistance is under 14°C/W. Both the package and the chips are compatible with automatic assembly techniques for high-volume low-cost manufacturing.

After the chips and capacitors are mounted and the packages sealed on the lead frame, the units are placed onto plastic carriers for lead protection. A special test fixture was designed to test the final parts in this carrier at full speed.

Electrical Performance

The G-link chips' power dissipations are both under 2.5 watts worst-case. The 20%-to-80% rise and fall times of the high-speed data outputs are under 200 ps. The chipset is specified from 110 Mbaud to 1.4 Gbaud under all conditions. The lockup time of the phase-locked loop including frequency acquisition is less than 2 ms.

Features and Applications

The features and flexibility of the G-link chipset make it ideal for a wide variety of applications. These applications range from computer backplane links a few meters in length to wide area networks 10 kilometers long. The low cost and high integration level of the G-link chipset make it attractive for systems requiring serial transfer rates up to 1.4 Gbaud. It can serve as a generic virtual ribbon cable or can be used to build complete networks and peripheral channels. The G-link coding scheme has been accepted by the Serial-HIPPI (High-Performance Parallel Interface) Implementors' Group, and by SCI-FI (Scalable Coherent Interface-Fiber), an IEEE standard.

This section describes the features that allow the G-link chipset to be applied to this broad range of applications. It also describes a few specific applications, including generic data transport, networking standards, and simplex applications.

Ease of Use

Since most computing equipment both sends and receives data, the great majority of these applications are full-duplex. The state machine controller included on the chipset takes care of all the details of starting up such a duplex link. The designer needs to be concerned with only two signals: ready for data (RFD) and data available (DAV). RFD is the signal the state machine provides to indicate that the link is ready for data transmission. DAV is a signal the user controls to mark the availability of data. At the receiver, this signal is recovered and used to discern the beginning or end of data transmission.

Some applications generate data in bursts or as packets. Such bursty data is handled automatically by the chipset. When no data is available to transmit, the user simply deasserts the DAV line at the transmitter. The link will transmit FF1 as an idle code to maintain link lock and framing. At the receiver, a deasserted DAV signal indicates that data is not being received. At the start of the burst of data, the user asserts the DAV line at the transmitter. The data is transmitted across the link and marked as valid data at the receiver by the receiver's DAV signal. Thus the DAV signals can mark the beginning and end of packets while adding no burden to the system design.

More complicated packet headers can be created using the control available (CAV) signal. This signal works like the DAV signal, but instead of marking the data as valid data words, it marks the data as special control words. A system designer can use these to send packet header information, link or system control information, or anything that needs to be treated separately from data. At least 2^{14} control words are available, so they can be used to indicate a large number of packet addresses or special functions. Few communication links have such a rich selection of nondata words for control and signaling.

Flexibility

Flexibility was a major goal of the G-link design. To make this a high-volume, low-cost part, the chips were designed to meet the needs of as many different systems as possible. As




Fig. 17. Serial-HIPPI (High-Performance Parallel Interface) system implemented with the G-link chipset.

described earlier, the G-link line code can accommodate various word widths. This is very different from block codes such as 4B/5B and 8B/10B, which have fixed word widths. The G-link chipset readily accommodates data words of width 16, 17, 20, 21, 32, or 40 bits. The chipset has two fundamental word sizes: 16 or 20 bits. In addition, the flag bit is available. Therefore, 17-bit-wide words can be accommodated by selecting 16-bit frames and using the flag bit as a 17th bit. 21-bit words can be transmitted similarly. 32-bit words are supported by sending them as two 16-bit frames in a row. In this case the flag bit is used to distinguish the first 16-bit frame (e.g., flag = 0) from the second 16-bit frame (e.g., flag = 1). It is a simple matter to build the off-chip 32:16 multiplexers and 16:32 demultiplexers since the flag bit automatically keeps track of the necessary frame ordering. 40-bit-wide words are supported analogously. The transmitter chip accepts either a full-frame-rate clock or a half-frame-rate clock for multiplication up to the serial clock rate. In other words, for an 800-Mbit/s data rate and 16-bit words, the chip will accept a 50-MHz frame clock. When 32-bit words are transmitted it accepts a 25-MHz frame clock. This saves the system designer the trouble of doubling the word clock outside the chip.

The G-link chipset supports a wide range of serial transfer rates ranging from 110 Mbaud all the way up to 1.4 Gbaud. This wide range makes it attractive for many types of data. Because the chipset requires no off-package tuned elements or adjustments, it can be digitally switched between data rates. This is unlike other systems, which require tuned elements and precise adjustments and operate over very narrow ranges of frequencies. Switching between data rates aids testing and debugging. It can also be used to establish a standard physical layer that spans several operating frequencies.

Generic Data Transport and Proprietary Channels

The most prevalent application of the G-link chipset is generic data transport. In these applications, the chipset acts as a point-to-point unswitched bus extender, or virtual ribbon cable. A great advantage of the G-link chipset is that it automatically handles startup and framing. Once the link is operating, the user can send data continuously, without having to insert extra framing characters or form special packets. Other links typically require that special framing characters be periodically inserted into the data stream. For systems transmitting data continuously for long intervals, periodically inserting these special characters can be difficult and inefficient. Other link chipsets do not have a built-in hardware controller that signals when the link is operating improperly. Without these signals the system designer must depend on upper-level protocols, resulting in uncertain time delays.

In many applications, a point-to-point unswitched bus extender is sufficient. In these applications, the G-link chipset

is all that is required and can form a complete communication link. The chipset can also be used in more complicated networks because it transports any data format across the link. Examples of standard data formats that can benefit from a point-to-point bus extender within private networks include SONET/SDH, Fiber Channel, and ATM data. SONET/SDH is a telecommunication standard that specifies data rates of 155 Mbits/s, 622 Mbits/s, 1.24 Gbits/s, and higher. Fiber Channel is an ANSI standard (X3T9.3) that covers a variety of data formats and rates. The IEEE 802.6 standard is an example of an ATM (asynchronous transfer mode) network.

The flexibility and ease of use of the G-link chipset enable it to fit a wide variety of applications. High-data-rate connections to disks and other peripherals are typical uses. These applications benefit from the very low overhead, simple operation, and high integration of the G-link chipset. For example the HP 27111A, introduced in 1988, is a fiber optic connection for disk arrays at 80 Mbits/s. With the tremendous increase in computing power and I/O rates in the last few years, the G-link chipset is well-suited for this type of application.

There is growing interest in using serial links for computer backplanes. Computer backplanes are typically jammed with hundreds of signals at data rates exceeding 100 Mwords/s. It can be difficult to control the skew on parallel data paths at high data rates. In addition, transmitting the data in parallel can require significant space. Serial links using optical fiber or coaxial cable may be the only way to transmit data without degradation by skew, loss, or reflections, while saving space.

Serial-HIPPI

In May 1991 the G-link chipset was accepted as the basis of the Serial-HIPPI standard. Serial-HIPPI is a specification for an 800-Mbit/s serial data link that has been agreed upon by over 40 vendors and users. Serial-HIPPI transmits data between HIPPI-PH nodes, up to 25 meters in coaxial cable, or 10 km with optical fiber. HIPPI-PH is an ANSI standard (X3.183-1991) for transmitting digital data in parallel between data processing equipment nodes. It is prevalent in supercomputing and high-end workstation environments. Fig. 17 shows a diagram of a complete Serial-HIPPI system using the G-link chipset. HIPPI-PH data consists of 44-bit-wide words at 25 Mwords/s. This data includes 32 data bits, 4 parity bits, 7 control bits, and the clock. Ahead of the G-link transmitter there is an additional circuit called the XMUX. This circuit reduces the data from 44 bits to 40 bits by replacing two control signals with the chipset's RFD, replacing the HIPPI-PH clock with the clock derived from the incoming serial data, and encoding three of the other control signals into two lines. The XMUX then multiplexes the data 40:20. This data is transmitted with the G-link chipset as 40-bit

125-Mbit/s FDDI data to 1.24-Gbit/s SONET data. The chipset can work in simplex systems, allowing its use for distributing video. Two widely accepted networking standards, Serial-HIPPI and SCI-FI, are tailored to the operation of the G-link chipset. The production volume made possible by this broad range of applications should make possible truly low-cost gigabit-rate data links.

Acknowledgments

We would like to thank Tom Hornak for his many contributions to this project, and his guidance throughout its development. Thanks to J.T. Wu for his contributions to the receiver chip design and layout. Kent Springer, Rasmus Nordby, Craig Corsetto, and Doug Crandall all had an early influence on the project. Hans Wiggers, David Cunningham, Steve Methley, David Sears, and Richard Dugan have been responsible for the G-link's success in the standards arena.

David Yoo, Jean Norman, Natalia McAfee, and Lie Lian-Mueller have all helped with the assembly and packaging of the chipset. Special thanks to Don Pettengill, Shang-Yi Chiang, and the HP B25000 process team, without whose excellent work and cooperation this project would not have succeeded.

References

1. D. Crandall, et al, *DC-Free Line Code for Arbitrary Data Transmission*, U.S. Patent no. 5,022,051, June 4, 1991.
2. R.O. Carter, "Low-Disparity Binary Coding System," *Electronic Letters*, Vol. 1, no. 3, May 1965, pp. 67-68.
3. C. Corsetto, et al, *Phase-Locked Loop for Clock Extraction in Gigabit Rate Data Communication Links*, U.S. Patent no. 4,926,447, May 15, 1990.
4. B. Lai, "A 3.5 Gb/s Fully Retimed Decision Circuit with Temperature Compensation," *Proceedings of the 1988 Design Technology Conference*, Hewlett-Packard, May 1988, pp. 296-303.


HEWLETT-PACKARD
JOURNAL

October 1992 Volume 43 • Number 5

Technical Information from the Laboratories of
Hewlett-Packard Company

Hewlett-Packard Company, P.O. Box 51027
Palo Alto, California, 94303-0724 U.S.A.

Yokogawa-Hewlett-Packard Ltd., Suginami-Ku Tokyo 168 Japan



89 Digital Network Monitor

Giovanni Nieddu



Gianni Nieddu is a system engineer at HP's Necsyste Telecommunications Operation. His professional interests include protocols and real-time embedded systems. He studied electronic engineering at the University of Padova (Padua), graduating in 1984. He joined Necsyste in 1987 and was responsible for the system architecture and software of the HP E3560 digital performance monitoring and remote test system. Born in Sacile, Italy, he served in the Italian Alpine troops for a year in 1985. He is married and enjoys mountain climbing and trekking, skiing, and games in general.

Fernando M. Secco



A system engineer with HP's Necsyste Telecommunications Operation, Fernando Secco studied electronic engineering at the University of Padova (Padua), graduating in 1982. He served a year in the Italian infantry in 1983 and joined Necsyste in 1989 after several years with Telettra SpA. He was responsible for the design of the peripheral units for the HP E3560 digital performance monitoring and remote test system. Fernando comes from Piazzola sul Brenta, Padova. He is married and has one child.

Alberto Vallerini



Alberto Vallerini is hardware system manager for the HP E3560 digital performance monitoring and remote test system. Previously, he was hardware system manager for the HP 3788A error performance analyzer. He joined HP's Necsyste Telecommunications Operation in 1988 and has a degree in electronic engineering from the University of Padova. Alberto was born in Lendinara, Rovigo. He did his military service as an official in the transmission service. He is married and enjoys reading, music, and theater.

103 G-Link Chipset

Chu-Sun Yen



Chu Yen was project manager for the G-link chipset at HP Laboratories. With HP Laboratories since 1961, he previously managed high-speed analog ICs, Ethernet transceiver and cable simulation, and infrared network projects, and was a design engineer on projects dealing with the HP 8405A vector voltmeter and the HP 35 calculator. A member of the IEEE, he has authored 20 professional papers on

circuits, phase-locked loops, instruments, and communications links, and is named as an inventor in three patents on a sampling phase-locked loop, the G-link phase-locked loop, and a dc-to-dc converter. He received his BSEE degree in 1955 from Taiwan University, his MSEE degree in 1958 from the University of Florida, and his PhD degree in 1961 from Stanford University. He is married and has three children.

Richard C. Walker



Rick Walker is a principal project engineer at HP Laboratories, specializing in phase-locked loop theory and high-speed circuit design. He joined HP Laboratories in 1981 and has contributed to broadband cable modem design, solid-state laser characterization, and the gigabit-link project. Born in San Rafael, California, he received his BS degree in engineering and applied science from the California Institute of Technology in 1982 and his MS degree in computer science from California State University at Chico in 1992. He has authored 16 professional papers and his work has resulted in six patents and two pending patents, all in the areas of high-speed links and circuit design. He's a member of the IEEE. He's also an advanced class amateur radio operator (W9GGV) and a private pilot, plays bass guitar and five-string bluegrass banjo, and cultivates several dozen kinds of carnivorous plants in a backyard greenhouse.

Patrick T. Petruno



Now R&D section manager for link products at HP's Communications Components Division, Pat Petruno was project manager for the G-link chipset. A native of Allentown, Pennsylvania, he attended Pennsylvania State University, receiving his BSEE degree in 1976 and his MSEE degree in 1978. After joining HP in 1978, he designed bipolar transistors, Darlington amplifiers, and digital circuits, and served as project manager for wideband amplifiers, AGCs, decision circuits, counters, and multiplexers. He has coauthored three papers on high-speed silicon circuits and participates in Serial-HIPPI and ATM (asynchronous transfer mode) standards activities. Pat is married, has two children, and serves as a Cub Scout den leader. His interests include astronomy, astrophotography, swimming, and softball.

Cheryl Stout



Cheryl Stout has been a member of the technical staff of HP Laboratories since 1983. She has done research and design for gallium arsenide and silicon high-speed multiplexers, optical receivers, and the gigabit-link chipset. She is a member of the IEEE and has authored conference papers on high-speed multiplexers and the G-link chipset. Born in San Jose, California, she received her BSEE degree from California State University at

San Jose in 1979 and her MSEE degree from the University of California at Berkeley in 1983. Before coming to HP, she developed optical communication products at Plantronics, Inc. Her interests include mountaineering and natural history.

Benny W.H. Lai



Engineer/scientist Benny Lai joined the HP Microwave Semiconductor Division (now the Communications Components Division) in 1981. He designed the HP HDMP-2003/4 decision circuits and the HDMP-2501 clock recovery data retiming circuit, and did circuit design and layout and high-speed testing for the G-link chipset. He has authored four papers on his designs and his work has resulted in one retiming circuit patent and two pending patents on CIMT coding and a unity-gain positive-feedback integrator. A graduate of the University of California at Berkeley, he received his BSEE degree in 1982 and his MSEE degree in 1983. He was born in Hong Kong, is married, and enjoys woodworking, gardening, and skiing.

William J. McFarland



HP Laboratories principal project engineer Bill McFarland received his BSEE degree from Stanford University in 1983 and his MSEE degree from the University of California at Berkeley in 1985. With HP since 1985, he has designed high-speed ICs for digital test instruments in silicon bipolar, gallium arsenide, and heterojunction transistor technologies, and has done research on bit error rate testers and pulse generators. As a member of the G-link project, he served as technical editor of the Serial-HIPPI specification. He is the author or coauthor of a dozen technical papers and is named as an inventor in two patents related to bit error rate testers. Bill was born in Milwaukee, Wisconsin. His interests include bicycling, playing guitar, and home brewing.