

A 2.488-Gbit/s Silicon Bipolar Clock and Data Recovery Circuit for SONET Fiber-Optic Communications Networks

Richard Walker

Cheryl Stout

Chu-Sun Yen

Lewis R. Dove

Adjustment-free clock and data recovery for 2.488-Gbit/s SONET applications is provided by a 1.77W, 3.45×3.45 -mm² chip implemented in a 25-GHz f_T silicon bipolar process. The chip has an on-chip VCO and operates from 2 to 3 Gbits/s over process, voltage, and temperature variations with a single off-chip filter capacitor. For network monitoring, a highly reliable loss-of-signal detector is provided. For good mechanical, thermal, and RF performance, a custom package was developed using HP's fine-line hybrid process.

SONET 2.488-Gbit/s transmission and switching systems, network backbones, and video transmission are among the many applications benefiting from inexpensive and robust clock and data recovery circuits. Clock and data recovery circuits are used in high-speed communications systems, typically long-span, single-mode fiber-optic links. Their job is to regenerate clean clock and data signals from arbitrary scrambled data inputs that have been corrupted by jitter and intersymbol interference.

To provide highly reliable clock and data recovery for 2.488-Gbit/s SONET data transmission, the HP CDR 2500 clock and data recovery circuit has been developed. Previous commercial solutions for this application have required multiple chips and GaAs processes.¹ The CDR 2500 was designed in HP's 25-GHz f_T silicon bipolar process,² and incorporates several new circuit ideas developed at HP Laboratories, namely an arbitrary-data phase detector, a reliable loss-of-lock detector, and a monolithic wide-range VCO circuit.

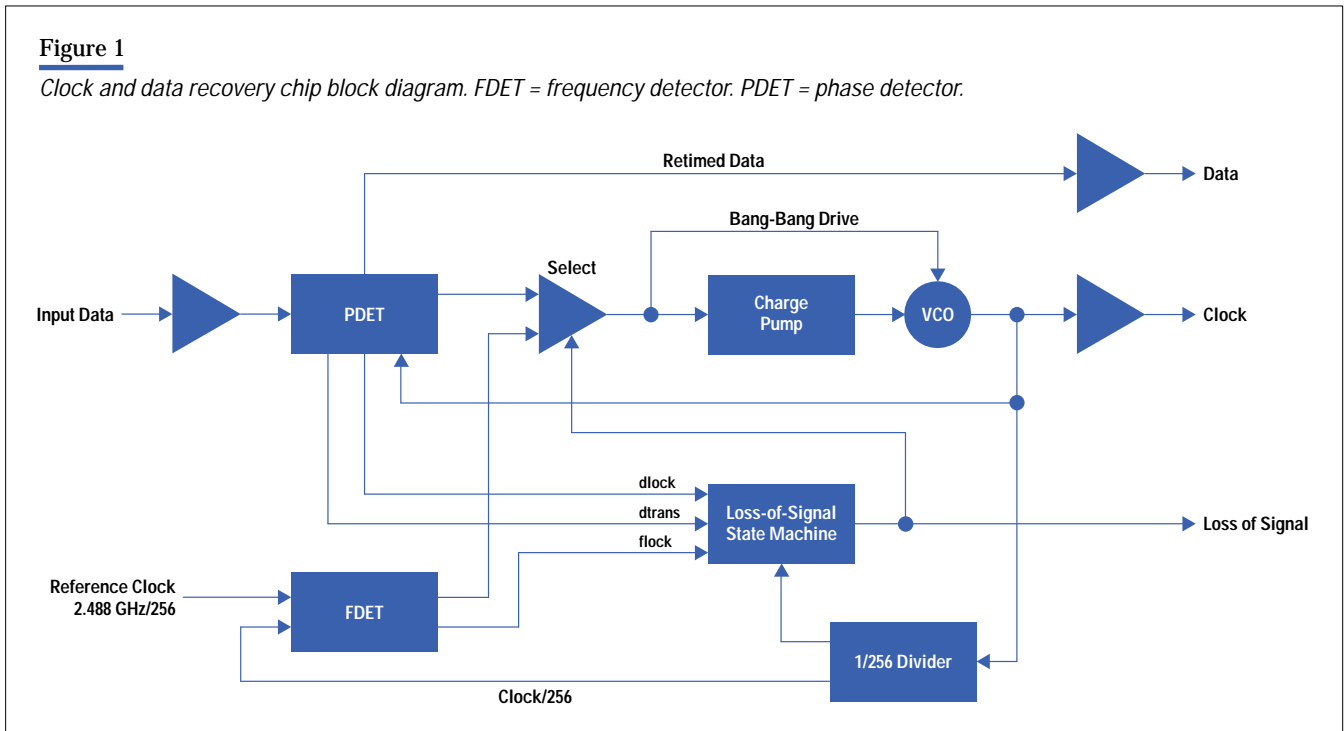
Building on techniques developed for the HP G-Link data communications chipset,³ the CDR 2500 has been optimized for telecommunications needs. The major differences between these two applications are how the data is coded for

transmission and how the circuit responds to jitter on the incoming data. In the G-Link application, a specific line code was used that guaranteed periodic transitions in the data. Such a code simplifies the clock recovery job. The telecommunications circuit, by contrast, must operate on scrambled data with random transitions.

Data communications links are typically much shorter than telecommunications links. The extra signal-to-noise ratio in a short link means that time jitter is much less of a problem than in a telecommunications link. Because of this, the HP CDR 2500 design required extra care to meet the more stringent jitter requirements of long-haul SONET applications.

Chip Block Diagram

A simplified block diagram of the chip is shown in **Figure 1**. **Table I** summarizes the chip characteristics. The input amplifier can be driven either single-ended or differentially and is ac-coupled and terminated on the associated thick-film hybrid. A Cherry-Hooper wideband amplifier⁴ at the input minimizes pulse width distortion with single-ended input signals.



The phase-locked-loop portion of the circuit includes two different phase detector blocks. To ensure that the on-chip VCO will operate at 2.488 GHz over variations in processing parameters, it has been designed with a 3:1 frequency tuning range. To avoid locking onto harmonics of the data signal, it is important that the VCO be kept close to the desired operating frequency. The first frequency/phase detector, labeled FDET in **Figure 1**, is used to frequency-lock the VCO initially to an external reference equal to 1/256 the desired bit frequency. The FDET block also produces a frequency lock error signal (flock), which is used by the state machine to sense when the VCO has reached the correct frequency. The second phase detector, labeled PDET, operates on random data and produces three signals: data transition detect (dtrans), data lock detect (dlock), and a tri-state bang-bang phase error signal.

Table I
Summary of Chip Characteristics

Parameter	Value
Guaranteed Frequency Range	2 to 3 Gbits/s
Supply Voltage	4.5V to 5.5V
Supply Current (nominal)	340 mA
Power Dissipation	1.77W
Case Temperature Range	0 to 60°C
Die Size (gate array)	3.45 × 3.45 mm ²
Number of Active Devices	3606
IC Process	25-GHz f _T Si Bipolar

The loss-of-signal (LOS) state machine block monitors the status outputs of the two detectors. On power-up, the state machine enables the frequency detector (FDET). When the VCO is frequency-locked to the reference and data transitions are detected on the data input, the phase-locked loop is switched to use the bang-bang phase detector. If the data lock detector stabilizes within an allowable time, then the loop remains phase-locked to the data. Otherwise, the VCO is retrained to the reference clock and the sequence repeats.

If the input data is interrupted or corrupted, the VCO needs to be retrained to the reference frequency input. Switching to the reference loop destroys the phase-locked condition between the VCO and the incoming data, causing errors in the retimed data output. In a telecommunications application, these errors could potentially interrupt service for thousands of telephone users. Therefore, a highly reliable algorithm must be used for determining loss of signal, so that the VCO is retrained to the reference only under high bit error rate conditions.

Loss-of-Signal Detection

The robust LOS detector is built by statistically processing the PDET data lock detector output. The data lock detector operates by monitoring the phase relationship between data zero crossings and the VCO output. Any rising or falling edge occurring more than ± 135 degrees away from the nominal location is flagged as a raw phase error.

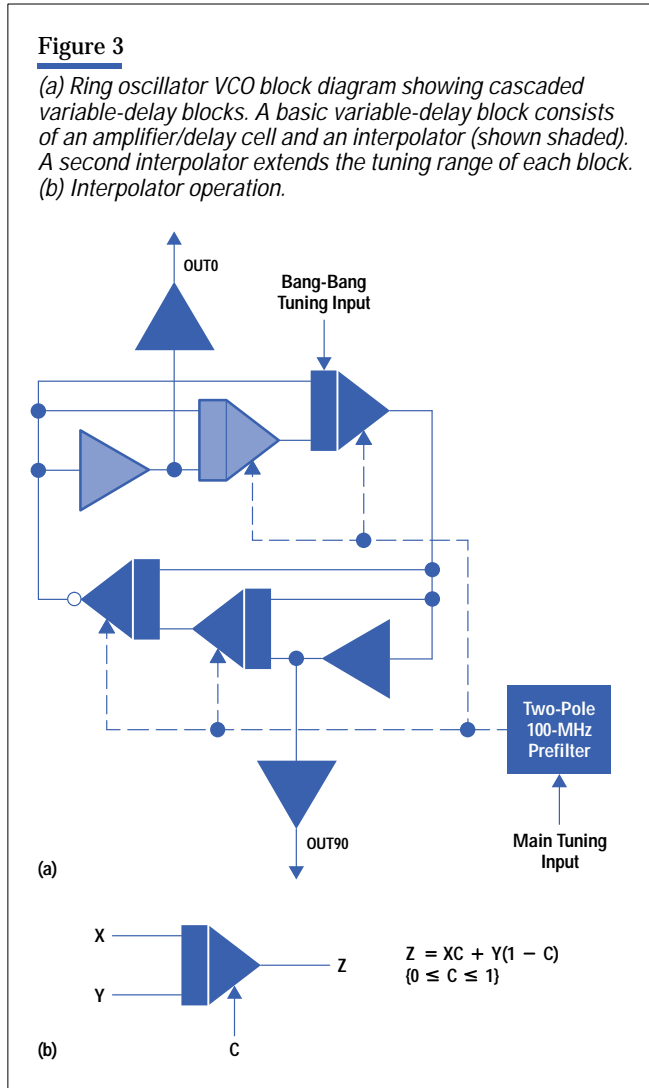
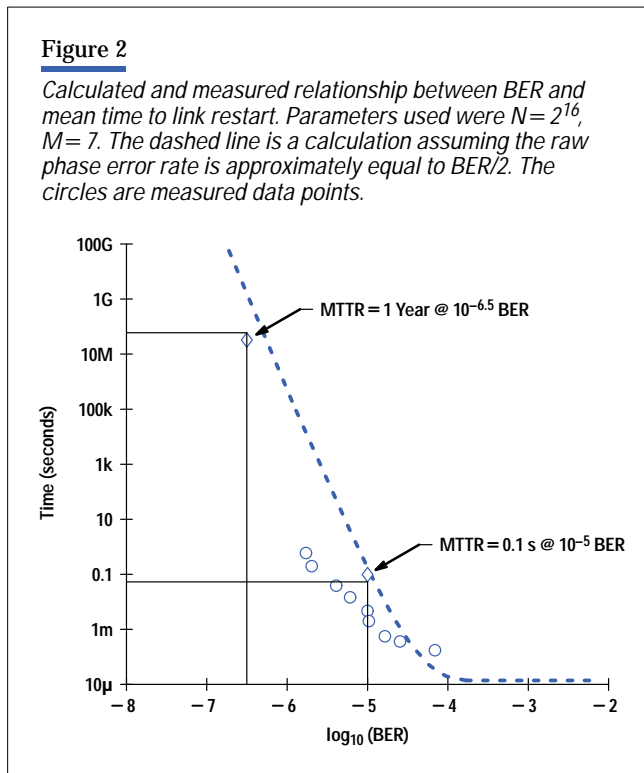
For typical noisy input signals, large input data phase excursions occur with approximately the same probability as the bit error rate (BER). Random data has half as many transitions as actual bits, so the raw phase error rate can be approximated as BER/2.

If we simply used raw phase errors to determine when to restart the link by relocking the VCO to the reference, then we would get a very unreliable and sloppy relationship between mean time to restart (MTTR) and BER. For a simple LOS algorithm in which the link is restarted on an isolated phase error event, the link would restart an average of once per second at a BER of 10^{-9} and every 10 seconds at a BER of 10^{-10} . This is a 1:1 relationship between BER and MTTR. In this simple system, an unrealistic BER of 10^{-16} would be required to ensure an MTTR of greater than one year between restarts. For system reliability reasons, we chose as a design target that the link should restart within 1/10 of a second at 10^{-5} BER, and should assert LOS less than once per year at 10^{-7} BER.

A way to steepen the LOS/BER relationship is to process phase error events in multibit bins and require that multiple consecutive bins each contain at least one error before asserting LOS. The current design groups phase error events into one of four programmable bin sizes (N), each requiring a different number of consecutive errored bins (M) before asserting LOS. Assuming that the phase error rate is approximately equal to half the BER, then if T_{bit} is the bit time,

$$MTTR \approx \frac{NT_{bit}}{(1 - (1 - BER/2)^N)^M}$$

The exponent M in the denominator of the MTTR expression sets the slope relation between phase error and MTTR. Setting M to 7 requires only one decade change in BER to change from an MTTR of one second to one year. Selection of one of the four possible states is programmable by two digital bond options, providing an LOS threshold between approximately 10^{-4} and 10^{-6} BER. The exact BER varies according to the application because of the system dependent relationship between phase error rate and BER. **Figure 2** shows measured MTTR versus calculated MTTR results for a system in which $M = 7$ and $N = 2^{16}$.



VCO

The VCO is composed of a cascade of two variable-delay blocks as shown in **Figure 3**. Each variable-delay block is composed of an amplifier/delay cell and an interpolator. A second interpolation section extends the tuning range while minimizing interpolation jitter. The shaded interpolator in **Figure 3** interpolates between a nondelayed input and an input

delayed by the shaded amplifier/delay cell. The shaded elements are recursively used as a delay cell for the unshaded interpolator (with the bang-bang input). This recursive connection minimizes the phase difference between the inputs of the interpolators.

The VCO outputs are taken from noncritical, highly buffered portions of the loop to minimize loading. The 3:1 frequency range of the circuit guarantees 2-to-3-GHz operation across process, temperature, and power supply variations. The main tuning input is prefiltered by two 100-MHz poles to reduce sensitivity to power supply noise. The wide-bandwidth bang-bang tuning input has 500 times less gain than the main tuning input, and is implemented by injecting small currents into the interpolation cell.

The phase-locked loop uses a bang-bang phase detector and a positive feedback charge pump similar to those described in references 3 and 5. The two-tuning-input VCO architecture, in conjunction with a binary-quantized phase detector, results in a VCO drive voltage equivalent to a first-order sigma-delta conversion of the instantaneous loop frequency error. By adjusting the relative gain of the two VCO tuning inputs, a trade-off can be made between jitter generation and jitter accommodation. Unlike a traditional analog loop, such a system cannot be characterized by a loop bandwidth. The ability of the loop to track incoming phase jitter is a slew-rate limited process, with an effective jitter bandwidth proportional to jitter amplitude. In practice, this behavior is ideal for the input of a SONET regenerator, for which a wide input jitter-tracking bandwidth is desirable to minimize sampling errors, and in which the overall system jitter transfer function will be set by a separate narrowband transmitter phase-locked loop. For this design, the bang-bang amplitude and charge pump time constant have been set to meet the SONET jitter tolerance specification. The resulting jitter generation, as calculated from the jitter spectrum in **Figure 4**, is 0.0049 UI rms (UI = unit interval = T_{bit}).

Figure 4

Clock phase noise measured under open-loop conditions and when locked to a $2^{23} - 1$ pseudorandom binary sequence (PRBS). Under locked conditions, the integrated value of the phase noise, filtered by a single-pole 5-kHz high-pass filter and a 100-MHz low-pass filter is 0.0046 UI, or 1.6 degrees rms.

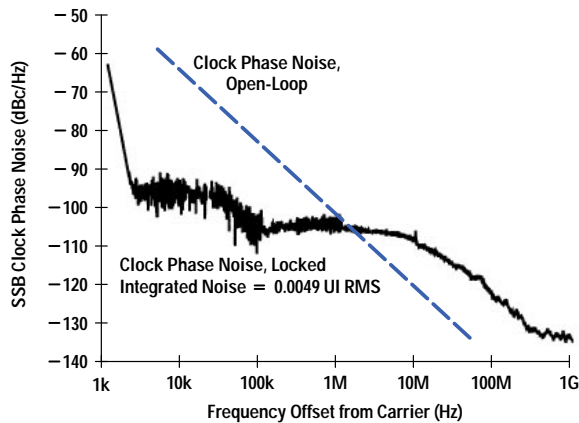


Figure 5

Top trace: Input $2^{23} - 1$ PRBS data with broadband noise added to achieve 10^{-4} BER. Middle trace: Recovered data eye. Bottom trace: Recovered clock. All signals are triggered from the BER tester clock input.

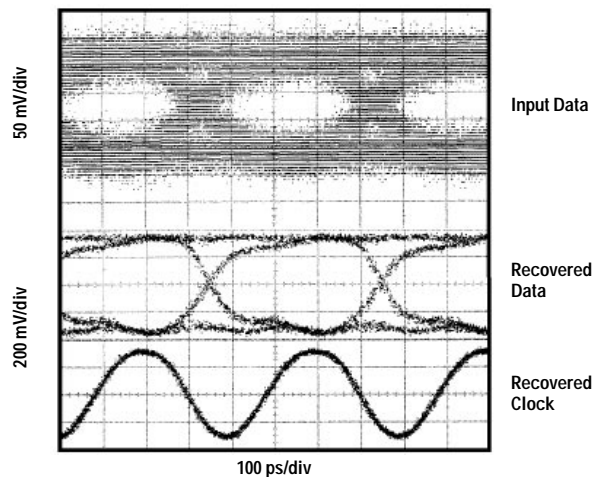
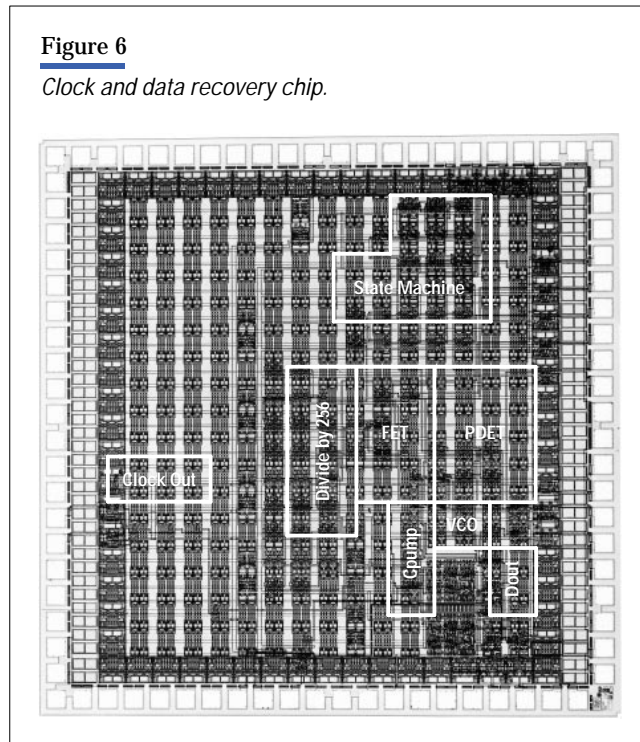


Figure 5 shows the recovered clock and data in the presence of noisy data. 3-GHz white noise was added to a $(2^{23} - 1)$ -bit pseudorandom binary sequence (PRBS) signal until the system was operating at 10^{-4} BER. The noisy input signal, the recovered data, and the recovered clock are shown in the top, middle, and bottom traces, respectively.

Figure 6 shows the die photo of the chip. The layout was done using a gate array methodology with fully differential ECL cells. Less than half of the array capacity was used in this design.



Package and Substrate

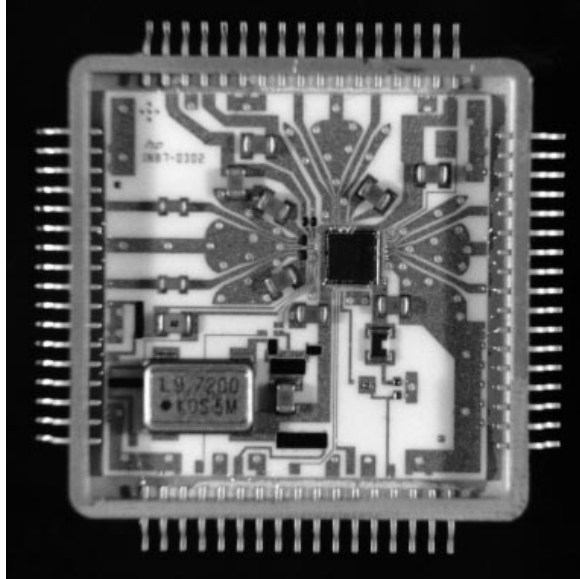
Figure 7 shows a picture of the assembled CDR 2500 multichip module. It consists of the HP-25 clock and data recovery IC attached to the 96% alumina substrate along with two diodes, a discrete transistor, 12 chip capacitors, and a chip resistor. The large packaged device in the lower left corner of the substrate is the reference oscillator. This is a commercially available surface mount crystal oscillator operating at $2.488 \text{ GHz}/256 = 9.71875 \text{ MHz}$. All devices are attached to the substrate using conductive epoxy. The two transistors visible on the hybrid are used to level shift the loss-of-signal output to be compatible with 5V CMOS. 0.001-inch gold wire bonding is used to connect to the die. The 1-in-by-1-in substrate is attached to a 68-pin Kovar package using conductive epoxy. The package leads are formed into a gull-wing configuration so that the package can be soldered to a printed circuit board. The module is tested at speed (2.488 Gbits/s).

The substrate is fabricated using proprietary HP enhanced thick-film processes. These processes make the high-speed performance of the packaged clock and data recovery circuit possible. They provide high-frequency performance and density rivaling what can be obtained using thin-film processing at a fraction of thin-film's cost.

Clock and data recovery circuits require excellent input return loss to prevent intersymbol interference caused by reflections at impedance discontinuities. The target for the CDR 2500 was to achieve better than 15 dB return loss up to 2.5 GHz. However, the discontinuities caused by the package's gull-wing lead and by the wire bonds that connect the substrate input and output lines to the package pins eat up much of the available return loss specification. Therefore, the return loss of the thick-film transmission lines and the input 50-ohm termination resistors had to be even better. That meant that tight control of the transmission line impedances was critical. To minimize high-frequency parasitics, it was important to use physically small, tightly controlled termination resistors.

Figure 7

Clock and data recovery multichip module.



The three HP proprietary thick-film processes used in fabricating the HP CDR 2500 are:

- Fine-line etched gold conductors
- Solid filled vias
- Small thick-film resistors.

The fine-line etched gold is the predominant pattern covering the substrate in **Figure 7**. Gold thick-film paste is printed across the substrate and is etched using a photolithographic process to produce the conductor pattern. The fine-line gold has several advantages over standard thick-film printing. It resolves very small dimensions—down to 0.002-inch lines and spaces. This can be seen in **Figure 8**, which shows a closeup of the area of the substrate around the pad for the HP-25 die. Several of the gaps between conductors leading to the die are 0.002 inch wide. This permits high pattern densities near the IC, making it possible to minimize wire-bond length and bringing the bypassing capacitors very close to the chip. Both of these factors help to minimize noise in the circuit. Having short, consistent wire-bond lengths is critical for achieving good return loss. Also, the fine-line etch process permits much better control of conductor line widths and spaces, needed to maintain accurate 50-ohm transmission line impedances. The etched fine-line conductors also have straight, well-defined side walls, which offer good high-frequency performance.

Solid filled vias can be seen in the ground regions in **Figure 8**. The vias are used to connect ground pads on top of the substrate to the backside ground plane. The solid filled via is formed by printing a solid plug of paste into the hole, filling the hole completely. This provides a very consistent, low-resistance ($\ll 25$ milliohms), low-inductance path to ground at RF frequencies. In **Figure 8**, there is a ground pad that has a solid filled via near the upper right corner of the pad for the IC. Note how small the gold pad is—only about 0.002 inch larger than the via itself. Two bond wires are bonded within 0.002 inch of the via.

A closeup of the input section near the IC pad is shown in **Figure 9**. The four black elements are thick-film resistors. These are the small HP resistors, which can be as small as 0.010 inch by 0.010 inch, very small for thick-film technology. The two lower resistors are 0.010-by-0.012-inch, 50-ohm termination resistors for the circuit's input lines. They are laser

Figure 8

A close-up of the clock and data recovery module thick-film circuitry near the HP-25 die pad.

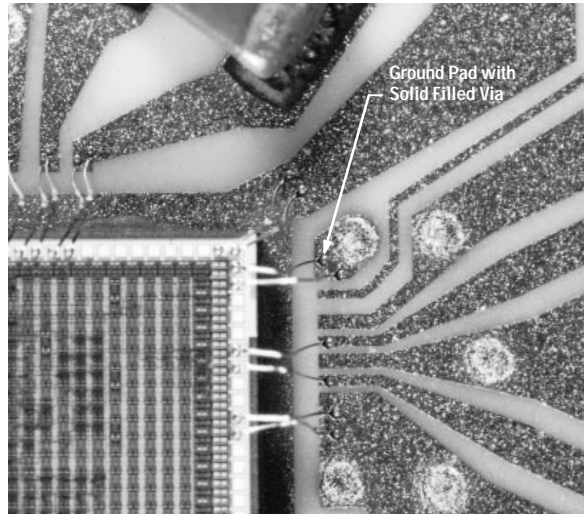
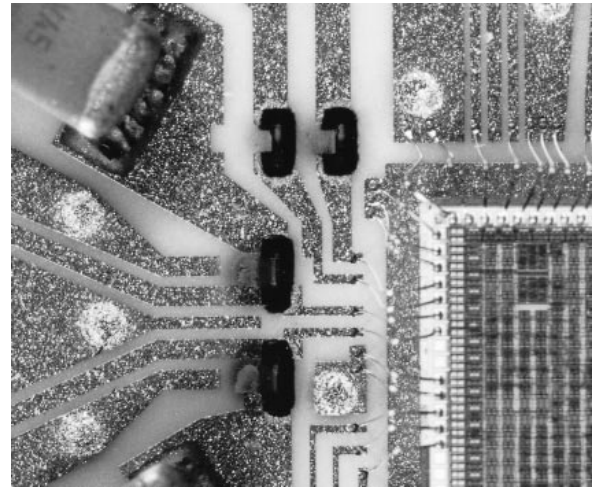


Figure 9

A close-up of the clock and data recovery module's thick-film input section showing small resistors (black elements).



trimmed to $\pm 0.5\%$ tolerance. Their small size, which minimizes their high-frequency parasitics, and their tight tolerance make them excellent terminations. The two resistors above the terminations are 0.010-by-0.012-inch 1-kilohm bias resistors. They are made small to maximize circuit density around the chip and keep bypass circuitry as close to the die as possible.

Substrate Technology Selection

Why is the HP proprietary thick-film on ceramic technology used for this application? Why couldn't a laminate technology, which is generally slightly less expensive, be used instead? Primarily, because it would be difficult to achieve the electrical performance. A laminate would have difficulty holding the tight transmission line tolerances that are required. If a laminate were used, the ten resistors now implemented using thick-film technology couldn't be integrated into the substrate. In general, this would increase cost and reduce the circuit density around the chip, thereby moving bypassing farther from the IC and degrading the electrical performance. For the input termination resistors in particular, using add-on resistors would add parasitics, making the return loss specification even harder to meet. Keeping the size of the circuit as small as possible is important, and using add-on resistors would probably necessitate an increase in the size of the device.

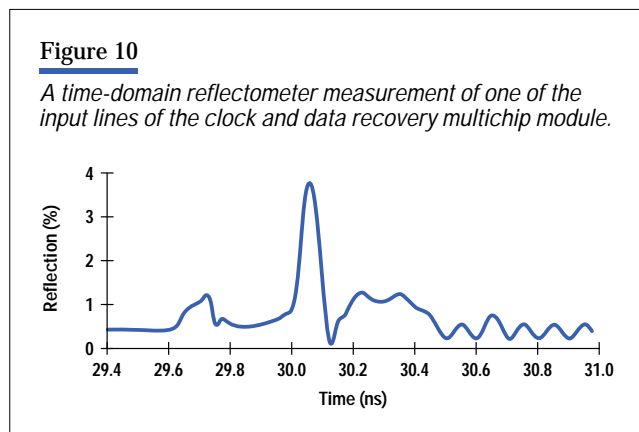
The HP-25 die dissipates 1.8W. For the original application of the HP CDR 2500, the circuit sits in a Kovar package in an environment with no air flow and no room to attach a heat spreader or a heat sink. Thus, the thermal performance of the device is important, and using a laminate instead of ceramic for the substrate would substantially increase the die's junction temperature. The die must be back-biased to -5.2V , so cutting a hole in the substrate and attaching the die directly to a pedestal brazed to the package is not an option. Attaching it to an electrically insulating heat slug is possible, but the added cost of this approach would negate some of the cost savings of the laminate.

Design

To determine that thick-film fine-line transmission lines could be controlled well enough to meet the CDR 2500's return loss specifications, a Monte Carlo analysis was run on the HP Microwave Design System (HP MDS). This analyzed the change in transmission line characteristic impedance as a function of the inherent variations in HP's fine-line conductor

process. Changes in conductor line width, substrate dielectric constant, and substrate thickness are the primary factors that affect characteristic impedance. HP MDS varies these parameters, and the result is a probability density function and a standard deviation for the resulting changes in characteristic impedance. The changes in characteristic impedance were put into a SPICE simulation of the clock and data recovery circuit's input. This confirmed that a fine-line solution could produce acceptable results in a manufacturing environment.

HP MDS was used to design the multichip module's critical transmission lines, after previously correlating the MDS models to measured data of the HP processes. The design of the HP CDR 2500 multichip module worked the first time it was prototyped. **Figure 10** is a time-domain reflectometer (TDR) display of one of the prototype input transmission lines, measured using a 160-ps pulse. The vertical scale is 2% reflection per division. The first inductive discontinuity (at 29.7280 ns) is caused by the test fixture used to probe the package. The second and largest inductive discontinuity is the wire bond leading from the package pin to the substrate. The region immediately following that (for nearly two time divisions) is the microstrip and grounded coplanar waveguide transmission line. It shows very little reflection, and there is no discontinuity at the 0.010-inch 50-ohm termination resistor.



The production version of the HP CDR 2500 replaces the package-to-substrate wirebonds with 0.001-by-0.005-inch ribbon bonds. This lowers the inductive discontinuity and further improves the module return loss.

Trends in Serial Communication

Long-haul transmission capacity has traditionally grown by expansion of bit rates. The last 20 years has seen data rates jump by over three orders of magnitude. The need for higher rates is motivated by the potential capacity of existing fibers and by the high cost of installing new fiber-optic cables.

Recent development of fiber broadband optical amplifiers and workable wavelength-domain multiplexing (WDM) techniques have somewhat changed this trend. Because WDM allows multiple high-speed channels to share the same fiber, there is less of a need for heroic efforts to extend the serial data rate of each channel. Pulse dispersion at high bit rates and the limitations of electrical circuits make it likely that individual link speeds will top out at 10 to 40 Gbits/s for the foreseeable future.

High-speed clock and data recovery circuits will still be needed for each wavelength channel, and the market demand for such circuits will continue to grow in relation to world data transmission needs.

As the industry becomes more mature, we will see increasing pressure on vendors to increase the integration level of their link ICs by including limiting amplifiers, demultiplexers, and payload processing circuitry on their chips. This is starting now at lower data rates (e.g., 155 Mbits/s), and will proceed up to higher rates as IC processes improve.

The relatively uncritical data communications needs will likely be met adequately with future CMOS processes. Gigabit-rate clock and data recovery circuits can already be built using multiphase sampling in current CMOS processes. Such architectures achieve high speed at the cost of systematic jitter caused by mismatched clock phases. For jitter-sensitive telecommunications applications, it is likely that bipolar processes will be needed for at least the next five to ten years. At the end of that time, bipolar implementations will likely be less desirable, the reason being that, although bipolar processes provide superior front-end clock and data recovery performance, they are not capable of the extremely high integration levels required for payload processing. A compromise approach would be to add relatively low-performance CMOS devices onto an existing high-speed bipolar process. Such a process would then be an ideal compromise between the precision high-speed, low-jitter front-end clock and data recovery requirements and the need for large-scale digital payload processing on the same chip.

Conclusion

The CDR 2500 is a step forward towards higher integration density for SONET 2.488-Gbit/s systems. HP's 25-GHz f_T bipolar process provides good jitter performance at a lower cost than previous multichip GaAs implementations for this data rate. The HP CDR 2500 takes advantage of proprietary HP thick-film technology that includes etched fine-line gold conductors, solid filled vias, and small resistors to provide excellent RF and high-speed digital performance at a lower cost than thin-film technology.

Acknowledgments

The CDR 2500 clock and data recovery multichip module was a collaborative effort between HP Laboratories, the HP Colorado Springs Technology Center (CSTC), the HP Fiber-Optic Components Operation (FCO), and the HP Optical Communications Division (OCD). Dr. Soo-Young Chai of Samsung (formerly of AT&T) provided valuable customer input for the design. The authors thank Jean Norman for her support in hybrid microassembly. Pat Petruno, Richard Dugan, and Virginia Graham of OCD were invaluable in allowing us to make prototype fab runs on the OCD Masterslice HP-25 gate array. We also wish to thank Teik Goh of FCO and Jack Casey, Steve Meyer, Jim Taylor, Steve Chaffee, Randy Fuchs, Don Griffith, and Albert Yeh of CSTC for their help in transferring this design into production.

References

1. H. Ransijn and P. O'Connor, "A PLL-Based 2.5-Gbit/s GaAs Clock and Data Regenerator IC," *IEEE Journal of Solid State Circuits*, Vol. 26, no. 10, October 1991, pp. 1345-1353.
2. W.M. Huang, et al, "A high-speed bipolar technology featuring self-aligned single-poly base and submicrometer emitter contacts," *IEEE Electron Device Letters*, Vol. 11, no. 9, September 1990, pp. 412-414.
3. C.S.Yen, R. Walker, P. Petruno, C. Stout, B. Lai, and W. McFarland, "G-Link: A Chipset for Gigabit-Rate Data Communication," *Hewlett-Packard Journal*, Vol. 43, no. 5, October 1992, pp. 103-116.
4. D Faulkner, "A Wide-Band Limiting Amplifier for Optical Fiber Repeaters," *IEEE Journal of Solid State Circuits*, Vol. 18, no.3, June 1983, pp. 333-340.
5. B. Lai and R.C. Walker, "A Monolithic 622Mb/s Clock Extraction Data Retiming Circuit," *International Solid State Circuits Conference Digest*, 1991, pp. 144-145.



Richard Walker

Rick Walker joined HP Laboratories in 1981, and is currently a principal project engineer. He develops and designs integrated circuits for gigabit-rate serial data transmission. He received a BS degree from the California Institute of Technology in 1982 and an MS degree from the California State University at Chico in 1992. His interests include tai'chi chuan, playing guitar, and raising exotic plants.



Cheryl Stout

Cheryl Stout is a hardware design engineer at HP Laboratories. She currently develops CMOS clock and data-recovery circuits. In her 14 years with HP Laboratories, she has designed gallium arsenide and silicon integrated circuits for gigabit-rate serial data transmission. She received a BSEE degree from the California State University at San Jose in 1979 and an MSEE degree from the University of California at Berkeley in 1983. Her interests include mountaineering and natural history.



Chu-Sun Yen

Chu-Sun Yen has done research and development at HP since 1961. He was project manager for both the HP CDR 2500 clock and data recovery circuits and the G-link chipset, and has authored 20 papers on circuits, phase-locked loops, instruments, and communications links. He holds a doctorate from Stanford University.



Lewis R. Dove

Lew Dove is an engineer-scientist at HP's Colorado Springs Technology Center, where he works on thick-film packages for high-frequency multichip modules. He has BSEE and MSEE degrees from the University of Arizona. He is married and has two sons, and his hobbies include flying, running, and skiing.

-
- ▶ [Go to Next Article](#)
 - ▶ [Go to Journal Home Page](#)