DALLAS SEMICONDUCTOR MAXIM

# Engineering Journal

Volume Forty-Six



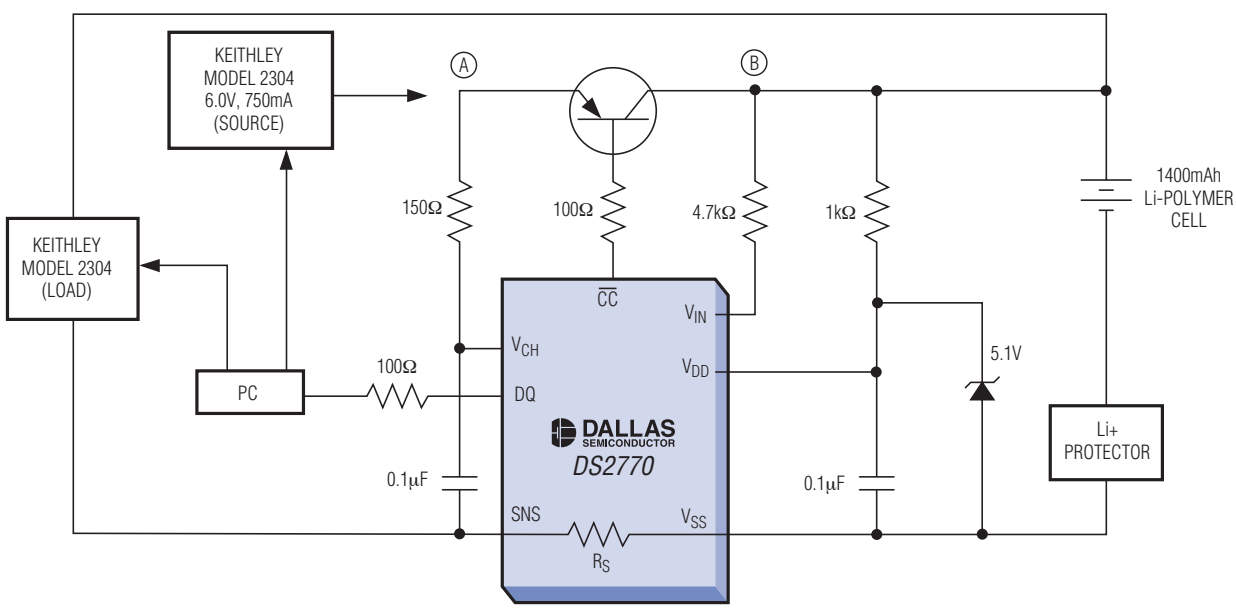


*The DS2770 pulse charger charges the cell 16% faster and ages the cell 4.6% less than a linear charger after 500 charge/discharge cycles at +25°C.*

# News Briefs

## MAXIM REPORTS REVENUES AND EARNINGS FOR THE FIRST QUARTER OF FISCAL 2003 AND DECLARES QUARTERLY DIVIDEND

Maxim Integrated Products, Inc., (MXIM) reported net revenues of $285.9 million for its fiscal first quarter ending September 28, 2002, a 19.4% increase over the $239.4 million reported for the same quarter a year ago and a 2.1% sequential increase over the $280.1 million reported for the fourth quarter of fiscal 2002. Net income increased to $73.2 million in the first quarter, compared to $61.3 million last year, a 19.4% increase. Diluted earnings per share were $0.22 for the first quarter, a 29.4% increase over the $0.17 reported for the same period a year ago.

During the quarter, cash and short-term investments increased $22.2 million after the Company repurchased 2.0 million shares of its common stock for $69.0 million and acquired $38.7 million of capital equipment. Accounts receivable increased by $5.6 million in the first quarter to $135.4 million due primarily to the increase in net revenues, and inventories decreased $2.9 million to $136.3 million.

Inventory reserves increased $3.0 million. Research and development expense decreased from the $72.0 million reported in the fourth quarter of fiscal 2002, or 25.7% of net revenues, to $71.1 million, or 24.9% of net revenues, in the first quarter of fiscal 2003. Selling, general, and administrative expenses remained relatively unchanged from the fourth quarter at $22.3 million, or 7.8% of net revenues.

First quarter bookings were approximately $270 million, a 13% decrease from the fourth quarter's level of $310 million. End market bookings for the first quarter were 11% below the fourth quarter's level. Turns orders received in the quarter were $129 million, an 8% decrease from the $140 million received in the prior quarter (turns orders are customer orders that are for delivery within the same quarter and may result in revenue within the same quarter if the Company has available inventory that matches those orders). Bookings decreased in all geographic locations and across most product lines.

First quarter ending backlog shippable within the next 12 months was approximately $219 million, including approximately $193 million requested for shipment in the second quarter of fiscal 2003. The Company's fourth quarter ending backlog shippable within the next 12 months was approximately $239 million, including approximately $210 million that was requested for shipment in the first quarter of fiscal 2003.

Jack Gifford, President, Chief Executive Officer, and Chairman, commented on the quarter: "Bookings during the first quarter were lower than those in the prior two quarters, indicating that third and fourth quarter bookings levels may have been in anticipation of expected growth, and not entirely reflective of demand in those quarters. Although design-in activities at our customers remain vigorous, customers do not yet have improved visibility of the demand for their products and thus continue to place orders for the near term. The level of turns orders received in the first quarter, 48% of net bookings, reflects this continued limited visibility at our customers.

"Maxim continues to be highly profitable in a challenging economic environment. As planned, gross margin improved this quarter, contributing to the improvement in operating margin. We expect our operating margin to continue to improve in the second quarter as we implement a variety of cost-saving measures."

Mr. Gifford concluded: "Given the stability of the Company's financial performance, its Board of Directors has declared a quarterly cash dividend of $0.02 per share. The Company is fortunate that its earnings and strong cash position permit the payment of a dividend without impairing our cash reserves or continued growth. Payment will be made on November 29, 2002 to stockholders of record on November 8, 2002."

Certain statements in this press release are forward-looking statements within the meaning of the Private Securities Litigation Reform Act of 1995. These statements involve risk and uncertainty. All forward-looking statements included in this news release are made as of the date hereof, based on the information available to the Company as of the date hereof, and the Company assumes no obligation to update any forward-looking statement.

Maxim Integrated Products is a leading international supplier of quality analog and mixed-signal products for applications that require real world signal processing.

# Charging efficiency vs. cell aging in Li+ pulse chargers

*The pulse-charging topology shares many of the same advantages as linear charging including simplicity, low cost, small size, and few external components. Its benefit over linear chargers is that heat is dissipated in the wall adapter instead of the charger itself. A perceived disadvantage of pulse charging a Li+ cell is premature aging resulting from exposure to the same current during the constant-current region and the pulsed region, where it is very near its maximum voltage. This article presents an overview of pulse and linear (CCCV) charging, followed by comparison of charge time, charge capacity, and cell aging between a CCCV charger and the DS2770, which features a Li+ pulse charger.*

## Linear charging

Linear charging is widely used for charging Li+/polymer batteries at currents typically less than 1.5A. The chargers are low cost, low noise, small in size, and require few, inexpensive, and small external components. The charge source is a low-voltage unregulated supply (typically 5V to 6V for single-cell applications). A high-side internal or external low-impedance sense resistor is typically used to accurately measure the charging current. The charger controls a pass element, depicted in **Figure 1** as an external P-FET, to regulate cell voltage or current. Figure 1 shows a typical implementation of a linear charging circuit and Li+ CCCV charge profile.

The charger should prequalify the cell to determine if a charger is present and if it is safe to subject the cell to a fast charge rate. If the cell voltage is below approximately 2.5V, it could be permanently damaged if a 1C charge is applied to it. Thus, the charger gently trickle charges the cell at a rate on the order of C/15 until its voltage increases to the prequalification threshold.

Once the cell is prequalified, the charger regulates the FET so the source delivers a constant current to the cell. If equipped with a safety timer, the charger starts the timer when fast charge is initiated. The drawback of linear chargers is most evident under fast charge with the cell voltage at a low level. If the unregulated charge source delivers 5V and a 750mAh cell is at 3V under a 1C charge, the FET is dissipating 1.5W. Figure 1's charge profile is a simple illustration; the voltage increase under constant

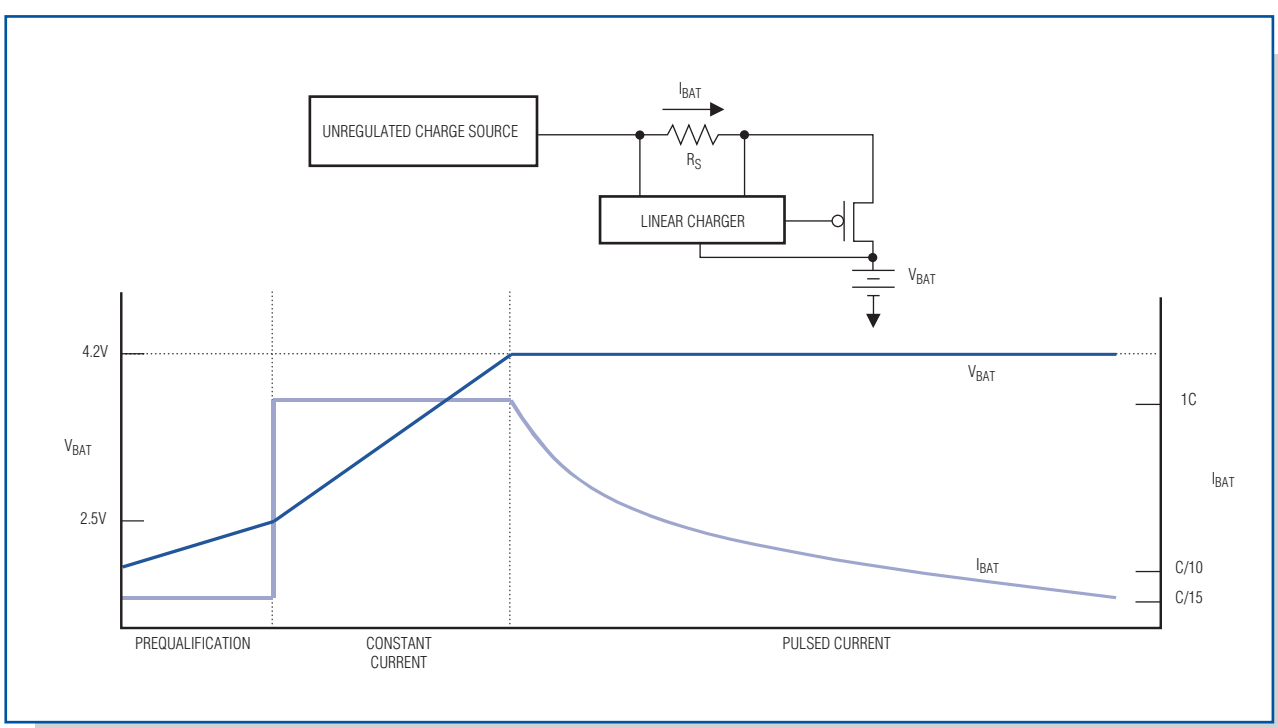


*Figure 1. A linear charger profile has a constant-current region, followed by a constant-voltage topoff.*

current is actually not linear. In reality, the cell's voltage will not be at 3V for very long, and the FET's power dissipation decreases as the battery voltage increases under constant current. The charger continues to supply the constant current until the cell voltage reaches a constant voltage ($V_{CV}$) threshold. This threshold is typically 4.1V or 4.2V.

When the cell reaches the constant-voltage threshold, it is at approximately 40% to 70% of its capacity. To "top off" the cell, the charger regulates the FET to supply a constant voltage equal to the $V_{CV}$ threshold. During this phase of the charge cycle, the charging current decreases over time. Constant voltage persists until the measured current drops to C/10 or C/15 or until the safety timer expires.

## Pulse Charging

Pulse charging is very similar to the CCCV method except for the charge source and the charge-termination method. It shares the same benefits as linear chargers but does not have the drawback of power dissipation in the battery pack itself. The pulse charger requires an unregulated, current-limited charge source. A typical pulse charger configuration and Li+ charge profile is illustrated in **Figure 2**.

Like the linear charger, the device first determines if a charger is present, which generally depends on whether the charge-source voltage is greater than the battery voltage. It then prequalifies the cell and provides a current-limited trickle charge path if the cell is less than the prequalification threshold.

After prequalification, the safety timer starts and the charger turns on the external FET or PNP transistor, charging the cell with a constant current equal to the current limit of the charge source. Because the source is current limited, there is minimal drop across the FET and thus very little power dissipation in the battery pack. The current limit losses are localized to the wall adapter as either an $I^2R$ loss in the secondary or as a coupling loss in the transformer.

After the cell voltage reaches $V_{CV}$, the pulse mode begins. Similar to the linear charger, the cell is at approximately 40% to 70% capacity following the constant-current region, and the pulse charge region "tops off" the cell. The charger keeps the FET conducting for $t_{PULSE}$ after the cell voltage reaches $V_{CV}$. It then turns the pass element off, shutting off charge current and pulling the charge source to its open-circuit voltage. It remains in this state until the cell voltage decays to $V_{CV}$. The charger again turns on the pass element for $t_{PULSE}$, then turns it off until the cell voltage decays to $V_{CV}$. This loop repeats until the pulse charge duty cycle [$t_{PULSE}$ / ($t_{PULSE}$ + $t_{OFF}$)] falls below typically 5% to 10%. The pulse charge is then terminated and the cell is near its full capacity.

The pulse-charge topology requires a source with a current limit equal to the desired charge rate. Some pulse chargers feature an optional low-impedance sense resistor to qualify the charge source; fast charge is disallowed if the limit is exceeded. Because the pulse-charge topology does not require current regulation like the linear charger, qualifying the charge source is the only purpose served by

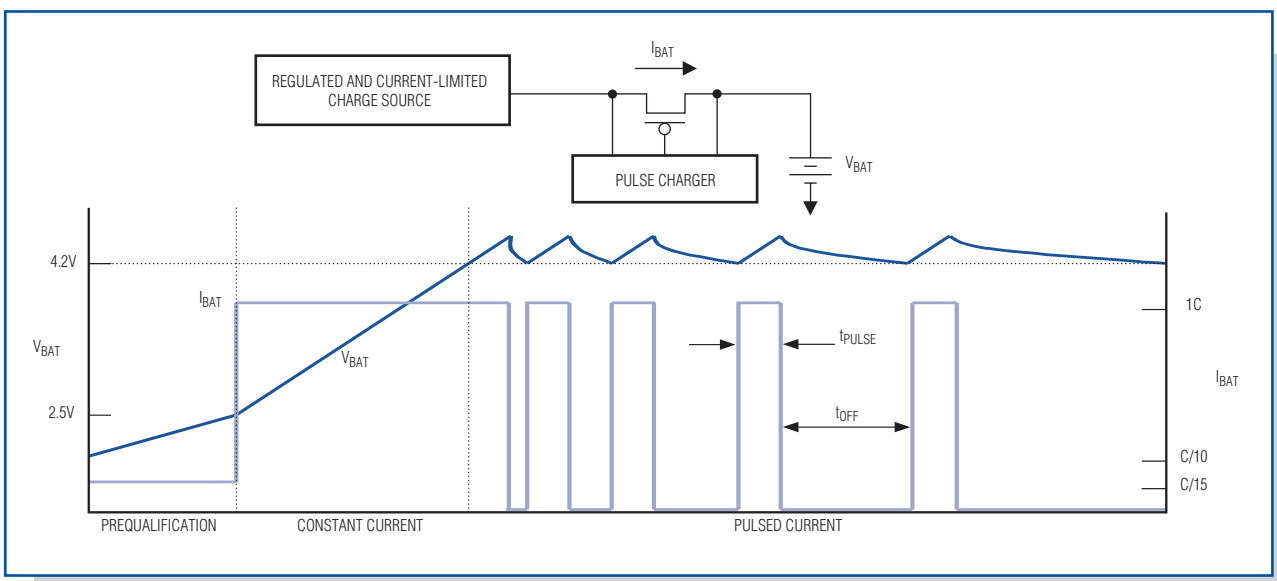


*Figure 2. A pulse-charger profile tops off the cell with pulses equal in amplitude to the constant-current level.*

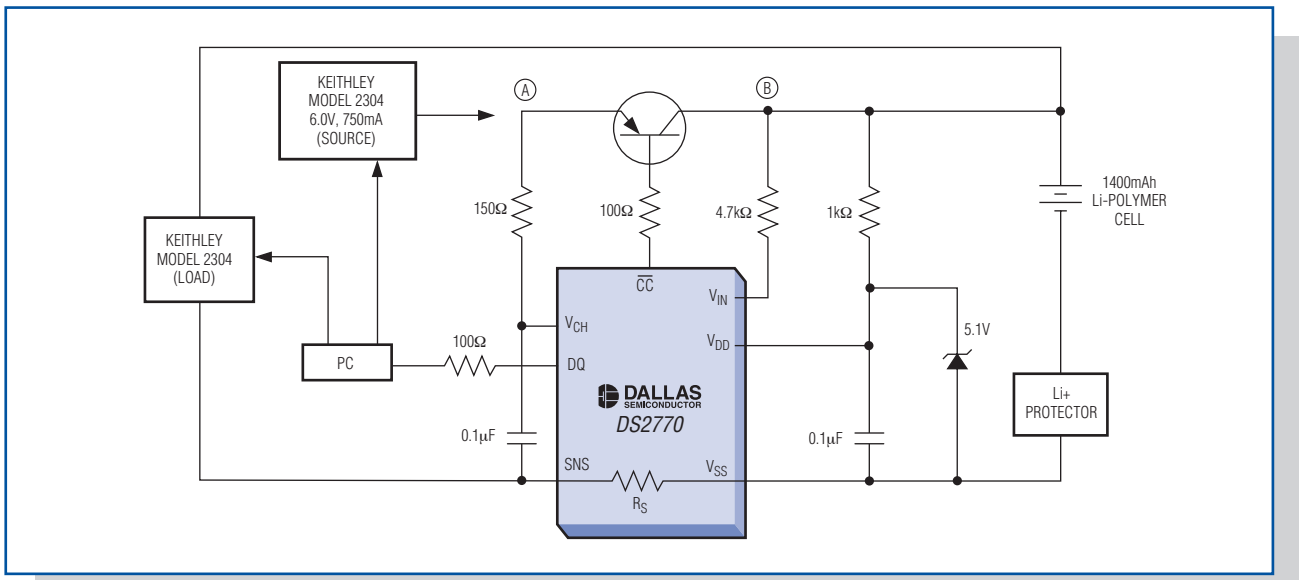


*Figure 3. The DS2770 is used for the pulse charger and for the coulomb counter in both the linear- and pulse-charge experiments.*

the sense resistor in a current-limited environment. Some pulse chargers do not offer current sensing, and rely on the Li+ protection circuit to open the battery pack if a non-spec charge source is used.

The general concern with the pulse-charging profile is that the charger subjects the cell to the full charge rate in pulses of width $t_{PULSE}$ (typically less than 1s) while the cell voltage is at $V_{CV}$, thus prematurely aging the cell. As Figure 2 indicates, this causes the cell voltage to increase slightly above $V_{CV}$ for $t_{PULSE}$ for as many pulses required before the 5% to 10% duty cycle requirement is met. Many have avoided using pulse chargers in the past because of the perceived premature aging effect, but does this profile age the cell any faster than a CCCV profile with equal peak-charge current? The next section describes a characterization procedure and presents results to challenge this concern about pulse chargers.

## Charge time and cell aging experiment

The DS2770 battery monitor and charge controller, which features a Li+ pulse charger, was used as a pulse-charger and a coulomb counter for the experiment. The coulomb count measured diminishing capacity as the cell cycled. **Figure 3** illustrates the test setup used for the experiment.

The charge source used was a Keithley Model 2304 power supply. It was set to an open-circuit voltage of 6.0V and a current limit of 750mA for the pulse-mode cycles. The Keithley was connected to node A in Figure 3 for the pulse-mode cycles. The DS2770 charged the cell at the current limit until the cell reached 4.2V, then transitioned into the pulsed phase. Pulse widths ($t_{PULSE}$) were 875ms (typ) and charge was terminated when $t_{PULSE}$ + $t_{OFF}$ = 14s (typ), a duty cycle of 6.25%. Refer to the DS2770 data sheet for detailed specifications.

For the linear-charge cycles, the power supply was programmed to deliver 750mA at 6.0V during the constant-current region. It was connected to node B for the linear cycles, thus bypassing the DS2770's pulse charger. The DS2770 measured the cell's voltage during the constant-current phase and its current during the constant-voltage region. When the voltage measurement reached 4.2V, the power supply was reprogrammed to supply a constant 4.2V to node B. Charge was terminated and the cell was considered fully charged when the battery current measured by the DS2770 dropped to 50mA.

A second Keithley 2304 was used to simulate the load for the discharge cycles. All discharge cycles were done at +25°C and had an identical profile, regardless of the mode used to charge the cell. The load current was set to 750mA DC until the cell voltage decayed to 3.4V. It was then dropped to 250mA DC again until the cell voltage fell to 3.4V. Finally, the load was reduced to 50mA, and the cell was considered fully discharged when the cell voltage fell to 3.4V under 50mA load. The difference in the DS2770's current accumulation register (ACR) between the start and end of the discharge period was

logged as the usable capacity of the cell for each given charge cycle.

## Charging time and usable capacity over temperature

A new 1400mAh Li-polymer cell was used to compare charge time over temperature for the DS2770 pulse charger and the discrete linear charger. Since a cell's capacity can vary drastically during its first few charge/discharge cycles, this cell was "broken in" by cycling it ten times before taking any measurements.

The experiment began by charging the fully discharged cell using the DS2770 pulse charger at +40°C. The time (using PC clock) between start and end of charge was logged. The chamber was then returned to +25°C, and the cell was fully discharged using the profile and "empty" criteria described earlier. The ACR difference from start to end of discharge was logged as the usable capacity of the cell under that charge profile at that temperature. The cell was then charged using the linear profile at +40°C, again logging charge time. It was discharged at +25°C, recording ΔACR. This process was repeated at charge-cycle temperatures ranging from 0°C to +40°C in 5°C increments, and the results appear in **Figures 4** and **5**.

Usable capacity is virtually the same for the two charging topologies at a given temperature, with the pulse charger performing slightly better at high temperatures and the linear charger performing slightly worse. The difference in charging time, however, is significant. At low temperatures charge time was 37% longer with the constant-voltage top-off than for the pulsed-mode top-off. At +40°C, the pulse charger was 14% faster. These results show that the pulsed-mode charger is significantly faster, but charges the cell with virtually the same usable capacity as the linear charger.

## Cell aging comparison between linear and pulsed chargers

With faster charge time demonstrated for pulse chargers, the next measurement to address is the cell's aging effect relative to linear-mode charging. Two brand new 1400mAh Li-polymer cells were cycled ten times to break them in. One cell was submitted to 500 charge/discharge cycles at +25°C under a linear profile and the other cell was charged by the DS2770. The discharge profile was the same as described for the charge-time experiment. Cell aging was based on decreasing usable capacity, as measured by the DS2770's ACR difference before and after each discharge cycle. **Figure 6** shows results of the experiment presented as a percentage of the usable capacity at the first discharge.

This experiment demonstrates that the aging effect from pulse-mode charging is actually less than that of the linear mode. The cell cycled with the CCCV profile had 4.6% less usable capacity than the pulse charger after 500 charge/discharge cycles at +25°C.

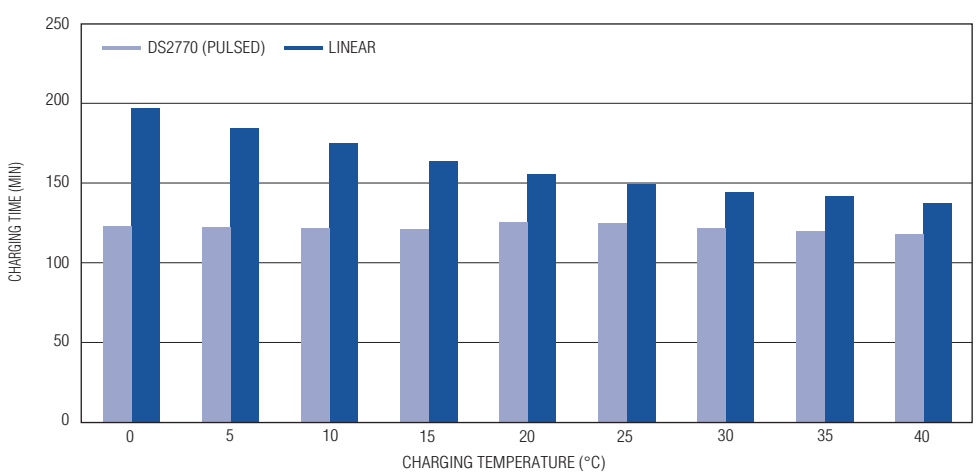


*Figure 4. The pulse-charge time is relatively constant with temperature and 16% faster than the linear charger at +25°C.*

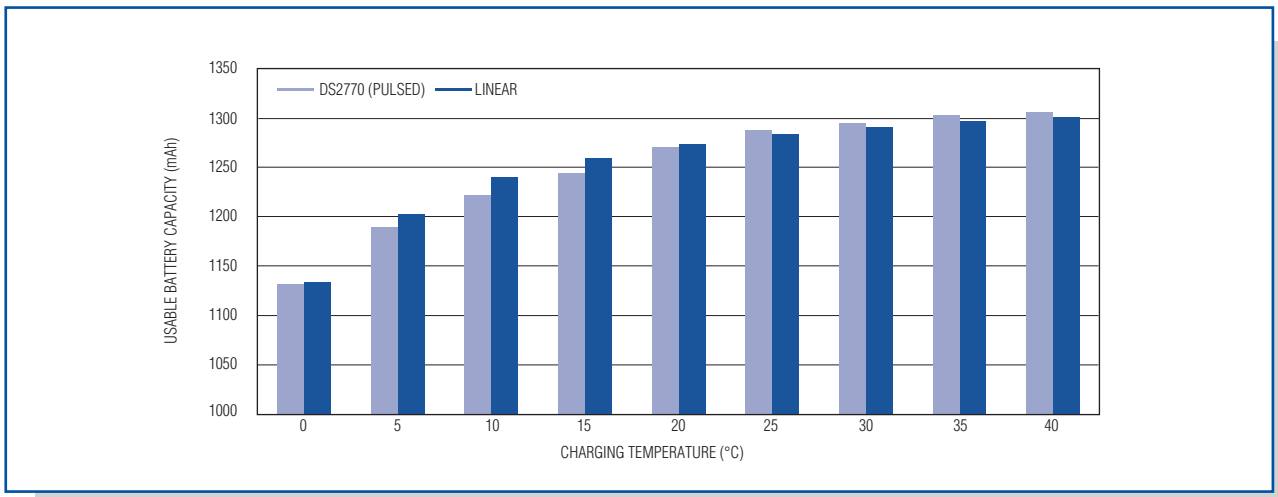


*Figure 5. Usable cell capacity for pulsed and linear charger over charging temperature is virtually identical.*

## Conclusion

Both linear and pulsed-mode Li+ chargers share many advantages including simplicity, low cost, small size, and few external components. An overview of each topology was presented, highlighting known and perceived disadvantages. The linear charger dissipated power in the pass element in the battery pack, requiring the pack designer to accommodate for the resulting heat. The pulse charger required a current-limited charge source, but that allowed the power to dissipate in the charge source and not the battery pack. A perceived disadvantage of the pulse-charge profile was cell aging as a result of pulsing 1C charge current with the cell at $V_{CV}$. Using the DS2770 battery monitor and charge controller, experiments demonstrated a) the pulse charger was 16% faster than the linear charger at +25°C with roughly equal usable capacity, and b) the pulse charger actually aged the cell 4.6% less than the linear charger after 500 charge/discharge cycles at +25°C.

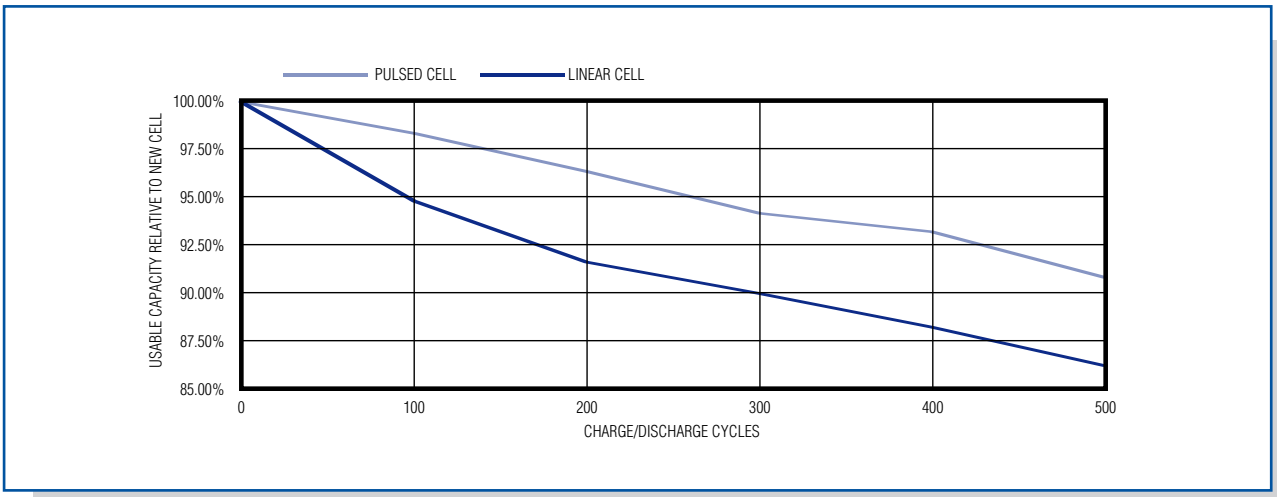


*Figure 6. Usable capacity is 4.6% less with linear charging than with pulse charging after 500 charge/discharge cycles.*

# Precision reference clock usage in clock and data recovery circuits

*Applications for clock data recovery (CDR) abound in telecommunications, optical transceivers, data and storage area networks, and wireless products. The benefits of CDR technology are increasingly important as designs require much greater bandwidth, and as allocation and spectral bandwidth usage increase. Moreover, vendors and their products are migrating from parallel to serial interfaces for system- and board-level interfaces.*

In recent years, the increased use of CDR technology has brimmed from the need to handle wider parallel bus widths across backplanes while managing clock and data skew at the receiver. Additionally, routing these signals can be difficult because they consume board space and power, and require multilayer routing schemes to manage signals and line termination. EMI generation from the use of high bit-width data buses is also a concern.

CDRs are extremely important due to the advent of new communication technologies, improvements in electrical signal processing, and the need to send multigigabit electrical signals across FR-4 and backplanes, optical, and wireless media. Communication techniques that combine clock and data prior to transmission are not new. The combination of clock and data ensure that the clock and data signals always arrive at the same time. The trick, however, is the separation of the clock and data at the receiver. This is accomplished by the CDR circuitry. Products that take data from a parallel to a serial format or vice versa are called serializers/deserializers (or "SerDes" for short). These products generally have CDR blocks to deserialize the serial data stream.

This article examines the component blocks of CDR needed to implement a successful CDR in high-speed serial-communication link applications. An overview of a typical high-speed serial-communication link is provided with respect to how data is transformed and recaptured over the link. Different CDR schemes are discussed with respect to the general CDR topology. Also, special attention is given to the reference oscillator's role on both the transmit and receive sides of the link.

## Clock and data recovery in high-speed serial communications

**Figure 1** provides a basic diagram of a high-speed serial-communications link. Parallel data (bits $b_1$,$b_2$,$b_3$..$b_n$) arrives at the transmission serializer at frequency $f_t$. Within the serializer, data is converted from the parallel format to a serial format. A serial-bit stream is developed with a minimum bit rate equal to n x $f_t$, where n is the total number of parallel data bits. The resulting frequency (bit rate) can be higher than $f_t$ depending on whether the data is encoded to meet channel requirements for bit error rate (BER) performance or to provide rich transition content at the receive-side CDR. Reed-Solomon forward-error correction (FEC) and 8B10B encoding are, respectively, each examples of channel encoding or creating rich transition content at the receive CDR. This serial data is then readied for transmission to the channel and sent through to the receiver, finally reaching the deserializer. This basic communications block applies to data whether it is transmitted through fiber, air, or across backplanes.

Timing (clocking) is critical in CDR applications. During system design, the designer determines how to drive data

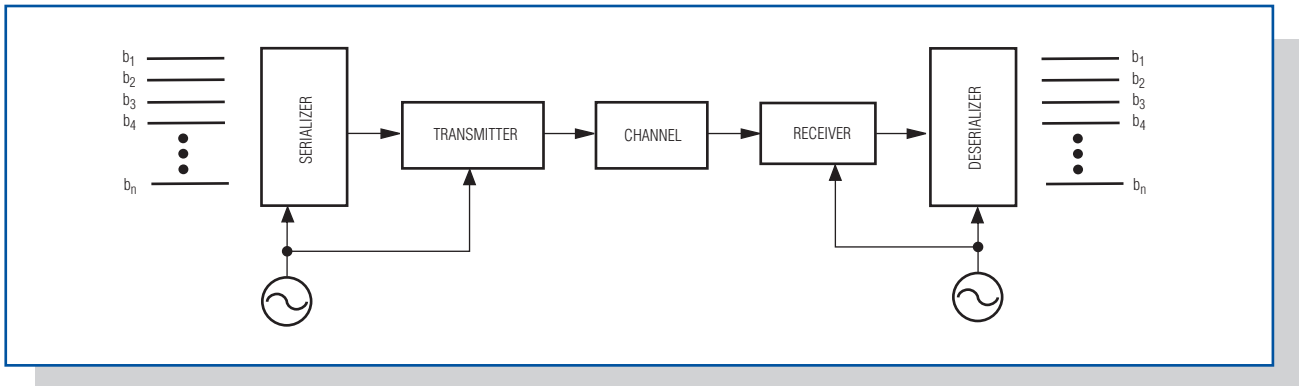


*Figure 1. Clock and data recovery form the basis for high-speed serial communications links.*

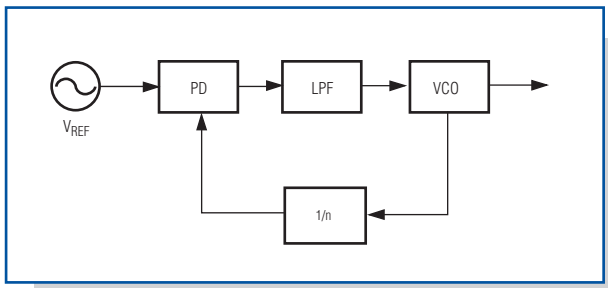


*Figure 2. Clock multiplication applications are driven by PLLs.*

from a parallel format to a serial format for transmission and reception through the channel with pending channel distortion of the transmitted signal. Minimizing the effects of the design's degradation on the data signal is important for protecting signal-to-noise ratio and maintaining BER performance. For example, in digital-transmission schemes across backplanes, the system's jitter performance is extremely important as high-speed electrical signals traverse various lengths (FR-4 and the backplane), thus causing signal degradation in terms of signal-level and time-variant distortions.

At the core of clock-data recovery is a phase-locked loop (PLL)-based circuit that in some cases could be digitally based. **Figure 2** is the basic PLL block diagram that could be used in the serializer or transmission side of the communication link shown in Figure 1. The PLL block includes a phase-frequency detector (PD), filter (LPF), voltage-controlled oscillator (VCO), and a divide chain (1/n). The divide chain is used to provide comparable frequency inputs to the PD. In this manner, the VCO's output is phase-aligned to the very stable reference input, $V_{REF}$. This PLL block's purpose is to multiply the reference frequency by a fixed amount (n), which would be the VCO's natural frequency. In most cases, $V_{REF}$ would be quartz based, providing a high degree of stability and accuracy with great phase-noise characteristics. Additionally, this reference may be temperature compensated or voltage compensated, depending on the required application or system requirements. In SONET-based applications, this reference may meet a certain stratum level (i.e., stratum levels 3, 3E, or 4).

At the receive side, the CDR PLL block takes on a slightly different look to address the need to retrieve both clock and data. As shown in **Figure 3**, the combined clock/data signal comes into the PLL block through a buffer that feeds two distinct paths. One path feeds the data-decision (DEC) block while the second path feeds the clock-recovery blocks. The clock-recovery blocks look very much like Figure 2's PLL block minus the 1/n block. The recovered clock from the VCO is used as a sampling input to the DEC, a feedback to the phase-frequency detector, and also to feed system-timing requirements downstream. In the case of Figure 1, this recovered clock is divided down to the parallel-clock frequency to drive the deserializer block.

## Reference oscillators in clock/data recovery

The reference oscillators shown in the diagrams illustrate a VCO being applied to the input of the oscillator. This voltage control is established by the LPF stage. Typically, VCOs or voltage-controlled crystal oscillators (VCXOs) can be used as the loop oscillator, as shown in Figure 3. Primary to its role, the loop oscillator is required to track the frequency deviation of the incoming clock/data. Additionally, it provides this clock to other components downstream from the CDR (deserializer). This is accomplished by the output of LPF, which drives the voltage-control input of the VCO or VCXO.

In CDR applications for telecom, wireless, and datacom, the incoming data signal plus clock should have relatively stable frequency characteristics. This assumes the transmit clock meets certain accuracy and stability specifications. On the receive side, the design precludes for the minimum and maximum accuracy/stability. In the event the transmit

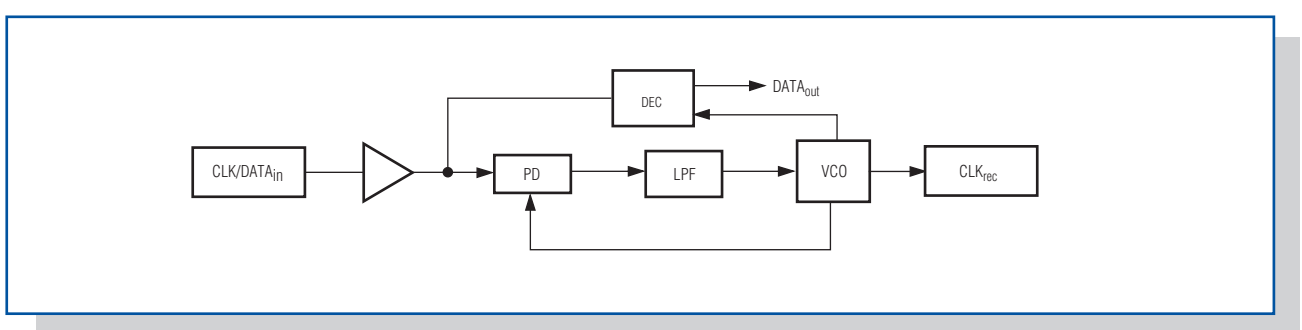


*Figure 3. Modifications to the basic PLL block are used to implement CDR circuits.*

clock frequency is expected to be with ±50ppm of a stated frequency, the receive clock would have frequency-adjustment capability of a minimum ±50ppm. For design purposes, however, consider a frequency adjustment capability somewhat larger than ±50ppm. This extended frequency-adjustment capability accommodates any additional signal-frequency distortion by the channel or disruptions in communications.

Although the PLL attempts to drive toward a static condition—meaning that frequency lock has been established—there could be conditions where the voltage-controlled input may move at higher rates than desired. The LPF bandwidth dictates the maximum rate at which the PLL can maintain lock. Ultimately, the role of the receive VCO (or VCXO) is to track and reproduce the recovered clock.

In the absence of a data/clock input into the CDR, the CDR is required for a specified time to provide a reference signal for any downstream communications requirements (i.e., the deserializer).

In some applications, a combination VCO/VCXO would be used. In **Figure 4**, VCO/VCXO allow for at least two benefits to the common CDR configuration. First, the addition of the VCXO allows fast regulation of the VCO frequency to match that of the expected clock/data signal. The VCXO frequency is chosen to match the expected clock frequency range. For example, a wideband VCO can require thousands of samples to lock onto the incoming data stream. The additions of the VCXO and lock-detection circuit ensure that the VCO maintains a certain operating frequency, and help provide a more predictable lock time in startup conditions. Second, the addition of the VCXO is helpful if the clock/data input has been lost over an extended period of time. In the absence of a clock/data signal, the system would refer to the very stable quartz-based oscillator ($V_{REF}$) to provide holdover until the clock/data signal has been recovered from a loss of signal (LOS). Holdover is a specification that applies to the reference clock's ability to hold a certain accuracy over a certain period of time (for example, ±4.6ppm over 24 hours).

## Conclusion

A variety of solutions address clock/data recovery and retiming, serializers and deserializers, clock generators, and TCXOs for communication applications. These devices allow designers to develop circuits for frequencies ranging from 10MHz up to 10GHz, and support applications ranging from GSM to OC-192 and above. As designs increasingly require greater bandwidth, CDR technology is ideal for telecommunications, optical transceivers, data and storage area networks, and wireless applications.

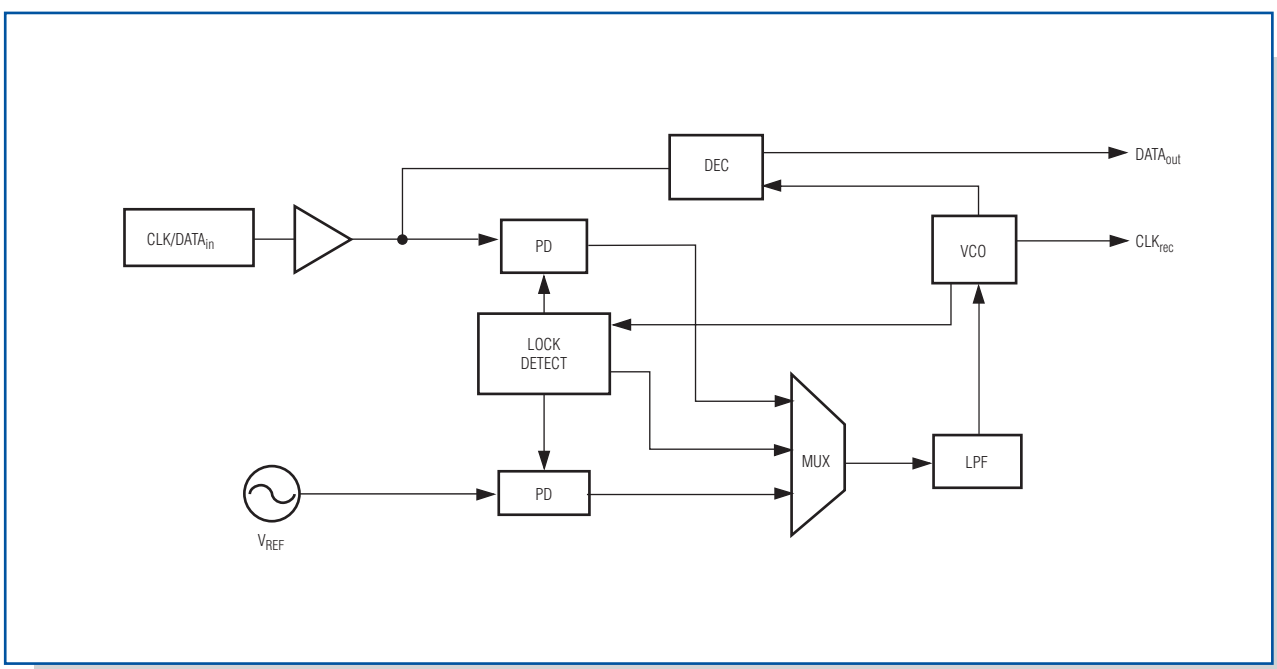


*Figure 4. A modification on the basic CDR block provides easy locking on the incoming clock/data stream.*

# Watchdogs improve system reliability—how to choose the right part

*Many circuit functions previously realized with dedicated hardware are now implemented in software, due in part to today's broad choice of low-cost microprocessors (µPs). While software is often the lowest cost and most flexible way to solve a problem, it forces the designer to take extra measures to ensure system reliability. While there is no such thing as a program without code errors, careful testing can reduce the number of errors to one to ten per 1000 lines of code. Therefore, designers must expect a minimum of 10 code errors in a typical control software program with 10,000 lines of code.*

*Desktop application software errors that cause a system crash are not critical since the user can reboot the system with only a minor loss of data. However, for industrial control software, the system must be able to recover from code errors without human intervention. This feature is critical for two main categories: systems that have high availability, such as servers, telephone systems, and production lines; and systems that must be highly reliable because a crash could lead to injuries, as with automobiles, medical instruments, industrial control, robots, and automatic doors. Even if neither of these criteria apply, system crash/recovery without user intervention (pressing reset or power cycling) is preferred. If a device recovers from an error without human intervention, the perceived quality of this device is good, as the user is unaware that something went wrong inside the device. A simple and effective method of achieving such improved system reliability is to use a watchdog.*

## The watchdog

The watchdog is a counter that must be cleared within the watchdog timeout period. If clearing does not occur, the watchdog generates a reset to cause system reboot or creates a non-maskable interrupt (NMI), causing a program branch to a fault-recovery subroutine. Most watchdogs are edge triggered. Therefore, either a rising or a falling edge on the watchdog input (WDI) will clear the counter. The WDI pin is connected to a processor I/O pin, which is toggled by the software (**Figure 1**). The command to clear the watchdog counter must occur within the main program loop (**Figure 2**). If the watchdog is not cleared, a reset occurs and the software branches to address 0000 (startup routine). Calculating the time it takes to execute the main loop is often difficult, as numerous subroutines might be called, depending on the inputs to the system. Therefore, the designer normally chooses a watchdog timeout that is much longer than the longest measured or calculated loop time. **Figure 3** shows the watchdog and reset signal for normal operation (watchdog is cleared within timeout period). In **Figure 4**, a reset is generated after the watchdog counter reaches the timeout. Industry-standard watchdog circuits have timeouts in the 100ms to 2s range, although there are adjustable and customized watchdogs covering a much wider range (30ms to minutes). If the execution time of the main loop is too long for the watchdog, the designer can implement multiple watchdog-toggle commands within different sections of the main loop or use a device with longer timeout.

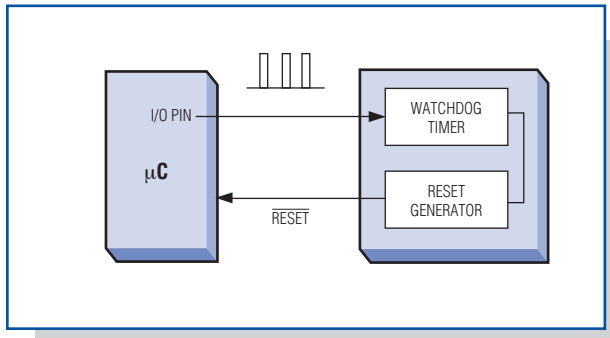


*Figure 1. The microprocessor clears the watchdog timer with a pulse on the WDI pin to prevent a reset.*

A technique that prevents the system from being stuck in a parasitic loop is to set the relevant I/O pin high at the beginning of the main loop, and to set it low in another section of the main loop. If the software gets stuck in a parasitic loop at the start of the main loop, the watchdog times out and the system recovers, as WDI remains high (**Figure 5**). If a low-high-low pulse is used (as in Figure 2), the watchdog will be cleared, but the system will remain stuck. A more sophisticated scheme might be necessary for programs with multiple tasks that require monitoring. Each task sets a flag, and the watchdog is only toggled if all flags are set. The duration of all tasks

must be shorter than the watchdog timeout period. Figures 2 and 5 might seem simplistic compared to actual programs, but they illustrate the relevant concepts. Other potential problems in more complex systems, such as memory leakage and stack overflow, should also be monitored. This is beyond the scope of this article, but is typically done by using suitable design procedures, performing a careful code review, and employing specialized software tools.

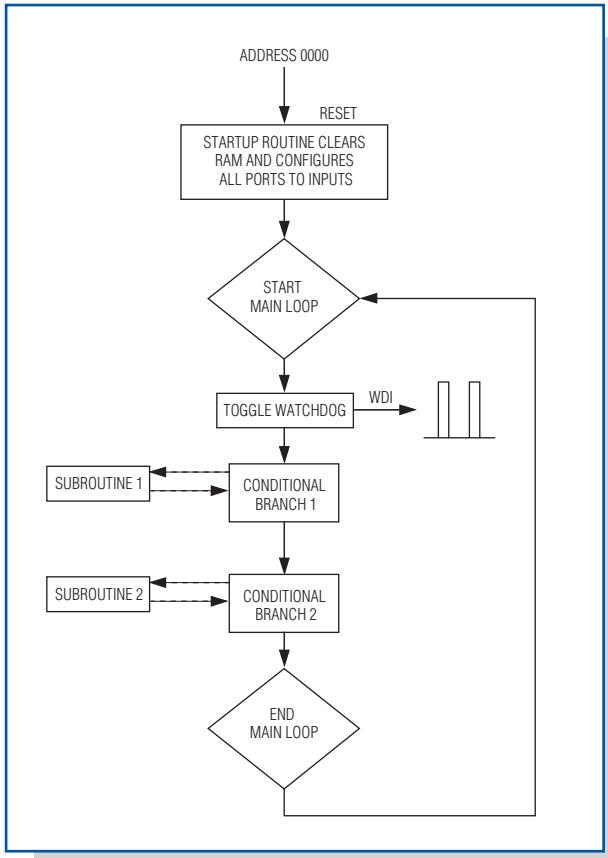


*Figure 2. This figure shows a typical program flow with the WDI signal generated within the main loop.*

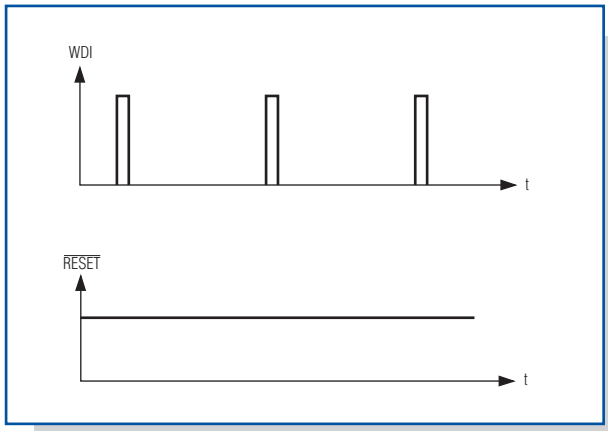


*Figure 3. If the WDI pin is always toggled within the watchdog timeout, no reset is generated.*

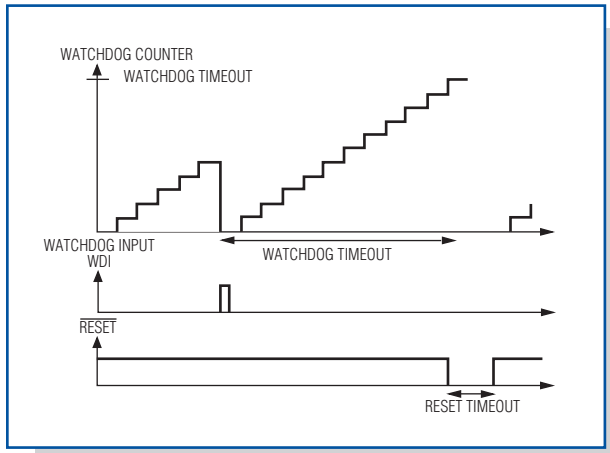


*Figure 4. As soon as the watchdog counter reaches the timeout value, a reset is generated.*

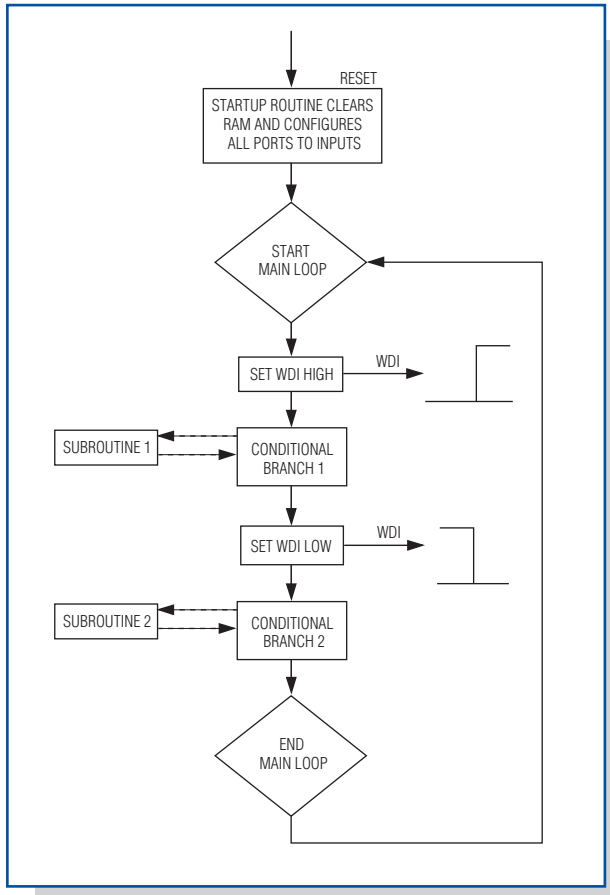


*Figure 5. An improved program flow has two separate watchdog-toggle commands, which generate a rising- and a falling-edge signal on the WDI pin. This prevents the program from being stuck in a parasitic loop.*

## Internal vs. external watchdog

Many µPs have an integrated programmable watchdog that can be disabled under software control. The internal watchdog is prone to code errors, so does not provide the same protection as an independent external watchdog. For safety-critical applications (i.e., automatic doors, medical devices, robots), the internal watchdog is unacceptable. Regulating bodies demand use of a separate, external watchdog. Thus, it is good practice to use an external watchdog to reduce the risk of critical system failures.

## Simple watchdog plus reset

Since a watchdog timeout normally resets the system, most watchdogs are integrated with a µP reset that also monitors the processor supply voltage. The reset is activated either by the watchdog or by an undervoltage condition. The MAX823–MAX825 family shown in **Figure 6** combines these two functions and is available with standard reset voltages, one nominal watchdog, one reset-timeout, and only 6µA current consumption. These devices are available in the ultra-small SC70 package.

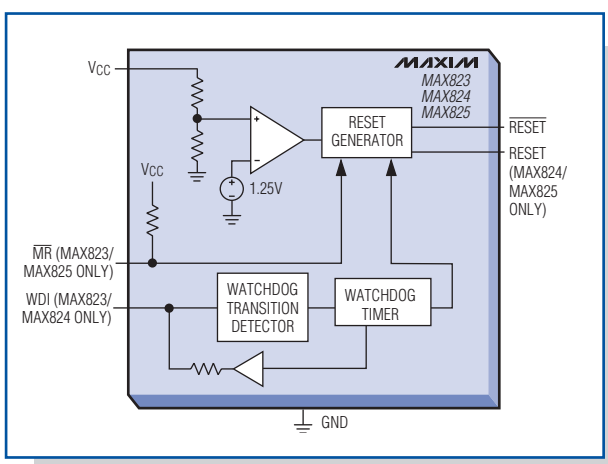


*Figure 6. The MAX823–MAX825 family integrates two popular functions: watchdog and reset.*

## Factory-preset watchdog families

The MAX6316–MAX6322 family offers a choice of 26 factory-preset reset voltages, four nominal-watchdog and four nominal-reset timeouts, and four output configurations (see **Table 1**).

## Capacitor-adjustable watchdogs

If the application requires a flexible watchdog timeout, the designer can use an adjustable circuit. The MAX6746–MAX6753 family offers either factory-preset or voltage-divider-programmable reset voltages, as well as external capacitor adjustment of watchdog and reset timeouts. **Figure 7** shows a typical operating circuit where:

- the reset voltage is determined by the voltage divider R1/R2,
- the reset timeout is determined by the capacitor to set the reset timeout ($C_{SRT}$), and
- the watchdog timeout is set by the capacitor to set the watchdog timeout ($C_{SWT}$).

**Figure 8** shows the watchdog-timeout range for $C_{SWT}$ values from 100pF to 100nF. With this wide range of available watchdog timeouts, the designer has a solution for any application. The MAX6301–MAX6304 family has basically the same features as the MAX6746–MAX6753 family, but is available in SO and DIP packages.

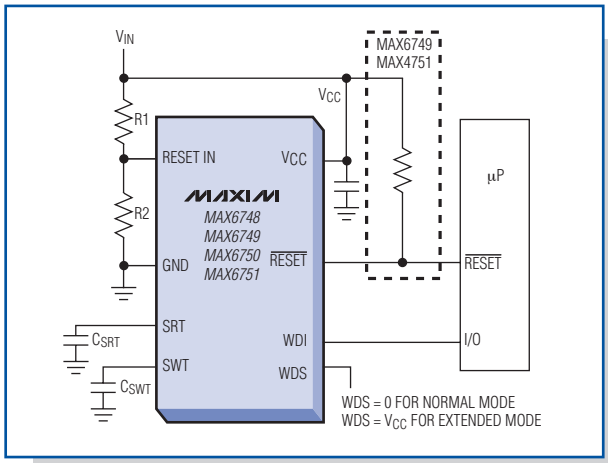


*Figure 7. This figure shows a typical application circuit for the capacitor-adjustable watchdog family MAX6346–MAX6353.*

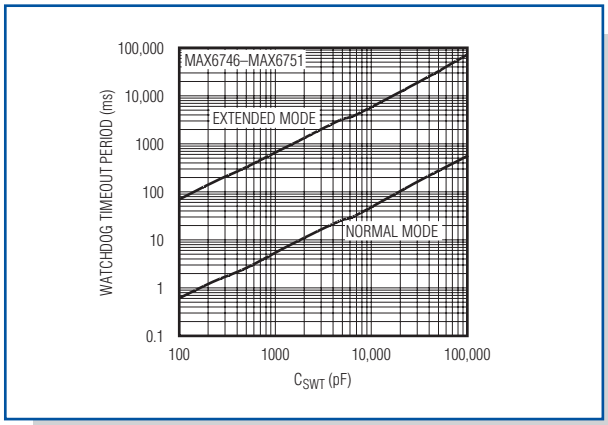


*Figure 8. This figure shows the wide range of available watchdog timeouts.*

## Pin-selectable watchdogs with longer startup/timeout

If the startup routine is long (see Figure 2), a watchdog with two different timeouts is desirable: a longer initial timeout and a shorter timeout for normal operation. The MAX6369–MAX6374 family has a pin-programmable startup delay selectable from 200ms to 60s and a watchdog timeout range of 30ms to 60s. Some versions offer a first-edge activation of the watchdog to provide a solution for even longer startup routines. For these chips, the watchdog is disabled during startup and is activated by the first edge from the relevant I/O pin of the μP.

## Watchdogs with multiple supply voltages

For systems with dual supply voltages, the MAX6358–MAX6360 family can monitor two standard voltages, and offers a watchdog with a long startup as well as a normal timeout. For systems with three supply voltages or that require both active-high and active-low reset functions the designer can use the MAX6721–MAX6729 family. These parts have a dual-mode watchdog with long startup plus normal timeouts. They monitor either two standard supply voltages (MAX6721–MAX6722) or two standard plus a third adjustable supply voltage (MAX6723–MAX6724). These are available with manual-reset input, power-fail comparator, dual reset outputs, and RESET and $\overline{\text{RESET}}$ outputs.

## Windowed watchdogs for ultra-high reliability

For ultra-high reliability, the designer can use the MAX6323/MAX6324 windowed watchdogs. With these parts, the pulse clearing the watchdog must occur within a well-specified time window. A valid pulse may come as early as 1.5ms after the last pulse or could arrive as late as 10ms after the last pulse (see Table 1 for additional ranges). With the MAX6323/MAX6324 the system recovers from parasitic loops, which can generate a fast-pulse train if the clear-watchdog command is within the loop. These pulses would clear a normal watchdog and no reset would be generated. This can be avoided with windowed watchdogs, as they require a minimum delay between watchdog pulses. Typical applications for these devices are anti-lock brake systems or other automotive circuits, industrial and medical applications where high safety requirements apply, or applications where system availability is critical.

## Conclusion

Since every software program has code errors, the designer must ensure that the system does not lock up. Noise and EMI can also affect data in the system and lead to unpredictable system behavior. A watchdog is a simple, inexpensive way to improve system reliability. An external watchdog protects the system from being stuck and resets the μP if WDI is not toggled within the watchdog timeout period. With today's wide choice of watchdogs, the designer is sure to find a device-requirement match.

**Table 1. Typical thermal resistances of MOSFET packages**

| Application | Family | Voltage Monitoring | Watchdog Timeout (min) | Reset Timeout (min) | Special Features |
|---|---|---|---|---|---|
| Simple plus reset | MAX823/ MAX824 | Factory-preset 2.5V, 3.0V, 3.3V, or 5V | 1.12s | 140ms | SOT23 or SC70 packages |
| Customized | MAX6316–MAX6322 | Factory-preset in 100mV steps 2.5V to 5V | 4.3ms, 71ms, 1.12s, 17.9s | 1ms, 20ms, 140ms, 1.12s | Push-pull, open-drain, or bidirectional output |
| Capacitor-adjustable | MAX6746–MAX6753 | Factory-preset, or adjustable by voltage divider 1.575V to 5V | 700ms to 70s in two ranges by 100pF to 100nF capacitor | Preset, or 0.5ms to 5s by capacitor | SOT23-8, min/max windowed option |
| | MAX6301–MAX6304 | | | | SO or DIP packages |
| Long startup, pin-selectable | MAX6369–MAX6374 | Dual factory-preset 1.8V, 2.5V, 3.0V, 3.3V, or 5.0V | 30ms to 60s; 200ms to 60s first-edge activation | Watchdog only | Dual mode, pin-programmable startup delay |
| Multisupply | MAX6369–MAX6360 | Dual fixed 1.8V, 2.5V, 3.0V, 3.3V, 5V; or dual fixed plus one adjustable | 1.6s normal | 100ms | Manual reset, power-fail comparator, dual reset, RESET plus $\overline{\text{RESET}}$ outputs |
| | MAX6721–MAX6767 | | 25.6s startup | | |
| Windowed | MAX6323/ MAX6324 Dual Mode | Factory-preset 2.5V, 3V, 3.3V, or 5V | 1.5ms to 719ms (min); 10ms to 1.3s (max) window | 100ms | Eight factory-trimmed options; timeout reset pulses accepted only within the defined window |

# Designing a pulse charger for use with unregulated power supplies

## Introduction

The *Charging Efficiency vs. Cell Aging in Li+ Pulse Chargers* article in this issue suggests that a current-limited and regulated power supply must be used with a pulse charger like the DS2770. Because the pulse charger does not regulate charging current, the current limit is indeed a requirement. However, by adding a few passive components, the DS2770 can support a variety of current-limited, full-wave rectified, unregulated charge sources.

## Operation

When the charge source is connected to the circuit between the CS and PACK- pads, the DS2770 detects the source and begins to charge the cell (following cell prequalification and trickle charging, if necessary). The voltage applied to the $V_{CH}$ pin must not fall below the battery voltage during charging, or the DS2770 assumes the source has been removed and will terminate the charge prematurely. **Figure 1** shows possible unregulated charge sources and the modified application circuit to prevent this.

Three modifications are required for the application circuit in the DS2770 data sheet to allow for unregulated supplies. Schottky diode $D_1$ replaces a 150Ω resistor to disallow capacitor $C_1$ from discharging through the charge source. Junction diode $D_2$ is added to provide at least 0.5V margin on $V_{CH}$ over $V_{DD}$ and to prevent the battery from discharging through the charge source. Capacitor $C_1$ exists in the regulated application schematic, but must be made larger to maintain the voltage on $V_{CH}$ for the entire period the

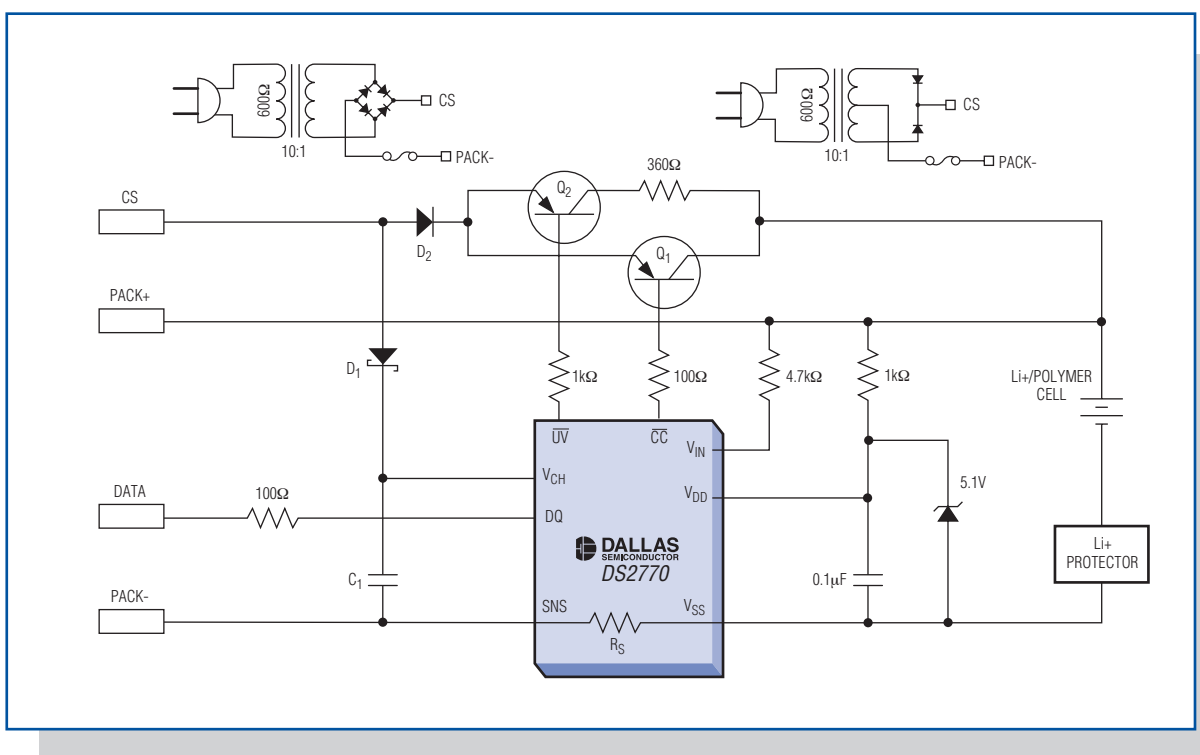


*Figure 1. By adding a few passive components, the DS2770 can support a variety of current-limited, full-wave rectified, unregulated charge sources.*

charge source is lower than the cell voltage. Design criteria for selecting $C_1$ follows.

While the source voltage is higher than the cell voltage, charge flows into the cell and the voltage on $V_{CH}$ is forced above the voltage on $V_{DD}$. While the source is lower than the cell voltage, capacitor $C_1$ maintains $V_{CH}$ at a level above the cell voltage. **Figure 2** shows the relationship between the charge source voltage, $V_{CH}$ pin voltage, and cell voltage during charging.

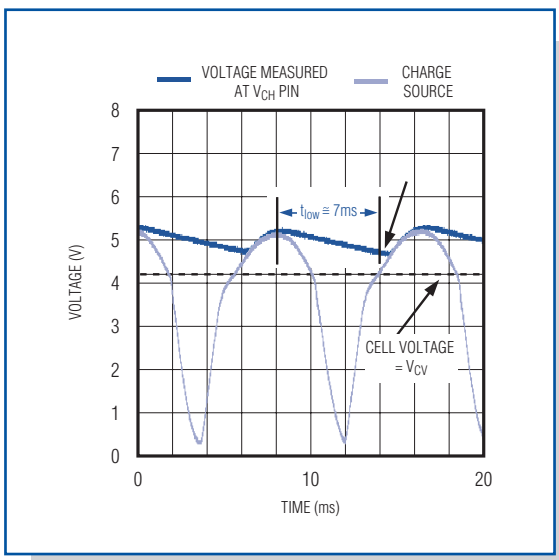


*Figure 2. The voltage on the $V_{CH}$ pin must remain above the maximum cell voltage.*

Like all pulse chargers, the DS2770 has two distinct phases during the fast-charge process: constant current at the charge-source current limit, and pulse-charge top-off. The primary concern addressed with this circuit is in the constant-current phase when transistor $Q_1$ is on continuously. This circuit prevents current from pulsing in this phase where constant current is desired. In the pulse-charge top-off phase, the circuit also maintains charging current while $Q_1$ is on for the pulses, which is approximately 100x the period of a 60Hz full-wave rectified supply (875ms typ for the DS2770). When $Q_1$ is off during the pulsed phase, $C_1$ charges up to the open-circuit voltage of the charge source, which is significantly higher than the cell voltage.

## Capacitor $C_1$ calculation

Capacitor $C_1$ prevents the voltage at the $V_{CH}$ pin from falling below the cell voltage for the duration of the charge source low period (time $t_{low}$ in Figure 2). During this time the 150µA load (data sheet max) through the $V_{CH}$ pin drains $C_1$. The minimum capacitance value required is derived starting from the standard equation:

$$I = C\frac{dV}{dt}$$

where dt is the low period ($t_{low}$), dV is the voltage difference from the $V_{CH}$ pin to the cell voltage at the start of $t_{low}$, and I is the internal load on the $V_{CH}$ pin. To solve for C, the equation can be rewritten as:

$$C_1 \geq I_{VCH}\left(\frac{t_{low}}{V_{D2} + V_{Q1} - V_{D1}}\right)$$

Solving for this example:

$$C_1 \geq 150\mu A \times \left(\frac{7ms}{0.7V + 0.2V - 0.2V}\right) = 1.5\mu F$$

1.5µF is the minimum value for $C_1$, not considering device tolerances. To ensure proper operation for a specific application, use worst-case tolerances for all components including the tolerance of $C_1$ itself and the worst-case timing for the width of $t_{low}$.

## Summary

To be used with lower cost, unregulated power supplies, pulse chargers like the DS2770 need very few passive component modifications to the application schematic for regulated supplies. A Schottky diode resistor replaces the charge source sense pin ($V_{CH}$). Adding a junction diode to the charge path prevents the battery from discharging through the charge source. A larger capacitor on $V_{CH}$ maintains the charge sense voltage at a higher level than the cell voltage for the entire full-wave rectified waveform generated by the charge source. This circuit maintains $V_{CH}$ greater than the cell voltage ($V_{CV}$) during the charge cycle's constant current phase and the pulsed phase.

# How do Dallas Semiconductor delay lines work?

In the early 1980s, Dallas Semiconductor was the first company to develop the all-silicon delay line (**Figure 1**). These ICs provided a smaller, more cost-effective alternative to the modular delay lines used at the time.

The first Dallas delay lines consisted of an RC-based ramp generator and a comparator circuit that transitioned the delay line output when it reached a certain voltage level on the ramp generator. Calibration was performed at the factory at wafer level by using a laser to blow a series of fuses until the desired delay was achieved. There was limited provision for temperature compensation.

The current generation of delay lines is more sophisticated. Maxim/Dallas Semiconductor's all-silicon delay lines now contain a novel circuit consisting of a voltage-controlled delay line (VCDL) and a compensation circuit to reduce delay variations caused by changes in process, temperature, and voltage (**Figure 2**).

Building a silicon delay line is not very difficult. Any logic gate has a propagation delay and can be used as a delay line. The difficulty is building a delay line that can be precisely set to a specific delay period that remains consistent over variations in process, temperature, and voltage. Stabilizing these delay times requires a compensation scheme independent of these parameters.

One can do this by using feedback to determine a delay error, and using that error to generate a corrective input back into the delay line. This requires a method to measure the delay error and to control the delay time. A simple way to control the delay time in a logic gate is to vary the supply voltage. In general, the higher the voltage, the shorter the delay through the gate.

More work is required to measure the delay time and determine the delay error. The simplest way to quantify delay time is to convert it to something more

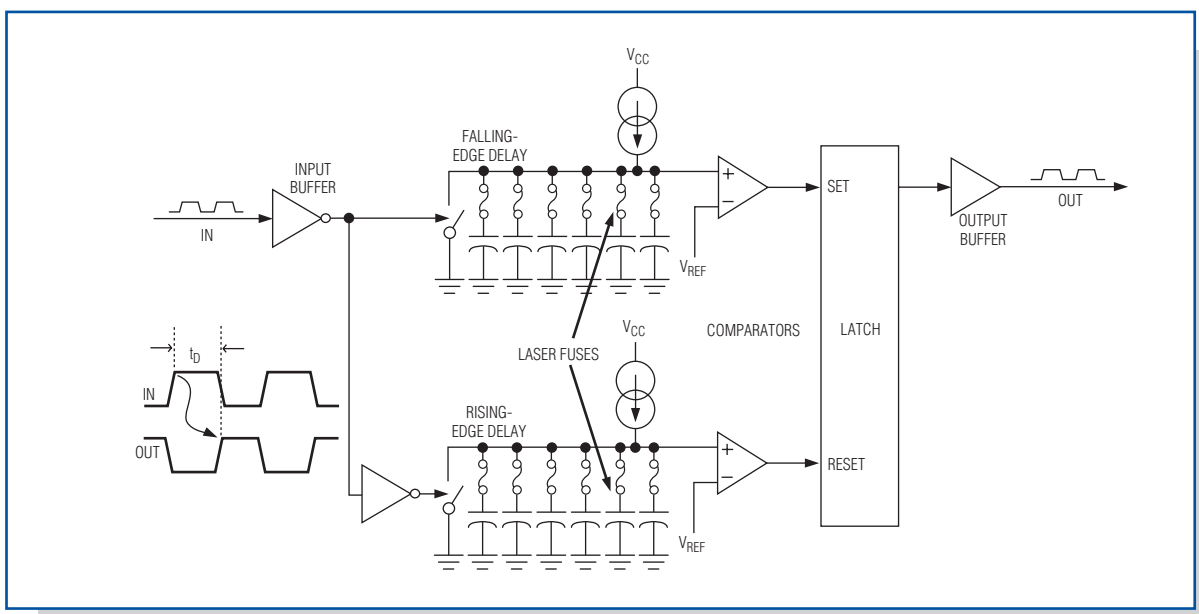


*Figure 1. The old-generation DS1000 delay line used a sample RC circuit to shift the input signal. Then, an output buffer re-drove the shifted signal to "square up" or remove any signal degradation that occurred during the filter/shift process.*

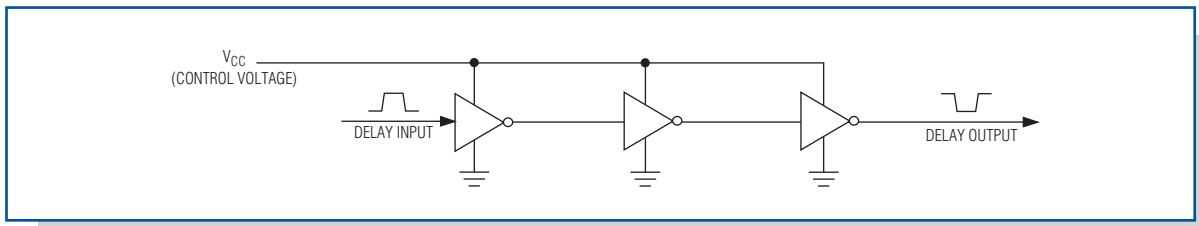


*Figure 2. A VCDL can be thought of as an inverter gate with a propagation delay inversely proportional to its supply voltage.*

easily measured, like frequency. If you take the output from the delay line, invert it, and feed it back into its input, you have an oscillator with a frequency 1/2td, where td is the total delay through the delay lines. In this case, we have a voltage-controlled oscillator (VCO) based on a VCDL (**Figure 3**).

With a precise frequency reference, a phase-locked loop (PLL) can be implemented to lock the VCO frequency to the reference frequency, giving it the same precision. This is called a delay-lock loop (DLL). Other references, such as voltage references, are available in the silicon world that can be calibrated at the time of production to very precise levels.

Given this knowledge, we have all the elements to build a compensated delay line. **Figure 4** shows a block diagram for a DS1135 3-in-1 delay line.

In the circuit (Figure 4), the oscillator output is fed back into a voltage-controlled resistance biased by a fixed-current source. The controlled resistance is actually a switched-capacitor circuit that has a DC resistance inversely proportional to the frequency of the feedback. As the frequency increases, the voltage drops. This voltage is compared to a fixed-voltage reference consisting of a matched current source and a fixed resistance ($R_{REF}$). $R_{FREQ}$ is calibrated to match the characteristics of $R_{REF}$ over temperature and voltage. The output of the comparator is filtered and provides the VCO's drive voltage. As frequency increases, $V_{FREQ}$ decreases with respect to $V_{REF}$, decreasing the drive voltage into the VCO, reducing the frequency. As the frequency decreases, the opposite happens, increasing the frequency. The stability of the frequency is equal to the stability of $R_{REF}$. $R_{REF}$ is a precision reference that is stable over voltage and temperature. The control loop forces $R_{FREQ}$ to equal $R_{REF}$. This control loop also negates changes in the circuit's delay portion caused by variations in process, temperature, and voltage.

But this circuit alone cannot be used as a delay line. It does make a stable oscillator and is the circuit used in Dallas Semiconductor's EconOscillators™. Fortunately, delay cells on the same piece of silicon have nearly identical characteristics to the delay cell used in the VCO (Figure 3). The control voltage fed into the VCDL (configured as a VCO) has an identical effect on these other independent delay cells. So even though they are operating open loop, the control voltage applied to them has the same effect as on the delay cell configured in the VCO. This provides compensation for changes caused by process, temperature, and voltage.

These individual delay cells can be daisy-chained together to make a tapped delay line like the DS1100, or they can be used independently as they are in the

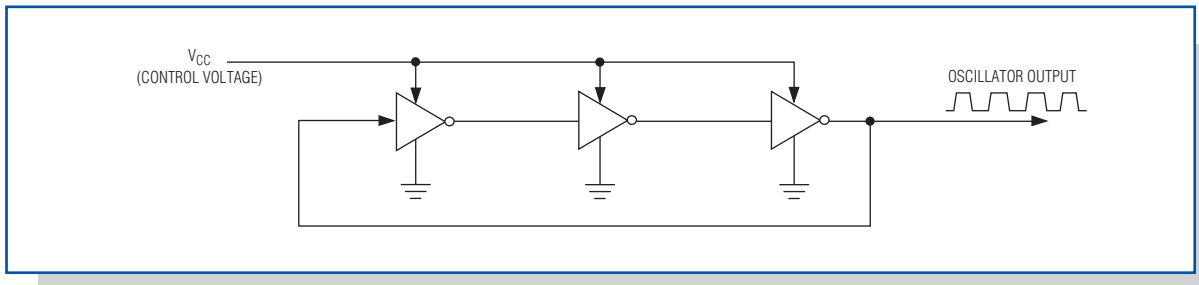


*Figure 3. VCO based on a VCDL. A VCDL can be converted into a voltage-controlled oscillator by inverting the outputs (relative to the input) and feeding this signal back into the input.*

EconOscillator is a trademark of Dallas Semiconductor, Inc.

DS1135. The DS1075 and DS1077 EconOscillators use the oscillator section only in conjunction with a programmable divider chain to give a customer-configurable all-silicon oscillator. Combine this circuit with a DAC controlling the VCO's control voltage and a full-scale frequency synthesizer like the DS1085 is produced.

The circuitry described above is patented by Dallas Semiconductor (U.S. patent 5,982,241, as well as other patents).

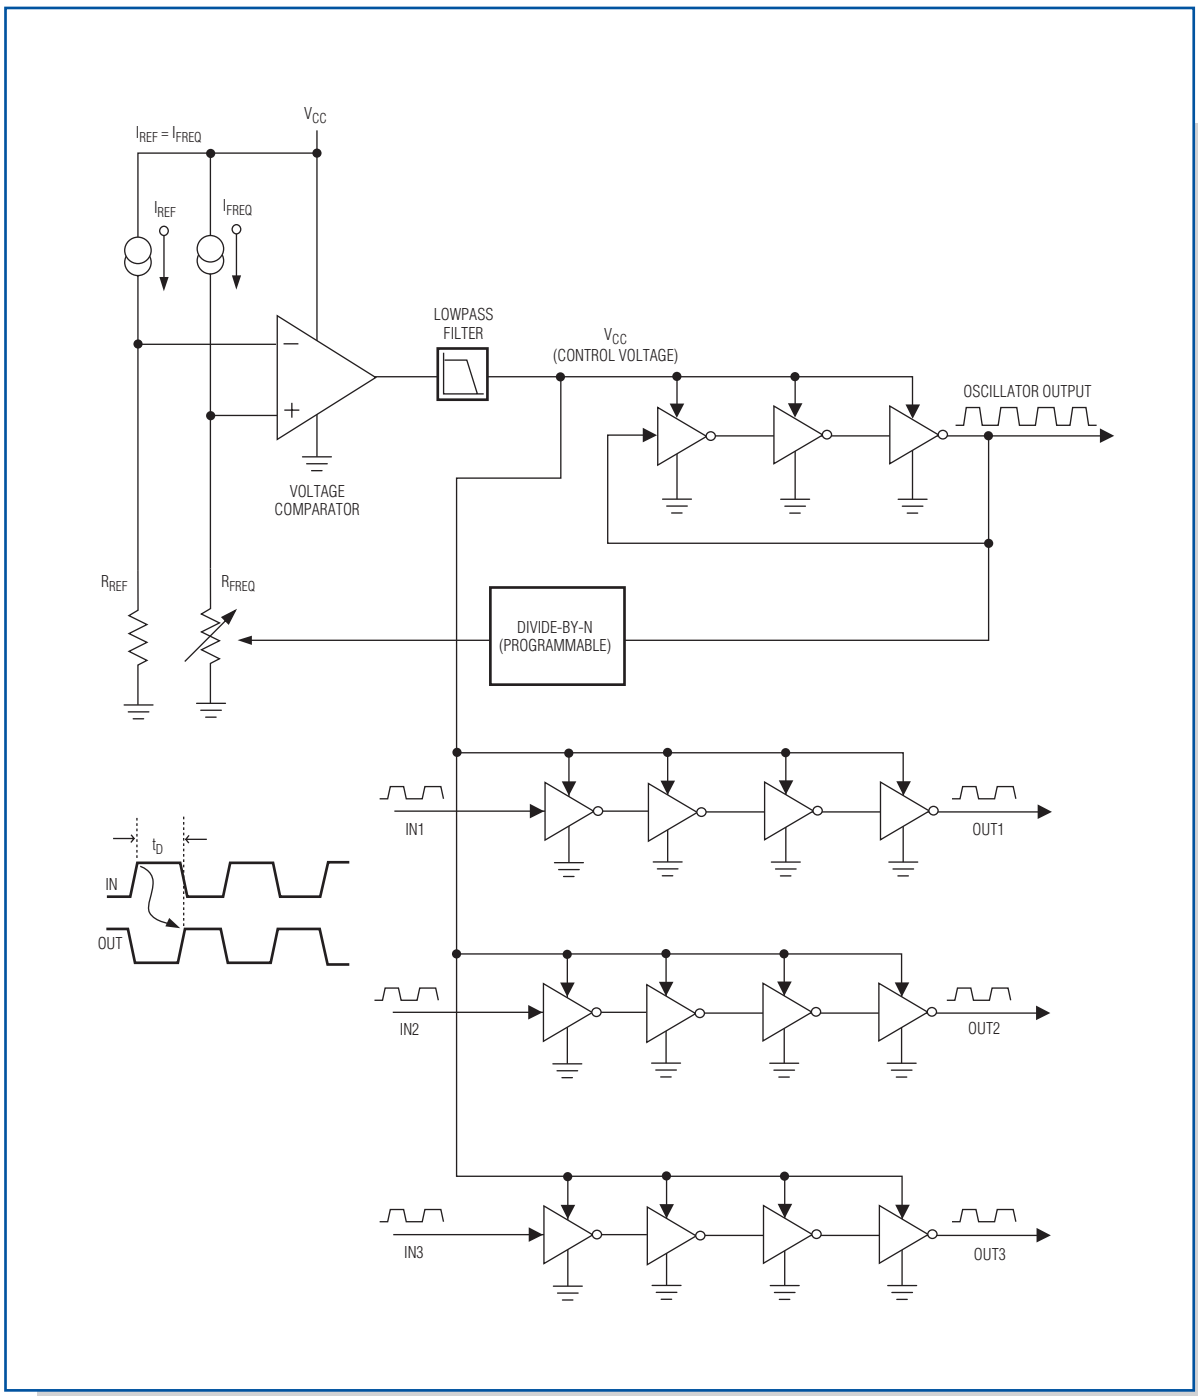


*Figure 4. Delay line (DS1135) based on delay-lock loop technology. Once the individual open-loop delay lines have been calibrated using the closed-loop reference as a master, they can be configured as independent, tapped, and other delay lines.*