Keywords: MAX6653, MAX6660, MAX6620, MAX6639, fan control, cooling, high-speed chips

APPLICATION NOTE 3173

Automatic Fan Control Techniques: Trends in Cooling High-Speed Chips

By: Kerry Lacanette
Jul 08, 2004

Abstract: Cooling fans are an important part of thermal management for high-power chips (such as CPUs, FPGAs, and GPUs) and systems. Unfortunately, their use can sometimes raise a system's acoustic noise level to the point where it is objectionable to the user. By measuring temperature and adjusting fan speed accordingly, the fan's speed (and noise level) can be minimized when temperature is low, but increased under worst-case conditions to prevent damage. This article describes two techniques for automatically controlling a cooling fan's speed.

High-speed chips tend to run hot. As they get faster, they get hotter. New generation high-speed digital chips use smaller processes that allow the supply voltage to be reduced, which helps somewhat, but the number of transistors increases faster than the supply voltage decreases. Power levels, therefore, continue to rise.

As chip temperature increases, performance can suffer. Parameters shift, maximum operating frequencies decrease, and timing can fall out of specification. From the user's point of view, the product is no longer operating properly when this occurs. The first reason for cooling high-speed chips, therefore, is to maintain good performance for the longest possible operating time and over the widest possible range of environmental conditions. The maximum allowable temperature for a high-speed chip to meet its parametric specifications depends on the process and how the chip is designed (how "close to the edge" the chip is operating), among other factors. Typical maximum die temperature values range from +90°C to +130°C.

Beyond the point where performance degradation begins, excessive die temperature causes catastrophic damage to chips. The maximum die temperature limit is usually well over +120°C and depends on such factors as process, package, and duration of high-temperature conditions. High-speed chips are, therefore, cooled to avoid reaching a temperature that could both degrade performance and cause irreparable damage.

A single cooling technique is rarely used with high-speed chips. Instead, combinations of techniques are generally necessary to ensure high performance and continued reliability. Heat sinks, heat pipes, fans, and clock throttling are commonly employed to cool high-speed chips. The last two, fans and clock throttling, can help solve the heat problem, but introduce problems of their own.

Fans can dramatically reduce the temperature of a high-speed chip, but they also generate a great deal of acoustic noise. The noise from a full-speed cooling fan is annoying to some consumers and is also becoming a target of government agencies concerned about the long-term effects of noise in the workplace. Fan noise can be reduced significantly by varying the fan's speed based on temperature; the fan can turn slowly (and very quietly) when temperature is low, and can speed up as temperature increases.

Clock throttling—reducing clock speed to reduce power dissipation—works by reducing system performance. When throttling the clock, the system continues to function, but at a reduced speed. Clearly, in high-performance systems, throttling should be done only when it is absolutely necessary—that is, when the temperature reaches the point where functionality is about to be lost.

Controlling fan speed or clock throttling based on temperature requires that the temperature of the high-speed chip is measured first. This can be done by placing a temperature sensor close to the target chip—either directly next to it or, in some cases, under it or on the heat sink. The temperature measured this way corresponds to that of the high-speed chip, but can be significantly lower (up to around 30°C), and the difference between measured temperature and die temperature increases as the power dissipation increases. Therefore, the temperature of the circuit board or heat sink must be correlated to the die temperature of the high-speed chip.

A better alternative is possible with a number of high-speed chips. Many CPUs, graphics chips, FPGAs, and other high-speed ICs include a "thermal diode", which is actually a diode-connected bipolar transistor, on the die. Using a remote-diode temperature sensor connected to this thermal diode, the temperature of the high-speed IC's die can be measured directly with excellent accuracy. This not only eliminates the large temperature gradients involved in measuring temperature outside the target IC's package, but it also eliminates the long thermal time constants, from several seconds to minutes, that cause delays in responding to die temperature changes.

The need for fan control forces the designer to make several key choices. The first choice is the method of adjusting the fan's speed. A common method of adjusting the speed of a brushless DC fan is to regulate the power-supply voltage of the fan. This approach works well for power-supply voltages as low as about 40% of the nominal value. However, there is a drawback. If the power-supply voltage is varied using a linear pass device, the pass device will dissipate power, thus reducing efficiency. The worst-case pass device dissipation will occur at about 50% of full speed, and will be approximately 25% of the fan's full-speed power dissipation. Better efficiency can be obtained using a switch-mode power supply for the fan, but this increases cost and component count.

Another popular fan-speed control technique is to power the fan with a low-frequency PWM signal, usually in the range of about 30Hz, whose duty cycle is varied to adjust the fan's speed. This is inexpensive because a single, small pass transistor can be used. It is efficient because the pass transistor is used as a switch. A disadvantage of this approach, however, is that it can make the fan somewhat noisier because of the pulsed nature of the power supply. The PWM waveform's fast edges cause the fan's mechanical structure to move (somewhat like a badly designed loudspeaker), which can easily be audible.

A third approach can be used with fans that have speed-control inputs. Such fans, usually referred to as "four-wire" fans (ground, power, tachometer output, PWM speed-control input), allow a logic-level PWM signal to control their speed. This PWM signal is usually in the 20kHz to 50kHz range, and the resulting fan speed is roughly proportional to the duty cycle.

Another fan-control design choice is whether the fan's speed is measured as part of the control scheme. In addition to power and ground, many fans are available with a third wire that provides a "tachometer" signal to the fan-control circuitry. The tachometer output produces a specified number of pulses (two pulses, for example) for each revolution of the fan. Some fan-control circuits use this tachometer waveform as a feedback signal that allows the fan's voltage or PWM duty cycle to be adjusted to give a desired RPM. A simpler approach ignores any tachometer signal and simply adjusts the fan's drive to speed up or slow down with no speed feedback. Speed control using this method is less precise, but cost is lower and at least one feedback loop is removed, simplifying the control system.

In some systems, it is important to limit the change rate of the fan speed. This is most critical when the system is in close proximity to users. Simply switching a fan on and off or changing speed immediately as temperature changes is acceptable in some environments. When users are nearby, however, sudden changes in fan noise are apparent and annoying. Limiting the rate of change of the fan's drive signal to an acceptable value (e.g., 1% per second) ensures that the acoustic effects of fan control are minimized. The fan speed still changes, but it does so without attracting attention.

The fan-control profile is another important design variable. Typically, the fan is off below a specific threshold temperature and then begins to spin at a slow rate (for example, 40% of full speed) once the threshold is exceeded. As temperature increases, the fan's drive increases linearly with temperature until it reaches 100% drive. The best slope depends on system requirements. A more rapid slope results in somewhat more consistent chip temperature, but fan speed has more variation as power dissipation changes from one moment to the next. If highest performance is the goal, the starting temperature and the slope should be chosen so that the fan reaches full speed before the die temperature is high enough to initiate clock throttling.

Implementing fan-control circuitry can be done in several ways. A variety of remote temperature sensors with up to five sensing channels is available that can detect the die temperature of the high-speed chip and transmit temperature data to a microcontroller. Fan-speed regulators with multiple channels of fantachometer monitoring can provide reliable control of fan RPM or supply voltage based on commands from an external microcontroller. For low cost and simple implementation, ICs are available with temperature sensing and automatic fan control included in a single package. Sensor/controllers also normally include overtemperature detection for clock throttling and system shutdown, thereby protecting the high-speed chips from catastrophic failure due to overheating.

Examples of some representative fan-speed control ICs are shown in **Figures 1**, **2**, and **3**. The MAX6620 in Figure 1 controls fan speed by generating DC supply voltages for up to four fans using external pass transistors. The MAX6620 uses tachometer feedback to force the fans' speeds to the selected values. In Figure 2, the MAX6653 controls a single fan by modulating the fan's power supply with a low-frequency PWM signal. The MAX6653 also senses temperature and adjusts the PWM duty cycle based on this temperature. Figure 3 shows the MAX6639, which measures temperature and controls up to two fans based on temperature. The MAX6639's PWM output frequency can be as high as 25kHz, so it can control a 4-wire fan as shown in the figure. The MAX6639 controls the fan's RPM as a function of measured temperature.
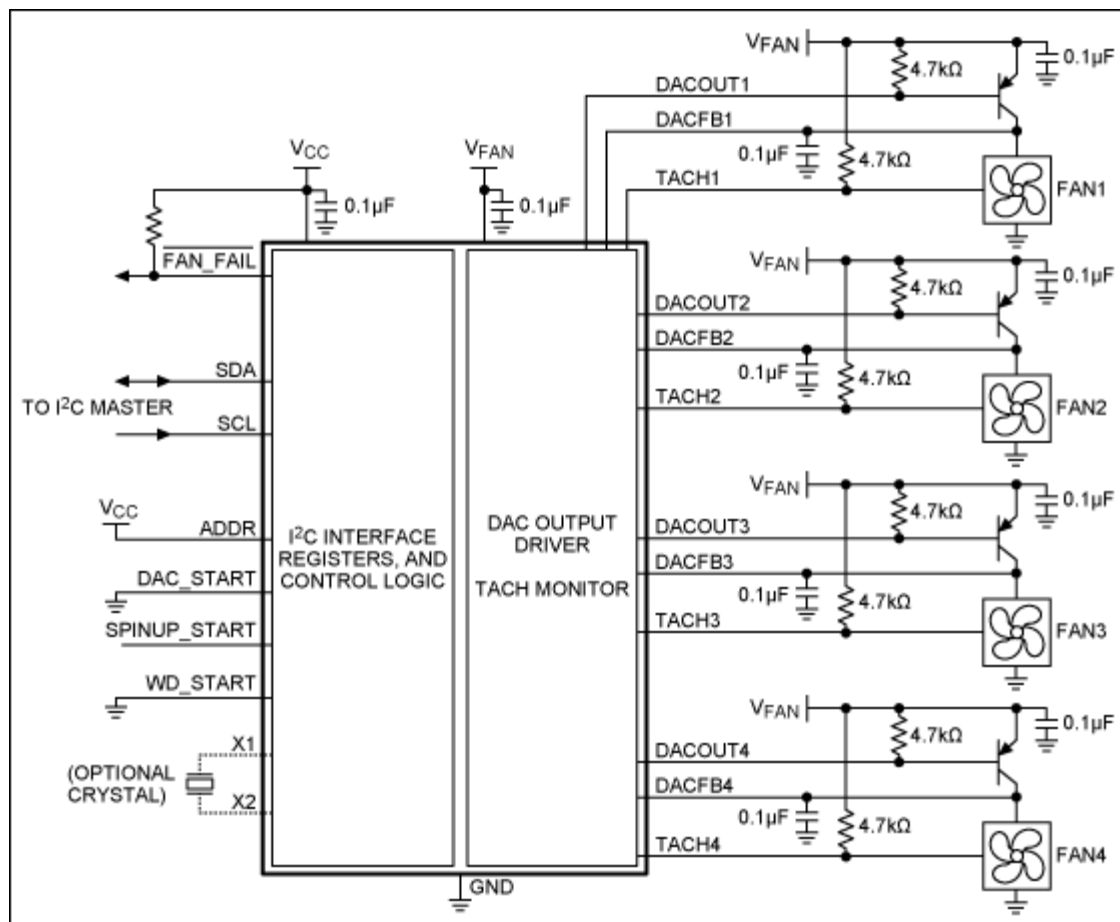
Figure 1. The MAX6620 shown here is a four-channel linear fan -speed controller using external pass transistors to supply variable power to the fan. The MAX6620 monitors tachometer signals from the fans and adjusts the power supply voltage so that the desired tachometer frequency is achieved.
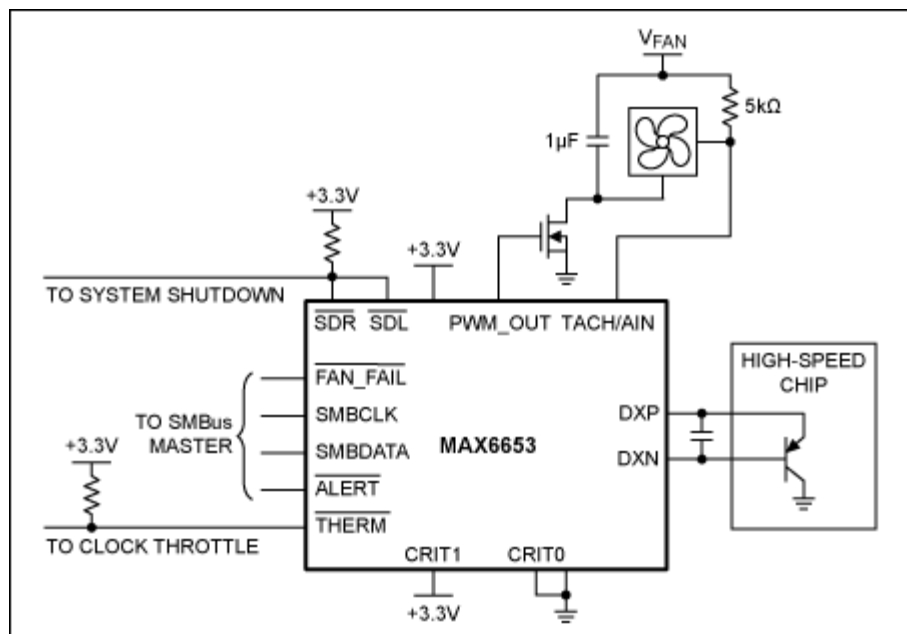
Figure 2. PWM-output temperature sensor and automatic fan-speed controller. Fan speed is controlled automatically based on temperature. Clock throttle and system shutdown outputs prevent a high-speed chip from reaching destructive temperatures. CRIT0 and CRIT1 pins can be strapped to supply or ground to select default shutdown-temperature thresholds, ensuring protection even when system software hangs.
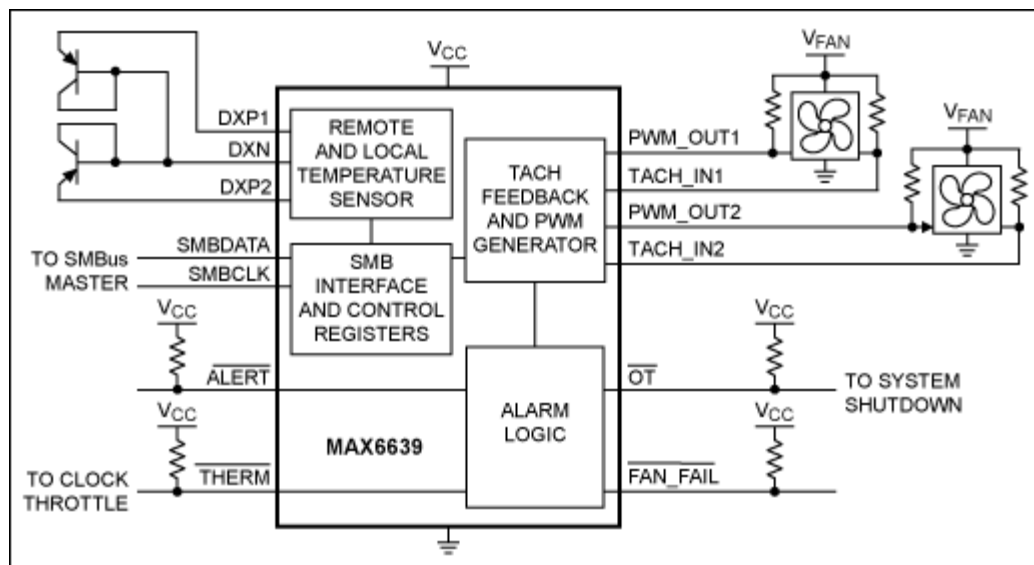

Figure 3. MAX6639 controlling two 4-wire fans. When controlling 4-wire fans, the PWM outputs have a 25kHz output frequency. The MAX6639 measures two temperatures and can control fan RPM based on these measured temperatures.

Related Parts		
MAX6653	Temperature Monitors and PWM Fan Controllers	Free Samples
MAX6660	Remote-Junction Temperature-Controlled Fan-Speed Regulator with SMBus Interface	Free Samples

More Information

For Technical Support: http://www.maximintegrated.com/support
For Samples: http://www.maximintegrated.com/samples
Other Questions and Comments: http://www.maximintegrated.com/contact

Application Note 3173: http://www.maximintegrated.com/an3173
APPLICATION NOTE 3173, AN3173, AN 3173, APP3173, Appnote3173, Appnote 3173
Copyright © by Maxim Integrated Products
Additional Legal Notices: http://www.maximintegrated.com/legal