Real-Time Control with C2000™

The Challenges of High-Frequency Power Electronics

By Matthew Watson for Mouser Electronics

(Source: Texas Instruments)

Though wide-bandgap semiconductors have existed for decades, the engineering community has only recently started leveraging the advantages offered by these materials. One such material is Silicon Carbide (SiC). Compared to Si, SiC has a 3x larger bandgap, 10x higher thermal conductivity, and lower RDS_{ON}, resulting in smaller and more efficient power converters. The improved thermal characteristics also allow for smaller heat sinks, which reduce the overall weight of SiC power management systems—a feature especially beneficial in electric vehicles (EVs) and other mobile electric systems.

However, SiC suffers from one particular drawback: The material achieves its greatest efficiencies when it operates at high frequencies (from 100s of kHz to as high as 2MHz). Therefore, creating systems capable of responding to sudden changes experienced by SiC—such as output shorting and brownouts—can be extremely difficult with standard off-the-shelf microcontrollers.

Worse, controlling the rise and fall times of SiC devices at these frequencies is virtually impossible with off-the-shelf microcontrollers, which lack sufficient speed to provide the required resolution. Even microcontrollers clocked at extremely high frequencies will struggle, as their hardware is not designed around real-time control. In response to a hardware and/or software event, the microcontroller can skew for a few clock cycles, generating incorrect rise and fall times on a SiC device, which in turn causes the SiC device to operate in its linear region for an extended period.

In addition, peripherals commonly found on microcontrollers often require buffering or waiting and thus cannot be used in real-time applications. For example, an analog-to-digital converter (ADC) reading will see a result stored in a register, which first must be checked by the central processing unit (CPU) to see if the result is ready; then, a second instruction is required to move the result into random access memory (RAM) or a CPU register. As such, hardware peripherals typically involve some level of latency unacceptable for real-time control of high-frequency systems operating in the nanosecond regime.

The C2000 Portfolio: Low-Latency Real-Time Controllers

Recognizing the challenges faced by real-time applications requiring low latencies, Texas Instruments has developed the C2000 real-time MCU portfolio (**Figure 1**), which offers engineers genuine real-time performance between peripherals, signal processors, and output controls.

To achieve this performance, the **C2000 MCUs** integrate processor cores, accelerators, and smaller secondary accelerators that enable specific low-latency tasks to operate independently. The main processor, the C28x, is a 32-bit DSP (Digital Signal Processor) featuring numerous accelerators, including a 32×32-bit multiplier, an atomic arithmetic logic unit (ALU), a 32- and 64-bit floating-point unit (FPU), a Viterbi, Complex Math, and CRC Unit (VCU), and timers; all of which are ideal for running generic applications (i.e., fully programmable).

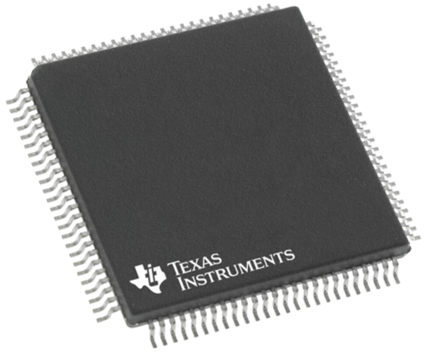

Figure 1: *TMS320F280x, TMS320C280x, TMS320F2801x DSPs. (Source: Mouser Electronics)*

Compared to the main C28x CPU, the secondary accelerators are simpler in operation but allow for high-speed, low-latency signal processing. For example, configurable logic blocks (CLBs) provide the C2000 MCUs with some level of complex programmable logic device (CPLD) capability by performing custom logic functions depending on the state of signals from peripherals. In contrast, the control law accelerator (CLA) is a 32-bit floating-point (math processor) programmable CPU that can read samples from an ADC in real time, process the data, and produce a result.

These secondary processing elements can be used directly with peripherals for both input and output (I/O), meaning they can operate independently from the main processor. Therefore, time-critical tasks requiring minimal latency can be distributed among the C28x DSP, the CLA, and CLB accelerators, allowing parallel computation, distributed application execution to improve the overall signal chain performance of the system.

Because these secondary accelerators can read peripherals at core speed, no buffering or register transfers are needed (i.e., data is processed as it is read). Thus, the C2000 MCUs can respond to signals in real time with zero latency; multiple cores and accelerators help distribute complex tasks across the device. As an element of comparison, measuring the latency of ADC reading, CPU computation, PWM update for a typical real-time control application, a C2000 real-time MCU running at 200MHz delivers the performance equivalent as a Cortex M7-based device running at 280MHz (**real-time signal benchmark document**).

Real-Time Applications Enabled by the C2000 Portfolio

Example 1: Real-time Control for Power Conversion

As previously discussed, wide-bandgap devices find their maximum efficiency when they operate at high frequencies (up to 2MHz). However, such devices should spend as little time as possible in their linear state (during which their channel is partly conducting) to avoid dissipating a large amount of heat (i.e., wasted energy). Correctly driving power devices with the appropriate rise and fall times along with nominal voltage levels is also important.

Figure 2: *Servers and data centers can reduce their energy usage and improve their compute density by power designs based on C2000 devices. (Source: Yanawut Suntornkij/stock.adobe.com)*

In such applications, the C2000 MCUs enable engineers to create systems not only capable of driving such devices but also making real-time changes to gate drivers for the greatest possible efficiency. For example, the C2000 MCUs can be used for detecting zero crossing to help minimize energy loss and for monitoring the output current to identify potential errors.

However, if operated too fast, wide-bandgap devices can suffer from electromagnetic interference and ringing, both difficult to efficiently eliminate. The C2000 MCU can monitor the output stage of a power device, detect ringing, and adjust the gate drivers, thereby maximizing power efficiency while improving electromagnetic compatibility characteristics.

Power designs based on the C2000 devices would be ideal for applications in which energy efficiency is essential. For example, EVs can extend their range by reducing the amount of power wasted, and servers and data centers can significantly reduce their energy usage and improve their compute density (**Figure 2**).

Example 2: Real-time Control of Stepper Motors

Another potential application for C2000 is in stepper motor control applications (**Figure 3**). Stepper motors are often found in precision engineering applications such as robotic arms, computer numerical control systems, semiconductor steppers, and pick-and-place production facilities, and such applications almost always require closed-loop motor controls. The ability of C2000 to control power drivers for stepper motor coils would not only boost energy efficiency but also help increase motor life.

Additionally, the C2000 could reduce the noise (harmonics and ringing) generated by stepper motors during operation. Silent operation of stepper motors is ideal for operators nearby and improves mechanical performance by providing more torque while reducing the stress induced on the stepper motor shaft. Smoother operation reduces the chance of missing steps, which can be detrimental to precision.

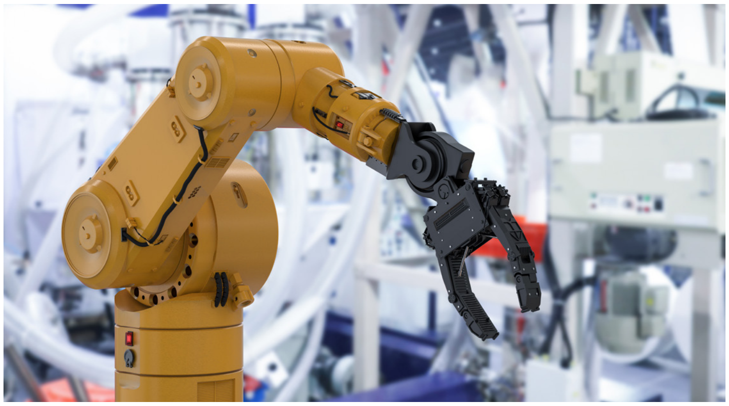

Figure 3: *Stepper motors can be found in precision engineering applications such as robotic arms. (Source: phonlamaiphoto/ stock.adobe.com)*

Lastly, with C2000 monitoring the current from each coil, the torque generated by the motor can be determined. If the motor stalls, the C2000 can immediately reduce the current going to the coils, reverse the motor, or outright stop the current operation.

Example 3: Motors in Electric Vehicles

The low-latency and real-time control abilities of the C2000 portfolio make it an ideal choice for EVs. To improve EV range, by far the best approach is to reduce any unnecessary weight. While the battery is one of the heaviest components on a modern EV, the motor is also a major contributor to weight (**Figure 4**).

With the C2000 as a motor controller in EVs, the motor can be dynamically controlled in real time to always provide the best operating current. As such, the C2000 could allow an EV to use a smaller motor to provide the torque needed to drive the EV at a reduced weight, thereby increasing driving range.

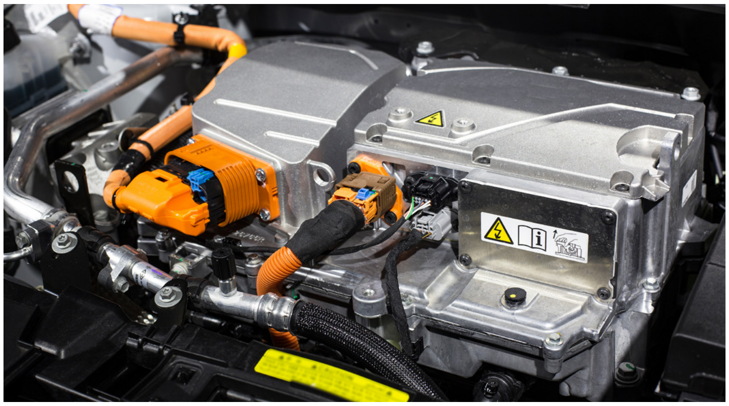

Figure 4: *Electric vehicle motor. (Source: bong/stock.adobe.com)*

Another benefit of using the C2000 in EV motor applications is that voltage and current spikes can cause damage over time, eventually degrading performance. Damage that persists for too long can require the motors to be outright replaced, which comes with a hefty price tag. The ability of the C2000 to control power devices in real time also introduces the possibility of more efficient regenerative braking systems capable of recharging onboard batteries while reducing wear and tear caused by charge cycles.

Example 4: Intelligent Operation and Diagnostics

While most of the examples so far have focused on motors and wide-bandgap semiconductors, the C2000 is also ideal for predictive maintenance applications (**Figure 5**). With the power of the C28x DSP core, the C2000 MCUs can execute on-device neural nets for predictive diagnostics. The low-latency peripheral design of the C2000 and the ability of its CLA CPU and the C28x processor to work together can create powerful edge-computing devices capable of processing inputs in real time.

Consider a two-tier system in which a neural net runs on the C28x and a danger monitor runs on the CLA. The C2000 portfolio can learn from an ongoing process to detect anomalies, while the CLA can compare the current state of a circuit and fire an immediate halt if it detects a dangerous condition. Thus, with a suite of processors and accelerators, the C2000 can handle multiple time-critical jobs simultaneously—a helpful capability for crafting catastrophic event–avoidance systems.

The Future of C2000 Applications

Energy efficiency will become one of the most important factors in future designs. The increasing cost of energy—combined with the growing market for portable electronic devices, the rise of renewables, and the move toward EVs—will force engineers to find every source of energy waste in their designs. Even a humble one-percent increase in efficiency can massively reduce an energy bill, especially for a data center or grid network.

At the same time, reliability will continue to be an important issue—one challenged by the shrinking nature of electronics. With more transistors packed into circuits than ever before, the chances of malfunction are increasingly more likely. Devices used in safety applications must be able to recover from errors and resume normal operation.

The C2000 portfolio of real-time MCUs from Texas Instruments provides engineers with a platform that can not only support numerous programming modes but also allows for designs to respond to changing signals in real time. Furthermore, with the ability to run on-chip accelerators independently from the main CPU and interface with I/O at core speed, engineers can create extremely advanced systems able to tackle even the most complex tasks in high-powered electronics. As SiC and other wideband-gap semiconductors like gallium nitride begin to dominate the market, such capabilities will become essential.

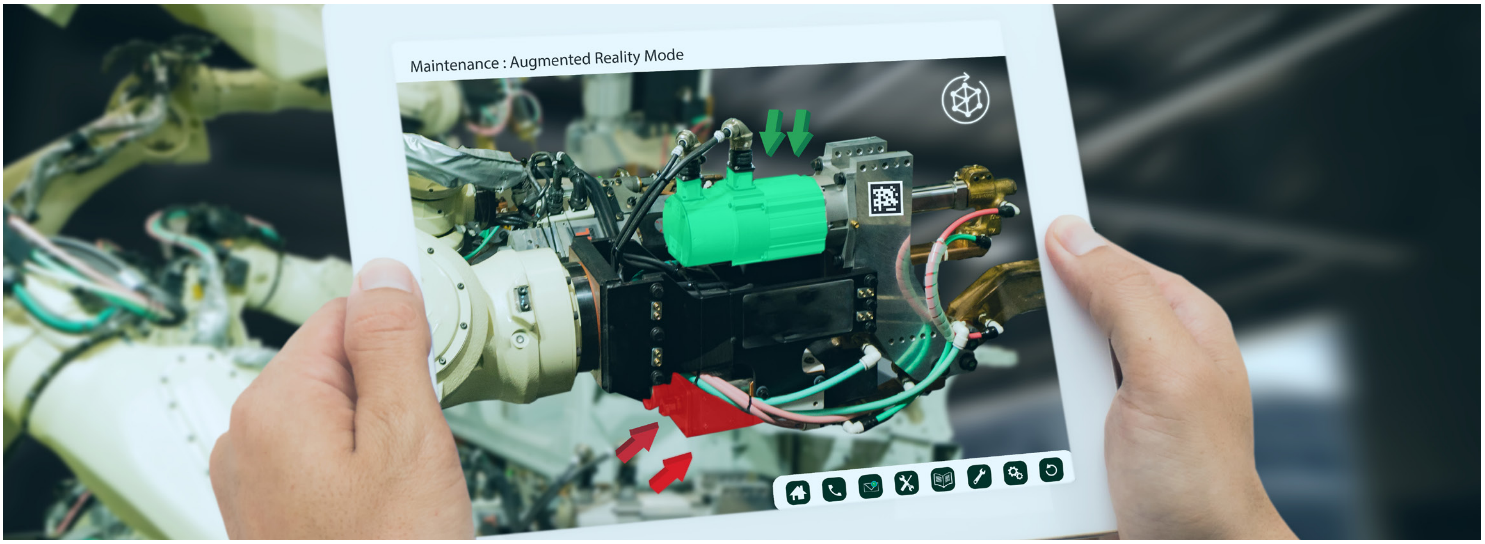

Figure 5: *The C2000 is ideal for predictive maintenance applications. (Source: Monopoly919/stock.adobe.com)*