

Industrial-Strength Features Make a Difference in Embedded Processing

David Katz and Rick Gentile, Analog Devices, Inc.



Many general-purpose signal processors are acceptable for industrial applications based on their computing performance and baseline connectivity features. However, there are important peripheral enhancements that can significantly upgrade the processor's capabilities for more demanding industrial systems. This article will examine two examples of these enhancements, in the realms of networking and motor control.

Ethernet interface

For traditional industrial applications, an Ethernet controller provides for basic network connectivity. The controller (MAC) is usually available on the same chip as the processor. This is usually coupled with an external PHY chip to complete the interface.

External MAC/PHY chips are also available, and they can often be connected directly to an asynchronous memory interface on the processor. While Ethernet MAC/PHY combination chips have



continued to decrease in price to the point where they are almost at cost parity to a standalone PHY chip, they generally can't compete with the transfer rate of an integrated MAC/external PHY solution. This is due to the fact that an internal MAC is usually tied to a system DMA channel, which can be set up to either transmit and receive data with minimal core processor

interaction. The internal MAC controller can typically achieve close to line speed, depending on the protocol.

Another important aspect of performance is the processor load required to achieve a given throughput. This is the portion of the overall performance that will differ the most between the internal and external MAC solutions.

In an industrial-type network, Ethernet can provide a basic system time using the Network Time Protocol (NTP). The synchronization across a controlled network for an NTP-based system is measured on a “human interface” time scale. While this protocol is suitable for general system clocking information, it is not precise enough in many industrial control systems where tighter synchronization needs to be in place.

To improve the precision, the IEEE 1588 Precision Time Protocol (PTP) standard was developed to be used in concert with the Ethernet controller and network stack to synchronize “local” clocks across a network with a master clock. That is, each processing or control node is synchronized with a master time reference driving the system.

By enforcing a precise clocking relationship across an industrial network, timed events can be synchronized to sub-millisecond levels. These timed events can include when analog/digital converters are sampled, when digital/analog converters are driven, or when I/O lines are activated for system control.

The IEEE 1588 PTP requires an exchange of specific packets to provide time information from two nodes. These packets are used to calculate the difference in the time and frequency between the clocks in each node. In addition, the protocol provides for a way to continuously adjust the clocks so they stay synchronized. The IEEE 1588 PTP protocol can be implemented either exclusively in software or through a combination of hardware and software. The hardware-based solution will provide for the best precision and therefore the greatest synchronization between nodes. With the hardware solution, the packet can be time-stamped as close as possible to its interaction point with the PHY. The result is lower jitter between the nodes.

PWM Units

A standard peripheral on microprocessors and DSPs is the general-purpose timer. This unit provides standard timer functionality based on one or more clock references that may be internal or external to the chip. At the pin interface, it also can provide width-capture or pulse-count capability, as well as single-ended pulse-width-modulated (PWM) output waveforms. These PWM outputs typically have programmable pulse width and period, and they can be utilized in many industrial control applications, including dc level generation and noise-resistant analog signal transmission (with appropriate lowpass filtering).

To be truly useful for ac motor control, however, several upgrades need to be added to the basic PWM functionality. Figure 1 shows a simplified motor control block diagram, where PWM outputs from the processor differentially drive high-side and low-side power devices to regulate the torque and speed of the motor. ADCs are

used to provide current measurement feedback to the processor so that the PWM duty cycles can be managed in a tightly sequenced closed loop system to control the motor.

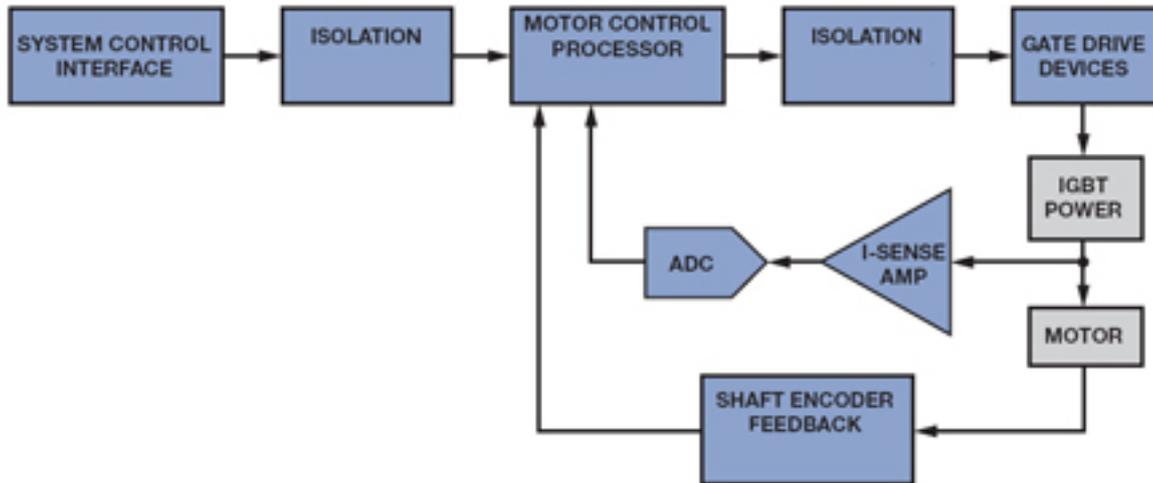


Figure 1: Simplified motor control signal chain

Compared with the PWM blocks on general-purpose processors, the PWM units for motor control have several enhancements. As implied above, motor control PWMs come in complementary pairs, in order to alternately drive the high-side and low-side power switch for a given motor phase. For a 3-phase ac motor, 3 PWM pairs are employed.

As shown in Figure 1, it is often important to provide isolation between the processor's PWM control and the gate drive devices for the power transistors. Commonly, this isolation is either performed via opto-isolators or pulse transformers. To this end, some PWM units feature a Gate Drive Unit that facilitates mixing the outputs with a high-frequency chopping signal for connection to pulse transformers, while also equipping the pin drivers with enough source and sink current to drive most opto-isolators.

Importantly, motor control PWMs must provide a guaranteed "dead time" between the deassertion of one power device and the assertion of its complementary pair on each phase. Otherwise, a dc short could occur through the power switch.

Moreover, there must always be a way to immediately and asynchronously disable the PWM outputs in case of an error condition where multiple output phases might be turned on simultaneously. This "PWM Trip" feature allows an external asynchronous signal to disable all PWM outputs regardless of the state of the processor clock.

Finally, while it is common to allow for general-purpose timers to be synchronously started, PWM timer synchronization takes on added importance for motor control. An internally or externally applied "PWM Sync" signal can be used to generate an interrupt (sometimes more than once per period) so that the processor can adjust duty cycles according to the control algorithm, and so that the ADC can acquire and

transmit the next current measurements.

By now, it should be evident that while many industrial applications might choose to use processors with a general-purpose peripheral set, it is wise to consider first which “industrial upgrades” would be beneficial for the application at hand. We’ve chosen to examine just two examples here – network connectivity and PWM capability – but parallels exist in many other subsystems, including memory structures and data conversion interfaces. By taking advantage of the added value of these augmented peripheral and system blocks, industrial products can be designed with improved robustness and system control.

Source URL (retrieved on 05/22/2015 - 7:53pm):

<http://www.ecnmag.com/articles/2009/12/industrial-strength-features-make-difference-embedded-processing>