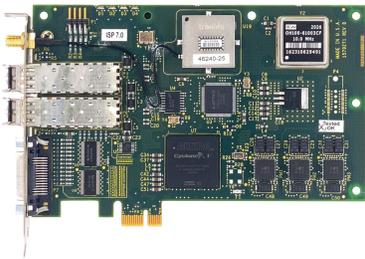


Real-Time Clock & Interrupt Module

Concurrent's Real-Time Clock & Interrupt Module (RCIM III) is a multifunction card designed for time-critical applications on Concurrent iHawk™ multiprocessing systems. The RCIM provides key functionality for single system as well as multiple system applications. Ideally suited for simulation, data acquisition and industrial control solutions, the RCIM includes a synchronized clock, eight programmable real-time clocks, and twelve input and output external interrupt lines. The RCIM mounts in a standard iHawk PCIe or PCI slot and is fully supported by Concurrent's RedHawk™ Linux® real-time operating system.



RCIM III

More Information:

concurrent-rt.com

wiki.simwb.com

info@concurrent-rt.com

800.666.4544

954.974.1700

FEATURES

- Deterministic event synchronization among multiple CPUs and systems
- Synchronized clock readable by all application tasks
- Eight interrupt-generating real-time clocks
- Clock synchronization to external time sources
- Twelve edge-triggered input external interrupt lines
- Twelve output external interrupt lines
- Twelve inter-CPU interrupts generated by real-time clocks or software
- GPS synchronization option
- Standard PCIe or PCI form factor
- RedHawk Linux and NightStar tool support
- Custom-engineered solutions

Real-Time Clocks

The RCIM III includes eight 32-bit, programmable, POSIX-compliant, real-time clocks. The clocks are programmable to several different resolutions which, when combined with a clock count value, provide a wide range of timing intervals. The minimum resolution is one microsecond. The real-time clocks are decrementing counters that generate an interrupt each time they reach zero. Real-time clocks may be one-shot or periodic. If periodic, the original clock value is automatically reloaded into the counter each time zero is reached and the clock again begins decrementing. This makes the real-time clocks ideal for running processes at a given frequency (e.g., 600 Hz).

External Interrupts

Each RCIM board has twelve input and twelve output external interrupt lines. Input interrupts are called Edge-Triggered Interrupts (ETIs), so named after their most common mode of operation. These lines permit users to provide their own interrupt sources to an iHawk system. The RCIM processes and delivers these interrupts to the host system and, if configured to do so, routes and delivers interrupts to all other RCIMs in the chain as distributed interrupts.

Each ETI can be configured independently of the others. An ETI may treat the incoming signal as an edge or level sensitive interrupt. If edge sensitive, it may raise an interrupt on either the rising or the falling edge. If level sensitive, it may raise interrupts for either the high or the low signal value. Configuration parameters are selected when the RCIM board and driver are initialized, and remain constant for the life of the driver. Applications can dynamically arm or disarm, enable or disable, select edge or level sensitivity and select the polarity for each ETI individually, as appropriate to the needs of the application.

Processor Synchronization

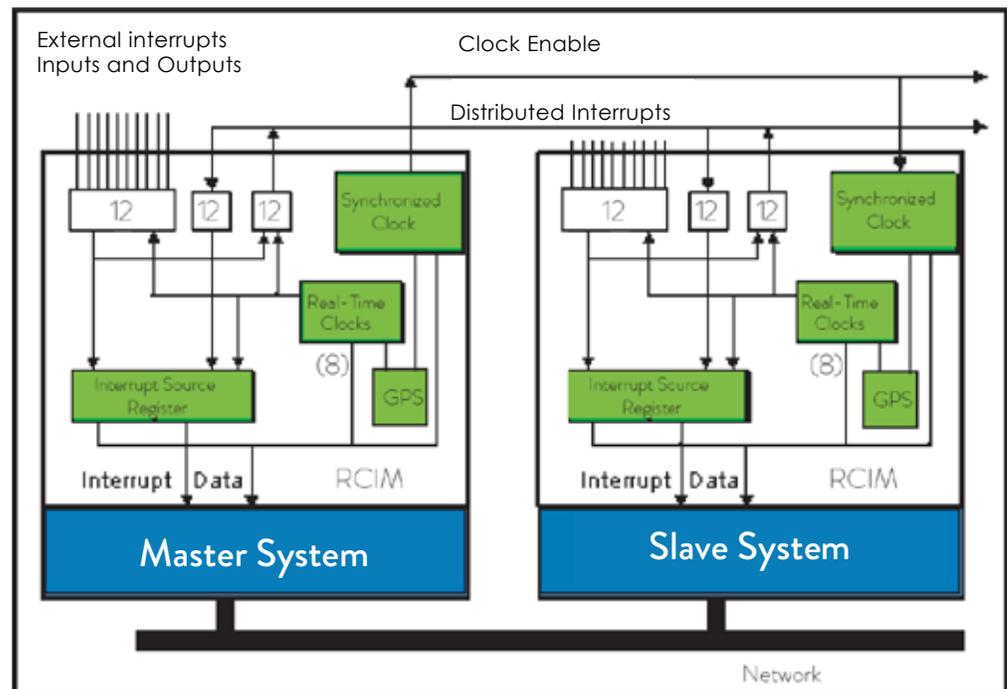
The RCIM provides fully deterministic event synchronization for applications that utilize multiple iHawk systems. The RCIM allows these systems to be chained together so that applications can direct both external and clock interrupts to other CPUs for fast inter-CPU event notification. RCIM interrupts can be triggered by hardware events or be generated by user software. The synchronized clock on each RCIM can be initialized by a designated master to provide a common, high-resolution time base for processes running on all connected systems. On-the-wire time stamps allow RedHawk Linux to provide for high resolution NTP synchronization between systems containing the PCIe version.

Synchronized Clock

The RCIM's synchronized clock is a 64-bit, non-interrupting, free-running counter that uses the POSIX 1003.1 format. The synchronized clock's high-order 32 bits are encoded in seconds, and the low-order 32 bits are encoded in nanoseconds. The clock ticks at 400 nanoseconds. Upon bootload of an RCIM-connected network, a designated master RCIM starts the synchronized clock on all RCIMs and issues an enable signal that simultaneously starts each clock. Application programs on all CPUs are then able to read a common time value. Individual RCIMs can be programmed to ignore the master synchronization signal and use their own synchronized clocks independently.

Distributed Interrupts

Distributed interrupts are signals sent from one RCIM to another via the RCIM synchronization cable. Each RCIM has twelve distributed interrupt inputs and twelve distributed interrupt outputs. Distributed interrupts can be generated by an RCIM's real time clocks, by its external edge-triggered input lines or by software. This feature allows an external event received by one CPU to be sent to one or more other CPUs.



Connectors and Cabling

The RCIM includes a serial synchronization cable for daisy chaining a master RCIM to one or more slave RCIMs. Any number of iHawk systems can be linked by RCIM serial synchronization cables. A connector is mounted on each RCIM for connection to external interrupts. Customized connectors and cable lengths are configurable upon request.

Optional GPS Synchronization

An optional, on-board GPS module is available for aligning the RCIM's synchronized clock to GPS standard time. One GPS-equipped RCIM can synchronize all iHawks in an RCIM chain. Antenna splitter modules for direct input of GPS signals to multiple systems are also available. Multiple RCIMs equipped with the GPS module can provide a common time base without any cable connections between iHawk systems. POSIX timers based on absolute GPS time can be used to simultaneously start the execution of programs on systems which are not physically connected.

Software Support

Concurrent's RedHawk Linux allows application developers to take full advantage of RCIM features. Users can read the synchronized clock, disable and enable individual interrupt lines, send interrupts to external locations, and direct incoming interrupts to user-specified application code. In multiple system applications, Concurrent's NightTrace analyzer tool can utilize the RCIM's synchronized clock to log system events using a single clock source. Concurrent's NightSim application scheduler can also make use of RCIM real-time clock interrupts to synchronize frequency-based scheduling among application modules running on multiple CPUs and systems.

Custom Engineering Support

Concurrent's Special Systems group is available to design and deliver solutions for customers who require features and functionality not included in Concurrent catalog products. Concurrent engineers can build special packaging, integrate third-party boards, develop drivers, and provide hardware and software designed to exact customer specifications.

SPECIFICATIONS:

SYNCHRONIZED CLOCK

Length: 64 bits (two 32-bit words)
 Resolution:
 • High-order 32 bits - 1 second
 • Low-order 32 bits - 400 nanoseconds
 Accuracy: ± 2.5PPM
 OCXO Freq. Stability Option: ±10 PPB

REAL-TIME CLOCKS

Number: 8
 Length: 32 bit
 Resolution: 1 microsecond (programmable to larger values)
 Accuracy: ± 2.5 PPM

LOCAL INTERRUPTS

12 External Edge-Triggered input
 12 External Output
 8 Real-Time Clock input

DISTRIBUTED INTERRUPTS

12 Input
 12 Output

INTERRUPT RESPONSE TIME

Interrupt to user process: <8 microseconds

PACKAGING

Form Factor: half-length
 PCIePerformance: PCIe x1
 Maximum cable length between RCIMs:3&,HIW

ENVIRONMENTAL

PCIe Operating Temp: 10° to 55° C
 PCIe Storage Temp: -40° to 70° C
 PCI Operating Temp: 10° to 40°
 PCI Storage Temp: -40° to 65° C
 Relative Humidity: 10 to 90%
 (noncondensing)
 RoHS compliant

©2022 Concurrent Real-Time, Inc. Information subject to change without notice. Concurrent Real-Time and its logo are registered trademarks of Concurrent. All Concurrent product names are trademarks or registered trademarks of Concurrent, while all other product names are trademarks or registered trademarks of their respective owners. The registered trademark of Linux is used pursuant to a sublicense from the Linux Mark Institute, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. All rights reserved. 121522