

**BR065R01 (submitted for p1394.1 committee vote):  
Clock sync protocols**

Dr. David V. James, Sony  
3300 Zanker Road, MS SJ3C1  
San Jose, CA 95134-4591  
Phone: 408 955-6295  
FAX: 408 955-4591  
Email: dvj@alum.mit.edu

**December 1, 1999**

**This contribution is one of several, presented for review and incorporation.  
An overall contribution, which provides an overall context for this and other contributions, is presented in BR047R10.**

This contribution proposes that all isochronous clocks are synchronized to the primary bus's cycle master. Two forms of clock are distributed, filtered and unfiltered. The unfiltered clock is the best-possible clock image, and doesn't pass through PLLs. The filtered clock is PLL'd before assertion on the bus, to eliminate cumulative jitter.



## 1.3 Net topologies

### 1.3.1 Hierarchical bus topologies

A configured Serial Bus net has one a prime-alpha portal (shaded black) and one or more alpha portals (shaded grey), as illustrated in figure 1. The primary bus has exactly one prime-alpha portal and the secondary buses have exactly one alpha portal. Each bus may have any number of additional nonalpha portals.




Figure 1—Bridged bus hierarchy

Within an active net, the bridge portal with the largest *portalID* identifier is selected to become the **prime-alpha portal** (often abbreviated as the **prime portal**). The 66-bit *portalID* value is a concatenation of a 2-bit software settable *preference* field and the portal's EUI-64 value.

NOTE—The *portalID* identifier does not include dynamic precedence information (such as the number of attached nodes or buses), as this would increase net refresh complexity and can sometimes be counter productive (see F.1).


The 2-bit *portalID.preference* field allows software to select its preferred prime portal, and the tie-breaking 64-bit *portalID.eui64* field guarantees the uniqueness of the *portalID* values. The assignment of *preference* values is expected to be performed by higher level software and is beyond the scope of this standard.

The bus with the prime-alpha portal is called the primary bus; other buses are called secondary buses. On secondary buses, the bridge portal that leads to the primary bus is called a **alpha portal**.

After completion of the net configuration process, any node in the net can be accessed by its unique 16-bit *nodeID* address. The *nodeID* address contains *busID* and *localID* components; in figure 1 these correspond to values {a,b,c} and {0,1,2,3} labels respectively. The bridge routing tables specify the *destination\_ID*-based routing paths for asynchronous request and response subactions.

### 1.3.2 Clock routing

To simplify interactions, all isochronous-capable nodes are expected to have synchronized timeOfDay clocks. The timeOfDay clock reference is generated by the cycle-master node on the primary bus, and distributed to other nodes, as illustrated by the dotted lines in figure 2 (the cycle-master nodes are cross-hatched).



**Figure 2—Clock synchronization**

On secondary buses, the clock-reference node is the alpha portal and clock distribution is more efficient if the alpha portal is also the cycle-master node. An alternate cycle-master node may be used, but only if the cycle-master node (c.3) has the capability of phase-locking to the clock-referencing (c.0) alpha portal.

## 1.4 Clock synchronization

### 1.4.1 Clock distribution

Although simple in theory, cascading phase lock loops, as illustrated in figure 3, complicates the dynamics of the overall system. This design approach should not be used.




Figure 3—Overly simple clock delivery

Instead, the clock estimate is forwarded through bridges, and PLLs are used to remove jitter from the local timer on each bus, as illustrated in figure 4. This reduces the dependencies of each bus clock from PLL dynamics of the others.




Figure 4—Desired clock distribution model

This design model mandates two time transmissions on the bus. A jitter-reduced signal marks the start of each isochronous cycle; in Serial Bus, the cycle-start packet transports this information. A second PLL source signal is repeated (with additional jitter errors) to the next bus with no PLL jitter filtering; the clockSync packet (see 8.4.2) serves this purpose.

### 1.4.2 Cycle-master synchronization

On remote buses, the clock-synchronization protocols proceed as follows. In every isochronous cycle a cycle-start packet is transmitted. During each cycle-start packet, all portals latch their current time value. During the following isochronous cycle, the clock-reference node transmits its previously latched value in a clockSync packet.

Other clock-slave nodes derive calculate an observed error as the difference between the clockSync packet and their latched value. That error is added to the free-running clock in the clock-slave nodes, to generate the clock-reference value that is distributed on the adjacent bus.

To minimize special routing requirements, the clockSync packets are “routed” by having alpha portals redistribute this information. The clockSync packet itself is not (strictly speaking) routed, but the observed bus-A time reference is either applied to bus-B or ignored, based on the portals’ source-routing tables.

### 1.4.3 Clock redistribution

To reduce the phase-lock whiplash effect, bridges are expected to redistribute a delayed clock reference and (when assuming the role of a cycleMaster) a phase-locked time in the cycle-start packet, as illustrated in figure 5. The intent is to redistribute an unfiltered (but delayed) clock reference, while phase-locking the cycle-start packet to reduce clock jitter effects.

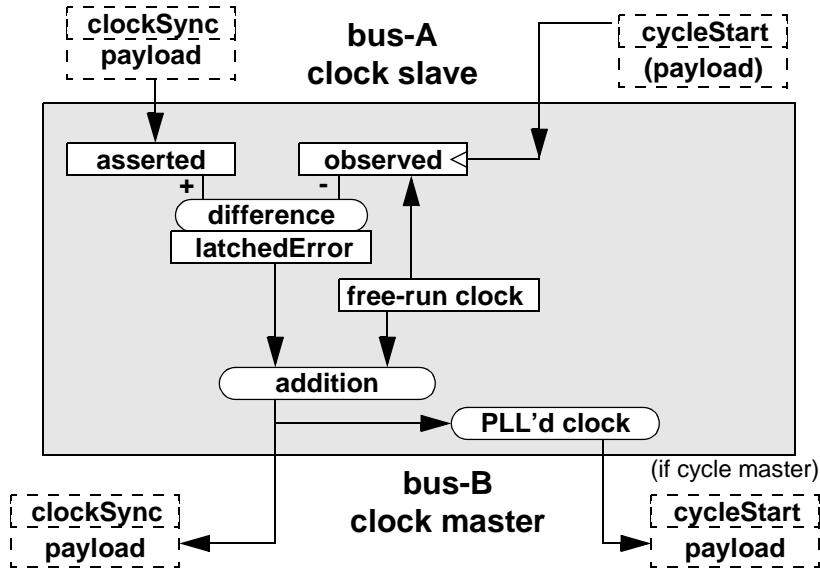


Figure 5—Clock redistribution model

For additional accuracy, the free-running clock value may be rate-adjusted, based on the long-term average of the *latchedError* value. Such rate adjustments shall not the free-run clock frequency to deviate beyond the specifications mandated by the attached bus standards.











