# Surreal interconnect requirements
# as submitted to the p1394.1 committee

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

**October 12 1999**

**This contribution is one of several, presented for independent review.
An overall contribution, which provides the context for multiple contributions,
is also provided in BR047R08.**

1394 may be applied to other physical interconnects, such as wireless media. To assist in the use of this an other media (such as infrared), there is a need to distinguish between physical layers properties (which may differ on other 1394-like interconnects) and mandatory services. For that reason, an annex is proposed.

# Annexes

# Annex B

(informative)

# Surreal bus (virtual 1394) requirements

## B.1 Design motivation

Several 1394 physical layers (other than broadcast copper) are being developed and desire to support 1394 device architectures. Although the wireless physical layer is primarily intended to support bridge attachments, node attachments are also possible, as illustrated in figure B.1.

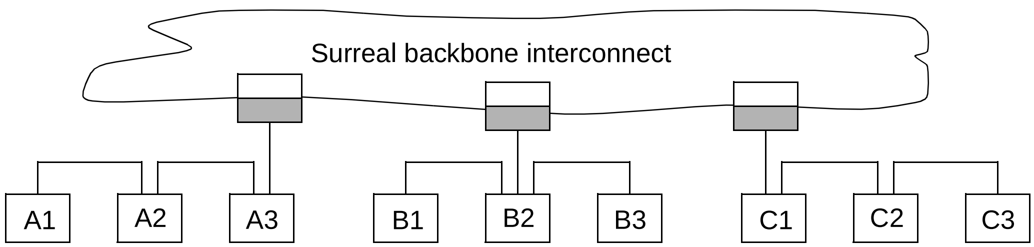


**Figure B.1—Wireless Surreal backbone**

Designing such a wireless physical layer could, in concept, be relatively simple: one simply supports 1394 services at the wireless interface. The existing 1394 specifications are much more than this; they represent combinations of higher level protocols (CSR register content and function, mid-level logical protocols (packet content interpretation), and a cable-specific lower-level physical model. As such, further clarification is needed to separate the 1394 transport services (the Surreal services) from the physical cable and signaling technologies.

This clause represents an attempt to summarize the Serial Bus services that alternative physical layers should provide. Surreal interconnects can present a consistent software interface and are easily bridgeable to other 1394 Surreal interconnects, including 1394, 1394a, and 1394b.

## B.2 Surreal interconnect properties

Mandatory and optional properties or a Surreal interconnect are listed in table B.1 and described in the remainder of this subclause.

**Table B.1—Surreal interconnect properties**

| Optionality | Property | Row | Description |
|---|---|---|---|
| mandatory | asynchronous formats | 1 | Request and response packet formats (selected components) |
| | isochronous | 2 | Isochronous packet format (selected components) |
| | independent streams | 3 | Request< response< isochronous transport precedence |
| | reliable transmission | 4 | Maximum effective error rates (see B.7) |
| | at least 512-byte payloads | 5 | At least 512-byte asynchronous-packet payloads are supported |
| | confirmed bus reset | 6 | Simultaneously observed by all local-bus nodes |
| optional | distinct confirmed events | 7 | Distince from bus reset: CSR interface change, phyId reuse, |
| | controlled courtesy events | 8 | Flow controlled courtesy events: ROM change, attach, detach |
| | header&data CRC | 9 | Should retain CRCs for uniformity andend-to-end coverage |
| | subtractive decode | 10 | Confirmed negative acknowledge when no consumer exists |
| | selective decode | 11 | Single/defined acknowledge when multiple consumers exist |
| | controlled events | 12 | Broadcast events are flow-controlled (not discarded) |
| | time-of-death | 13 | Packets transport a time-of-death, reducing timeout times |
| | pipelined | 14 | Multiple packets send before first is acknowledged |
| | concurrent | 15 | Starts other packet transmissions before first one completes |

**Row 1:** Primary components of Serial Bus asynchronous packets shall be transported. Secondary fields, including *data_length*, CRCs, padding, and local-retry fields, may be different in size or form.

**Row 2:** Primary components of Serial Bus isochronous packets shall be transported. Secondary components, which may differ on other Surreal interconnect, include length, CRCs, padding, and channel number fields.

**Row 3:** The delivery of isochronous packets shall not depend on asynchronous packet loading; the success of asynchronous response packets shall not depend on asynchronous request-packet loading

**Row 4:** The effective packet-error is assumed to be low, to reduce the impact of inefficient higher level retry protocols. Actual error rates may be higher, if ECC or fault-retry protocols reduce the perceptable effects.

**Row 5:** Up to 512-byte asynchronous-packet data-payload sizes shall be supported by the interconnect.

Optional features of the interconnect , which are expected to simplify interface designs, include the following:

**Row 7:** Robust confirmed event indications reduces the need to invoke frequent bus resets (on Serial Bus, a bus reset is the only confirmed broadcast event, so disruptive resets are invoked more often than necessary).

**Row 8:** Retaining Serial Bus CRC formats and data-padding allows bridge to be build in ways that support end-to-end data-payload coverage.

**Row 10:** A subtractive decode protocol returns a negative acknowledge (as opposed to no acknowledge) when no nodes respond to the packet-header specified address.

**Row 11:** A selective decode protocol selects one consumer node, when multiple nodes are enabled to acknowledge the packet-specified address. This transient condition can occur, for example, when bus-bridge routes are changed.

**Row 12:** Event-queue overflows are avoided by flow controlling broadcast event transmissions.
**Row 13:** Supplementing packets with time-of-death labels reduces worst-case timeout delays.
**Row 14:** Pipelining allows a node to send multiple packets before acknowledgment packets are returned.

**Row 15:** Multiple packets may be sent concurrently, meaning transmission of a second packet may begin before transmission of a first packet completes. As examples, concurrence may be possible over nonoverlapping cable segments, radio frequencies, or time-slots, depending on the interconnect media.

## B.3 Bus reset services

The broadcast nature of the 1394 physical layer has no provisions for supporting reliable confirmed broadcasts, other than a bus reset. For that reason, a bus reset is used for signaling events as well as reseting (possibly stuck) device interfaces. Surreal interconnects may support a finer-grained set of events, to reduce the need to reset bus interfaces.

### B.3.1 Self-id services

Surreal interconnects shall provide notifications of topology changes and (when changed) shall provide all attached nodes with *phyId* addresses of the currently attached nodes; these *phyId* addresses shall be within the range of 0-to-63 (inclusive).

### B.3.2 Supplemental selfId information

In addition to attached phyId information, Surreal interconnects may provide supplemental information, listed below (a leaf-limited bridge is defined in B.2):

1) Presence. A 1-bit link-on/off state indicates whether the node’s minimal CSRs can be accessed. This is unnecessary if the link is automatically activated when accessed.
2) Class. The following classification services simplify bus-bridge configuration services.
   a) Legacy. This is a legacy node (writes should be posted, read/locks rejected)
   b) Aware. This is a bridge-aware node; remote requests are processed normally.
   c) Portal, type-A. This is a leaf-side portal of a leaf-limited bridge.
   d) Portal, type-B. The top-side portal of a leaf-limited bridge.
   e) Portal, type-C. Either side of an non-leaf-limited bridge.

### B.3.2.1 Other event indications

Other event notifications that are implied by a 1394 bus reset, but could be signaled in other ways, include the following:

1) ROM change. The contents of ROM (as visible through the node's CSR space) have changed.
2) Reclaim. Isochronous resources may be inconsistent with attached nodes.
Nodes are normally reclaim their resources in a timely fashion (1 second), or lose them.
3) Attached. A node has been attached and now responds to a previously unused phyId address.
4) Detached. A node has been detached and no longer responds to a previously used phyId address.
5) Changed. A node has been assigned to a previously used phyId address.

## B.4 Special control

In low-power mode a node may not capable of responding to CSR accesses. If such idle-node states are allowed, a link-on indication (that activates a known-to-be or thought-to-be node) should be provided.

## B.5 Packet formats

### B.5.1 Asynchronous packet formats

Most of the 1394 packet contents are interconnect independent, but a small portion of these packets is not, as illustrated in figure B.2.

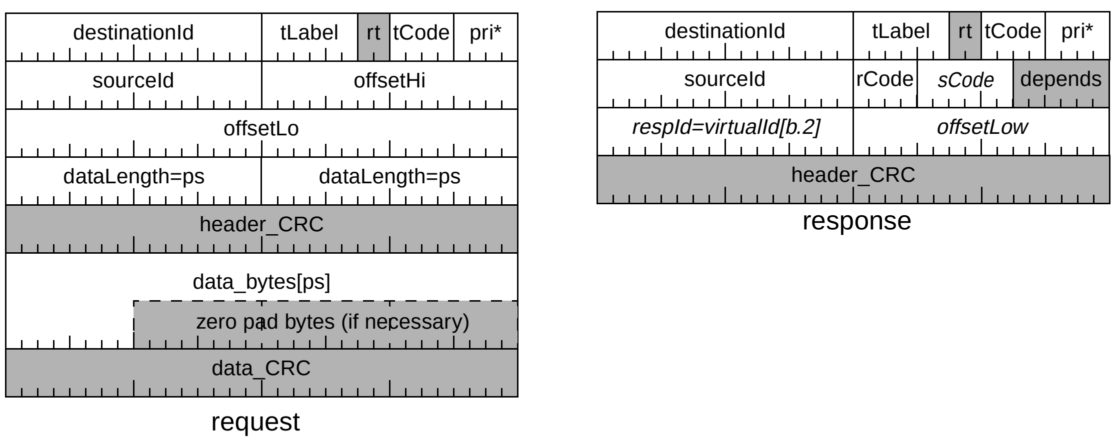


**Figure B.2—Request and response subactions**

Within these 1394 packets, the shaded portions are not part of the Surreal services. Further details are provided in the remainder of this subclause.

The 2-bit ***rt*** (retry type) field is used to signal busy-retry phases and is updated on a bus-local basis. Although flow-control services shall provide equivalent functionality, this is not a Surreal field.

The 4-bit ***pri*** (subaction priority) field should be used to classify the packet's priority, so that higher priority packets receive preferential services. Although Surreal services mandate the transport of this field, the granularity of priority services is not mandated (cable 1394 supports only one level).

The 32-bit ***header_CRC*** and 32-bit ***data_CRC*** fields are used on cable 1394 to verify the packets correctness, to reduce the possibility of misinterpreting corrupted packets. Although Surreal services mandates maximum levels of packet loss and corruption (see B.7), the encodings of EDC and ECC codes are not specified. Although physical layers should provide separate header and payload CRCs, this is not mandated and a single packet CRC may also be used.

The zero-pad bytes, as well as their absence (or presence) at the beginning of the packet are not mandated. A physical layer may provide initial padding to maintain convenient address alignments, or (as is true for 1394) may provide trailing padding to maintain convenient packet-size alignments.

### B.5.2 Isochronous packet formats

#### B.5.2.1 1394 isochronous packet content

Most of the 1394 isochronous packet contents are mandated by Surreal services, but small portions of these packets are physical layer dependent, as illustrated in figure B.3.

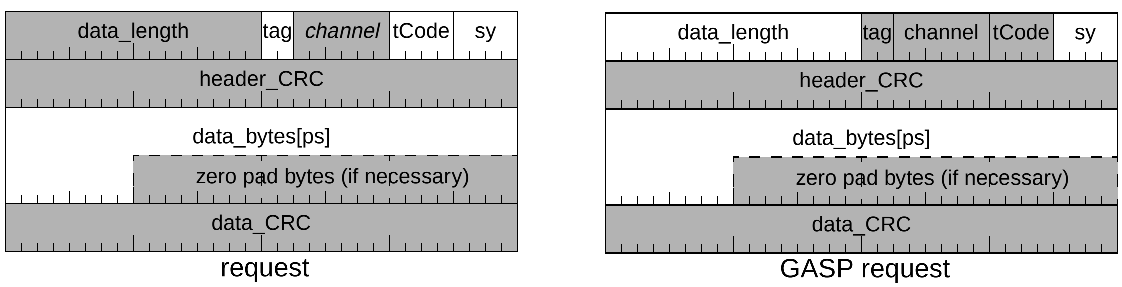


**Figure B.3—Isochronous and GASP packets**

Within these 1394 packets, the shaded portions are not part of the Surreal services. Further clarification is provided in the remainder of this subclause.

The packet length must be specified, but a smaller field or special end-of-packet indications are viable alternatives to the 16-bit *data_length* field. A *channel* number field is required, but this field may be anywhere between 4 and 16 bits in size.

The 32-bit ***header_CRC*** and 32-bit ***data_CRC*** fields not mandated, as described in B.5.1. The zero-pad bytes, as well as their absence (or presence) at the beginning of the packet are not mandated, as described in B.5.1.

#### B.5.2.2 CIP-header mandated content

A certain class of isochronous packets, identified by a 1-valued tag field, has a standardized CIP (common isochronous packet) header. A portion of this header sometimes contains a time stamp, that must be modified as these packets pass through bridges (see D.2.1).

A Surreal interconnect may transport a relative (as opposed to absolute) *time-stamp* fields in CIP-headers. This adds no complexity to the heterogeneous bridge and simplifies the design of homogeneous bridges.

The CIP header also contains a 6-bit *sid* (source identifier) field. On single bus systems, this is used to identify the source of the isochronous packet. The Surreal services mandate this functionality, but interconnects may elect to prepend packets with this information to avoid modification of payload content.

Again, this adds no complexity to the heterogeneous bridge and simplifies the design of homogeneous bridges.

#### B.5.2.3 GASP content

On 1394, a class of packets (called GASP, for general asynchronous stream packets) have an isochronous format but are transmitted using asynchronous rules. Since these packets may be lost (or merged) under heavy loads or transient errors, higher levels are expected to compensate for packet losses. Therefore, the acceptable error rates are independently specified for GASP packet transmissions (see B.7).

Events are a special easily merged type of GASP packet. When possible, interconnects should flow-controlled event transmissions, so that information losses associated with GASP-packet merging is avoided.l

## B.6 Packet delivery services

### B.6.1 Directed asynchronous addressing

On 1394, packets with remote-bus addresses are broadcast to all possible portals and address recognition is done by the consuming portal, which creates the acknowledge that is returned to the producer. Other Surreal interconnects may use a local address resolution protocol (LARP) to discover physical-layer routing information, allowing subactions to be routed over non-broadcast-capable interconnects. The LARP may use an inefficient broadcast (such as a sequence of directed writes), since LARP results can be cached for later reuses.

### B.6.2 Asynchronous payloads

Surreal interconnects, even those lower than 100Mbs, shall support 512-byte asynchronous-packet data payloads, regardless of transfer rate. This ensures a minimal guaranteed end-to-end transfer size through bridges.

Surreal interconnects may support larger payload sizes, up to the packet-format-constrained size of 65535 data bytes. Payload sizes longer than an isochronous cycle are allowed, but (if implemented) the interconnect should provide facilities for interleaving asynchronous and isochronous traffic.

### B.6.3 Isochronous services

Data delivered, with constant delay, on 125µs intervals (or multiples thereof)

Timely delivery in the presence of arbitrary request and response traffic.

1) Clock source.
2) Clock-slaving capability.
3) IRM (isochronous resource manager). The IRM services, which are available to nodes and portals, provides a pass/fail indication and channel number to isochronous resource requests.

## B.7 Effective packet-transmission error rates

Effective error rates of the media (after performing physical-level error ECC and retries) should be less than specified in table B.2. In computing these rates, a 50/50 mixture of 512-byte read and write transactions shall be assumed.

**Table B.2—Error rate constraints (suggested)**

| Traffic | Error type | Value | Row | Description |
|---|---|---|---|---|
| Asynchronous | detected-payload | -TBD- | 1 | Effective data_error response rate |
| | detected-header | -TBD- | 2 | Effective asynchronous packet-loss rate |
| | undetected | $10^{-20}$ | 3 | Effective asynchronous packet-corruption rate |
| GASP packets | detected-payload | -TBD- | 4 | Effective indication-lost rate |
| | detected-header | -TBD- | 5 | Effective payload-loss rate |
| | undetected | $10^{-20}$ | 6 | Effective false-indication rate |
| Isochronous | detected-payload | -TBD- | 7 | Effective data_error interpretation rate |
| | detected-header | -TBD- | 8 | Effective isochronous packet-loss rate |
| | undetected | $10^{-20}$ | 9 | Effective isochronous packet-corruption rate |

### B.7.1 CSR requirements

CSRs: 1212r compliant.

Bus_Info_Block—Similar, but distinguished from other Serial Bus variants. Coding is TBD.

TBD—Document which CSRs are required and why.