

IsoGrid Protocol Specification

Version 0.100

What follows is a free and open proposal for a new open network protocol with a grid topology.

1 Licensing & Legal

This protocol specification (*The Protocol*) is freely available for all to use as is.

The Protocol, legally speaking, is completely Public Domain: CC0

However, just because you have a legal right to do something, does not mean you *should* do something. The right thing to do, morally speaking, is not codified in law.

The Protocol was released to further the following socioeconomic goals:

- Lower barriers to entry in markets for goods and services that rely on networks
- Empower individuals to improve their lives
- Increase individual freedom

In particular, I believe *The Protocol*, when widely implemented, will further the above goals.

If, in the 10 years following the release of this version of the specification, you want to implement a change to *The Protocol*: You **MUST** make a good faith effort to ensure that your changes to *The Protocol* do not undermine the above goals.

The simplest way to do this is to openly declare your intended changes at the IsoGrid Forum, and see if the community agrees.

Implementing a change to *The Protocol* that undermines the above goals **MUST** be considered a form of corruption; akin to taking more than your fair share from a commons. It **MAY** be legal, but you **MUST** expect negative social consequences if/when this comes to light.

I hereby release these moral conditions for all uses of this version that follow 10 years after this specification version is first released into the public domain.

2 Introduction

This proposed specification attempts to meet the requirements of a new network protocol with a grid topology.

2.1 Definitions

Term	Description
Node	A device that interacts with other devices on The IsoGrid using the IsoGrid Protocol
Link	The protocol running on a physical wire or wireless line that connects two adjacent nodes
Switch	A node with multiple links that is able to route data using the IsoGrid Protocol
Word	The atomic unit of transmission across the network. 16 bits of data, 1 bit of parity
Slot	A 1 word wide logical division of the link's bandwidth, delivered isochronously 1 word at a time. Each available slot on a link can be allocated to switch a single connection stream at any given moment.
Input Slot	A slot on an input link. Has a matching output slot on the node on the other end of the link.
Output Slot	A slot on an output link. Has a matching input slot on the node on the other end of the link.
<i>IsoConn</i> / Connection	End-to-end Isochronous Connection. A stream of words that flows isochronously from a source to a destination across a pre-defined subset of the nodes that comprise The IsoGrid
Frame	A Link-Layer logical aggregation of individual isochronous slots to be sent together across a single link. Note: Frames are NOT to be thought of as network packets; they do NOT route across the IsoGrid network.
Advertise	To make something publicly known to any network participant that asks

2.2 Layers

There are 3 layers relevant to this specification:

Layer 0: Physical Layer

- Options include, but are not limited to:
 - Ethernet, ATM, USB, etc.
 - Tunnels through other networks (like the TCP/IP Internet, etc.)

Layer 1: Link Layer

- Defines how nodes directly communicate with each other across links
- The IsoGrid does NOT mandate a globally-required protocol at this layer
- The IsoGrid does impose a something generic requirements at this layer

Layer 2: Network (IsoGrid) Layer

- Defines how the network uses source routing for Isochronous Connections (*IsoConn*)
- An *IsoConn* is switched at the word level, one word at a time
- Defines how nodes exchange credits so as to allocate scarce resources

Higher Layers: All the normal protocols you would expect to run on networks

2.3 How it works

The IsoGrid is a mesh network that supports routing of one-way isochronous streams (an *IsoConn*). In order to provide isochronous streams, the IsoGrid runs with a synchronized frequency, very similar to the way an electrical grid runs on [Utility Frequency](#) (except much higher frequency). The source provides a series of route instructions to be used at each hop ([Source Routing](#)). Each switch along the way uses its route instruction to establish the connection. The header that starts a connection defines the length of the *IsoConn*, and how many credits to send per word. The IsoGrid is optimized to support nodes fracturing the data, applying a forward error correction code, and sending the data over many paths across the network.

2.4 Protocol Secondary Limitations

No network is without limits. The design of a protocol standard necessitates making tradeoffs in order to meet the requirements. Here are some of the known limitations due to the design of the IsoGrid Protocol:

- A single connection can use no more than the fastest available slot along a route
 - However, an endpoint can create multiple connections such that nearly 100% utilization with connections is possible
- low-rate connections have Higher Latency
- Half-Duplex links not supported
- Connections through nodes that move will have non-trivial buffering/timing requirements
 - The faster it moves, the more demanding the requirements
- TODO: Add more as new limitations are identified

3 Isochronous Word Format

The atomic unit of transmission across the IsoGrid network is called a *word*. A word is 16 bits of data, with an additional 1 bit for parity. The parity of a word is defined to be EVEN if all 17 bits of the word have an even number of 1 bits. The parity of a word is defined to be ODD if all 17 bits of the word have an odd number of 1 bits.

Valid words meant for *IsoConn* establishment and *IsoConn* payload MUST have odd parity, and are called CONN_ODD words, or abbreviated as COWORD.

Valid words meant for non-*IsoConn* communication MUST have even parity, and are called SIG_EVEN words, or abbreviated as SEWORD.

The following table shows common words types and their expected parity

Type of word	Value	Parity	Description
NoData	Link Defined. Suggested: 0x0000	SIG_EVEN	This isn't necessary, but a Link Layer protocol MAY find it useful to have a designated word that indicates no data is available on a slot.
Data	Node Defined	SIG_EVEN	Data sent from one node to the neighbor node using a Link Layer protocol. Not part of an <i>IsoConn</i>
<i>IsoConnInit</i>	IsoGrid Protocol Defined Sent by Source Node	CONN_ODD	These words initialize the payment and initial word count of an <i>IsoConn</i> . The structure is specified in this document, and the values are produced by the source node.
<i>IsoConnRoute</i>	Switch Defined Sent by Source Node	CONN_ODD	These words are defined and advertised by each switch to describe each step of a route that an <i>IsoConn</i> is going to take through the network.
<i>IsoConnHeader</i>	Source Defined	CONN_ODD	These words are sent by the source as a series of headers to the payload of an <i>IsoConn</i> .
<i>IsoConnPayload</i>	Source Defined	CONN_ODD	These words are sent by the source as the payload of an <i>IsoConn</i>
<i>IsoConnFooter</i>	Sent by switch	CONN_ODD	These words are sent by each switch along the route as the final parts of the <i>IsoConn</i> .
Lost/Corrupted <i>IsoConn</i> *	<i>LostWord</i> (0x0000)	SIG_EVEN	If an <i>IsoConnRoute</i> , <i>IsoConnHeader</i> , or <i>IsoConnPayload</i> word is lost or corrupted along the route through the network, that word MUST be replaced

			with <i>LostWord</i> to signal this fact to the destination.
--	--	--	--

4 Link Layer Protocols

There are many possible ways to define a protocol for the Link Layer. The IsoGrid Protocol does not mandate any specific protocol or implementation at the Link Layer. As such, it is NOT necessary that everyone in the world agree to any standard protocol(s). Deciding on a Link Layer protocol is entirely a local decision between two neighbor nodes.

That said, the IsoGrid Layer does impose some non-trivial requirements on the Link Layer below it:

- The Link Layer MUST meet the Slot Isochronous Standard below
- The Link Layer MUST meet the Slot Frequency Standard below
- The Link Layer MUST be able to detect the start of an *IsoConn* in a slot. See section © below

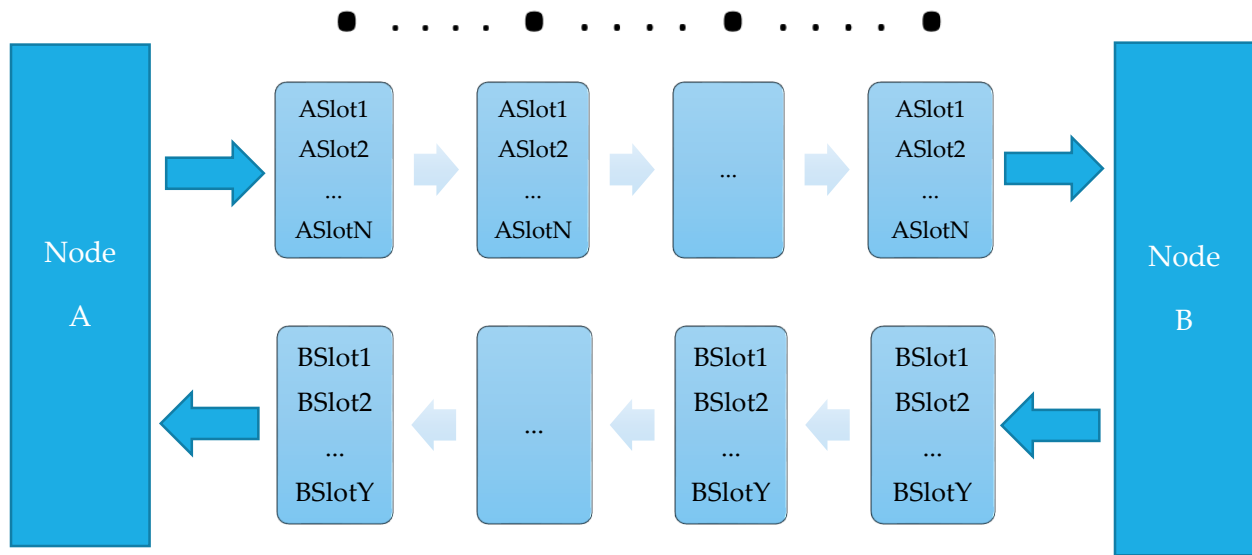


Figure 1. Link Layer frames exchanged between two neighbor nodes

Figure 1 shows a generic full duplex link exchanging isochronous frames between two neighbor nodes. Notice how the words on a slot arrive at well-defined periodic intervals and the input slots are unrelated to the output slots. Also, notice how the number of slots going one way does not have to match the number of slots going the other way (though typically they probably will).

4.1 Slot Isochronous Standard

The IsoGrid Protocol requires the existence of isochronous slots on all links across the IsoGrid. An isochronous slot is a 1 word long logical division of the link's bandwidth, delivered isochronously 1 word at a time.

This means that the words belonging to a slot MUST be sent across the link and arrive at well-defined 2^n word / second frequencies (where n MUST be a non-negative integer). A Link Layer protocol MUST define some sort of isochronous frame format logically divided up into slots:

- A frame MAY contain any number of words
- A frame MAY have some or all of the words broken into 2, 4, or 2^n slots such that each slot appears every other frame, or every 4th frame (and so on)
- Concurrently, A frame MAY have some or all of the words aggregated into a smaller number of slots, such that 2^n words arrive for a slot every frame.
- A node MUST be able to identify the slots of a valid frame that arrives on one of its links
- A node SHOULD be able to identify the slots of a valid frame even immediately following missed or corrupt frames
- A node MUST be able to assign an exact total ordering and count to every valid frame it receives
- Electronics MUST be fast and stable enough such that skipping a frame, or somehow seeing two when there was only one, MUST be extremely improbable
- A node MUST advertise the best-case and worst-case word-corruption rates of its output links
- A node MAY use error-correction codes to ensure the error rate meets the advertised value
- A frame MAY contain hashes to determine validity
- A node MUST be able to detect corruption of frames or words that occurs on its input links
 - Bit error(s) detected in a slot assigned to an *IsoConn* MUST be re-transmitted as *LostWord*.
 - A slot MUST have fewer than 1 undetected corruption of a word in every 1024^6 words
- Words from entirely failed links MUST be assumed to be *LostWord*.

Some example frame format protocols are:

- Statically sized frame:
 - Y ordered words, each their own slot
 - Y ordered words, Y/2 slots occur every frame, Y half-rate slots occur every other frame
- Negotiation at initialization-time decides a static size of the frame, in words
- Negotiation at any time can dynamically change the size of the frame
- Negotiation at any time can dynamically change the size of the frame and/or change the logical allocation of slots

When a slot does not have an allocated *IsoConn*, the sending node MAY use that slot to send non-*IsoConn* data to its neighbor (1 word at a time per slot). If an *IsoConnInit* arrives on a slot currently being used for non-*IsoConn* data, the node MUST prioritize the *IsoConn* over the non-*IsoConn* data.

Some possible non-*IsoConn* uses include (but are not limited to):

- Probe for connection availability
- Request/confirm a link reservation
- Send routing information (perhaps based on physical location)
- Synchronize clocks/frequencies (if used/needed)
- Switch actual packets
 - Computational packets
 - Physical location based routes?
- Update link state
 - Link drop announcements
- Transfer/exchange Credits
- Check credit statistics

Non-*IsoConn* data MAY be ignored/dropped if uncorrectable errors are detected in a link layer frame.

The frequency of word arrivals for an isochronous slot MUST be a power of two words/second. For example, $2^{13} = 8,192$ words/second, or $2^{14} = 16,384$ words/second.

But at very high frequency, the precise definition of a second is relevant. The definition of a second on the IsoGrid is provided by the Link Slot Frequency Standard.

4.2 Link Slot Synchronized Frequency Standard

The IsoGrid system frequency standard is TCG: Geocentric Coordinate Time

In order to provide isochronous streams, the IsoGrid runs with a synchronized frequency, very similar to the way an electrical grid runs on [Utility Frequency](#) (except much higher frequency).

Nodes MUST attempt to lock their link output frame rate using the TCG definition of a second. However, there are a number of ways that nodes MAY meet this requirement.

4.2.1 Stationary Nodes

If a node is stationary relative to its neighbors and has just a single link, the node MAY directly clock its output frame rate synchronized to its input frame rate. The link for these edge nodes MAY be arbitrarily long-lived.

If a node is stationary relative to some number of neighbor nodes, the node MAY tune an oscillator with respect to those input frame rates.

The tuned oscillator will be used to set the output frame rate. The neighbors that receive this as an input will use it to tune their own oscillator, which will be set the frame rate of the input links of its neighbors (including back to the first node). In this way, a stabilizing feedback loop will work to synchronize the entirety of The IsoGrid. The links for these stationary nodes MAY be arbitrarily long-lived.

To illustrate how this might be implemented within a node, here are a few examples:

- Combine all the input waveforms and use that combined waveform to tune the output frequency oscillator
- This waveform feedback could also be a digital process, where the 'fullness' of buffers determines the 'waveform' phase shift.
 - Buffers filling up: Advance the waveform
 - Buffers getting empty: Retard the waveform

This frequency synchronization strategy establishes a fundamental tradeoff between the following:

1. Higher link frame rate
2. Smaller buffers at each hop
3. Longer distance links
4. Higher tolerance for clock drift and clock skew

Longer links have more frames in transit. Higher frequency links also have more frames in transit. The more frames in transit, the more buffer is needed to accommodate clock instability versus ideal TCG.

Stationary nodes that have an amazingly stable TCG input SHOULD bias their output frame rate to attempt to match the TCG frame rate. This is intended to pull its neighbor nodes closer to TCG. The IsoGrid as a whole is reliant on the collective work of all nodes with TCG inputs to ensure the entire network locks to the TCG frame rate over time.

4.2.2 Feedback-Mediated Link Frame Rate Synchronization

The IsoGrid protocol's concept of using input frame rates to drive output frame rates is referred to as "Feedback-Mediated Link Frame Rate Synchronization" or FMLFRS.

Given a Frame Rate, a Link Latency, and the measure of the clock's short-term stability, it's possible to specify the minimum buffer required to compensate for clock drift and skew.

Here is a basic mathematical expression that expresses the fundamental tradeoffs precisely.

<i>Rate</i>	Frame rate of the link, in frames / second
-------------	--

<i>Distance</i>	The round-trip link distance, in meters
C_{medium}	Speed of information travel in the transmission medium, in meters / second
<i>Latency</i>	The round-trip latency of the link := $\text{Distance}/C_{\text{medium}}$
<i>ClockStability</i>	Clock stability of the node, expressed as a fraction. For example: 1 part in a million --> 0.000001
<i>MinBuffer</i>	Minimum possible buffer required to accommodate clock drift and clock skew, in number of frames

$$\text{MinBuffer} = \text{Rate} * \text{Latency} * \text{ClockStability}$$

In practice, the required buffer will be larger, but it seems reasonable to expect that it's within two orders of magnitude. If a clock is stable to 1 part in 500,000 (ie. A simple quartz clock), then even if it takes 100x the *MinBuffer*, the underlying latency of the link is only increased by 0.02%.

4.2.3 Clock examples

A quartz clock, for example, typically has a short term stability of 1 part in half a million. This means that a link with one frame of buffer can have up to 500K frames in transit (round-trip) after which feedback-mediated link sync is impossible. For a 100km link, this leads to a maximum theoretical 50 mega-frame / second rate (with only 1 frame of buffer).

4 frames of buffer would allow 4 times the number of frames in-transit.

Here, a frame travelling round-trip across a single link is defined to be in-transit up until the frame is able to be used in the frequency feedback mechanism of its sender.

Clearly, quartz clocks alone aren't stable enough for use with extremely high-rate, long-distance connections, where link sync can be lost before the frequency feedback loop is able to correct the issue. A node MAY mitigate this issue by using larger buffers. Since the number of buffered frames at the end of a link would be quite small compared to the link itself, quartz clocks are likely to be good enough for most nodes for at least the next decade.

A GPSDO clock, on the other hand, typically has a short & long term stability of 1 part in 300,000,000,000. With this clock, a link with one frame of buffer can have up to 300G frames in transit (round-trip) before feedback-mediated link sync is impossible. For a 1000km link, this leads to a maximum theoretical 240 Tera-frames / second rate.

4.2.4 Dealing with bad clocks

Since most nodes rely on their neighbors to collectively lock to the TCG frame rate, a node with a misbehaving clock will have local impacts. It could potentially cause a group of neighbors to lose link sync frequently. Since this is a local issue, it can be dealt with at the local level by the

affected neighbors making the choice to stop using the bad clock as a clock synchronization source. That way, the bad clocked node alone has the consequences of the bad clock. There is no need for a Time Cop :-)

4.2.5 Moving Nodes

Nodes that move, but that stay 'near' a starting position, MAY compensate for the movement with buffers; either logical buffers, or physical buffers (in the form of a longer link distance). In so doing, they MAY provide arbitrarily long link connectivity.

However, nodes that move arbitrarily long distances, MUST have transitory links. These nodes MAY pre-compute the buffer requirements for a transitory link. As the node continues to move, it can only maintain the link for so long before the buffers are exhausted.

For a node that is moving axially between two other nodes, it MAY consider clocking the output frame rate based on the opposing input frame rate. Doing so could allow a smaller buffer.

5 Credits and Micro-Payments

In order to avoid a "tragedy of the commons", the IsoGrid allows for exchanging credits between neighbor nodes so as to pay for the data sent or processed. This allows payments to be made among neighbors instead of having to have a centralized payment processor.

Each node owns its outbound links (but not really its inbound links), its CPU hardware, its storage, etc. Each node SHOULD charge credits for use of those resources. Each pair of neighbor nodes SHOULD agree on the settlement instrument that will represent the value of a credit between themselves, and payment SHOULD be by simple agreement (there is no required third party).

Any settlement instrument is possible, but here are some examples:

- electricity
- bitcoins
- cash
- check
- ACH transfer
- propane
- gasoline
- gold
- water

While any settlement instrument is possible, the value of one credit MAY be scaled such that 10 credits is approximately equal to the fair price to transmit 1 word across a link.

Neighbor nodes MAY choose to have a maximum per-diem settlement, to provide a backstop against possible software bugs or security vulnerabilities in the system.

6 IsoGrid Headers

Each *IsoConn* has a sequence of headers that follow the routing envelope. There is no globally required header.

Each header starts with one of two things:

1. An 8 word *ProtocolIdentifier* value, where bits 15, 11, and 10 of the first word are all cleared
2. A 1-word *ProtocolAlias* value, where bit 15 is set

Bits 12:14 of these values indicate the length of the header's data section which follows directly after the *ProtocolIdentifier* or *ProtocolAlias* value:

Bit Sequence	Header Data Length	Comment
000	0	Useful for 'flag' type headers
001	2	
010	8	
011	16	
100	32	
101	64	
110	128	
111	1024	

There is no central authority for assignment of *PrototcolIdentifier* values. Instead, anyone MAY create a cryptographically random 128 bit number, clear bit 15, 11, and 10 of the first word, and use it as a *ProtocolIdentifier*.

A *ProtocolAlias* is used by agreement between the source and destination in place of an actual *ProtocolIdentifier*. This is just to allow for decreased overhead and latency. There is no global requirement that a *ProtocolAlias* maps to any specific *ProtocolIdentifier*, and it's conceivable that some endpoints might not recognize any *ProtocolAlias* values.

The end of the header sequence is marked by a header end word:

Name	Value
HEADER_END_COWORD	0x8000

6.1 Scenarios

To illustrate the potential usefulness of headers, here are some example scenarios:

- Flagging the method by which a payload is encrypted with Public Key Encryption
- Holding session identifiers to support symmetric encryption
- Describing individual streams of Forward-Error-Corrected data, ie. 16 streams were sent, which one is this?
- Defining the encapsulated upper layer protocol
- Describing how long the payload data is
- Etc.

7 Isochronous Connections

An example IsoConn across 2 switches (3 links) might look like this:

Source Sends	Switch A Sends	Switch B Sends	Destination Processes as Payload
*	*	*	*
CO – IsoConnInit0	*	*	*
CO – IsoConnInit1	*	*	*
CO – IsoConnRoute_A	*	*	*
CO – IsoConnRoute_A	CO – IsoConnInit0	*	*
CO – IsoConnRoute_A	CO – IsoConnInit1	*	*
CO – IsoConnRoute_B	CO – IsoConnRoute_B	CO – IsoConnInit0	No (IsoConnInit)
CO – IsoConnRoute_B	CO – IsoConnRoute_B	CO – IsoConnInit1	No (IsoConnInit)
CO – IsoConnHeader0	CO – IsoConnHeader0	CO – IsoConnHeader0	CO – IsoConnHeader0
CO – IsoConnHeader1	CO – IsoConnHeader1	CO – IsoConnHeader1	CO – IsoConnHeader1
...
CO – IsoConnHeader9	CO – IsoConnHeader9	CO – IsoConnHeader9	CO – IsoConnHeader9
CO – HeaderEnd	CO – HeaderEnd	CO – HeaderEnd	CO – HeaderEnd
CO – Payload0	CO – Payload0	CO – Payload0	CO – Payload0
CO – Payload1	CO – Payload1	CO – Payload1	CO – Payload1
SE - <i>SetInBandSignal</i>	SE - <i>SetInBandSignal</i>	SE - <i>SetInBandSignal</i>	No (<i>SetInBandSignal</i>)
CO – Payload2	CO – Payload2	CO – Payload2	CO – Payload2
CO – Payload3	CO – Payload3	CO – Payload3	CO – Payload3
CO – <i>InBandSignal0</i>	CO – <i>InBandSignal0</i>	CO – <i>InBandSignal0</i>	No (<i>InBandSignal</i>)
CO – Payload4	CO – Payload4	CO – Payload4	CO – Payload4
CO – Payload5	CO – Payload5	CO – Payload5	CO – Payload5
CO – Payload6	CO – Payload6	CO – Payload6	CO – Payload6
CO – <i>InBandSignal1</i>	CO – <i>InBandSignal1</i>	CO – <i>InBandSignal1</i>	No (<i>InBandSignal</i>)
CO – Payload7	CO – Payload7	CO – Payload7	CO – Payload7
CO – Payload8	CO – Payload8	CO – Payload8	CO – Payload8
CO – Payload9	CO – Payload9	SE – <i>LostWord</i>	No (<i>LostWord</i>)
CO – <i>InBandSignal2</i>	CO – <i>InBandSignal2</i>	CO – <i>InBandSignal2</i>	No (<i>InBandSignal</i>)
CO – Payload10	CO – Payload10	CO – Payload10	CO – Payload10
CO – Payload11	CO – Payload11	CO – Payload11	CO – Payload11
CO – Payload12	CO – Payload12	CO – Payload12	CO – Payload12
CO – <i>InBandSignal3</i>	SE – <i>LostWord</i>	SE – <i>LostWord</i>	No (<i>InBandSignal</i>)
...
CO – Payload79	CO – Payload79	CO – Payload79	CO – Payload79
CO – Payload80	CO – Payload80	CO – Payload80	CO – Payload80
CO – Payload81	CO – Payload81	CO – Payload81	CO – Payload81
CO – <i>InBandSignal26</i>	CO – <i>InBandSignal26</i>	CO – <i>InBandSignal26</i>	No (<i>InBandSignal</i>)
CO – Payload82	CO – Payload82	CO – Payload82	CO – Payload82
SE – Filler (NoData)	SE – Filler (NoData)	SE – Filler (NoData)	No – Filler (NoData)
SE – Filler (NoData)	SE – Filler (NoData)	SE – Filler (NoData)	No – Filler (NoData)
*	CO – IsoConnFooter_A	CO – IsoConnFooter_A	No (Footer)
*	CO – IsoConnFooter_A	CO – IsoConnFooter_A	No (Footer)

*	CO – IsoConnFooter_A	CO – IsoConnFooter_A	No (Footer)
*	*	CO – IsoConnFooter_B	No (Footer)
*	*	CO – IsoConnFooter_B	No (Footer)
*	*	*	*

'*' Indicates the data word isn't part of this specific IsoConn (it could be anything).

Read the rows as a type of timeline: A sending switch only has access to the previous and current rows.

In the above example, the payload being sent to the destination is exactly 83 words, with a 10 word header. The source chose to use 2 words for the *IsoConnInit* sequence, with the *IsoConnWordCount* set to exactly 128 words. Switch 'A' handles a 3 word long *IsoConnRoute* that points to Switch 'B'. Switch 'B' handles a 2 word long *IsoConnRoute* that points to a link heading to the final destination endpoint. An *InBandSignal* is setup for every 4 words, and starts 3 words after the *IsoConnSetInbandSignal* word. Switch 'A' inserts its 3 word footer at the end of the 128 words it sends. Switch 'B' inserts its 2 word footer at the end of the 128 words it sends. In this way, with an IsoConn with 128 *IsoConnWordCount*, all nodes process exactly 128 words. Notice in this example, that the link between Switch 'A' and Switch 'B' lost or corrupted Payload9, and the destination receives *LostWord* instead of Payload9. Also, the link between the source and Switch 'A' lost or corrupted InBandSignal3, and thus Switch 'A', Switch 'B', and the destination won't receive InBandSignal3; receiving *LostWord* instead. This rate of loss isn't expected, but shown merely as an example.

For a slot not allocated to an *IsoConn*:

- A word sent with SIG_EVEN parity indicates that the node is sending data directly to its neighbor using some Link Layer protocol
- A word sent with CONN_ODD parity indicates that the sending node is attempting to allocate the slot for an *IsoConn*.
 - This first CONN_ODD word MUST be interpreted as the first word of an *IsoConnInit* sequence. The sequence MUST be between 1 and 19 words long, depending on the desired precision of the data values.

All attempts to start an *IsoConn* MUST begin with an *IsoConnInit* sequence.

The *IsoConnInit* sequence passes along two values:

1. *IsoConnWordCount*: A floating point value describing the initial maximum number of words to be sent on the *IsoConn* stream.
 - a. Each node MUST remember this value for the *IsoConn* on this slot
 - b. Each node MUST begin counting the words of the *IsoConn* stream with the first word that follows a valid *IsoConnInit* sequence
 - c. This value is used to know when the node should send the *IsoConnFooter* words.
 - d. Minimum expressible value: 128 words
 - e. Maximum expressible value: ~4 zebiwords

- f. If the *IsoConnWordCount* is greater than the allowable maximum, as advertised by the switch, the switch MUST drop the *IsoConn* connection request.
- 2. *IsoConnPaymentCredits*: A floating point value describing the number of credits per word to pay for the *IsoConn* stream.
 - a. The switch receiving this value MUST deterministically subtract the exact amount of credits that covers the node's advertised cost per word of handling the *IsoConn* stream
 - b. If, after the subtraction step, the credits would be negative, the switch MUST drop the *IsoConn* connection request.
 - c. If, after the subtraction step, the value is greater than the advertised amount of credits the switch is willing to transfer per word, the switch MUST drop the *IsoConn* connection request.
 - d. In all other cases, the switch MUST scale the remaining credits appropriately at the advertised scale factor, and MUST forward the resulting value in the *IsoConnPaymentCredits* field of the next *IsoConnInit* sequence over an available output slot in the appropriate outgoing link.

The following table describes the meaning of the bits in these words:

Word	Value	Parity
IsoConnInit1	<p>14:15 (2 bits) (Number of additional <i>IsoConnInit</i> words after IsoConnInit1)</p> <p>OPEN: If the above is 0, use a very compact format for payment and assume initial word count = 1024.</p> <p>A floating point value representing the number of words that are being paid for: 8:13 (6 bit) exponent, biased by 7 (such that the minimum expressible value is 128) 0:7 (8 bit) a 9 bit fraction (with an assumed initial 1 <u>'hidden bit'</u>)</p>	CONN_ODD
IsoConnInit2	<p>A floating point value representing the number of credits to transfer to the neighbor per word of <i>IsoConn</i>. 9:15 (7 bit) exponent, biased by -7 0:8 (9 bit) a 10 bit fraction (with an assumed initial 1 'hidden bit')</p>	CONN_ODD
IsoConnInit3	16 bits of additional precision for credits	CONN_ODD
IsoConnInit4	<p>0:11 (12 bits) additional precision for credits</p> <p>12:15 (4 bits) Number of additional <i>IsoConnInit</i> words after IsoConnInit5</p>	CONN_ODD

IsoConnInit5	16 bits of additional precision for credits	CONN_ODD
IsoConnInit6	16 bits of additional precision for credits	CONN_ODD
IsoConnInit7	16 bits of additional precision for credits	CONN_ODD
IsoConnInit8	16 bits of additional precision for credits	CONN_ODD
IsoConnInit9	16 bits of additional precision for credits	CONN_ODD
IsoConnInit10	16 bits of additional precision for credits	CONN_ODD
IsoConnInit11	16 bits of additional precision for credits	CONN_ODD
IsoConnInit12	16 bits of additional precision for credits	CONN_ODD
IsoConnInit13	16 bits of additional precision for credits	CONN_ODD
IsoConnInit14	16 bits of additional precision for credits	CONN_ODD
IsoConnInit15	16 bits of additional precision for credits	CONN_ODD
IsoConnInit16	16 bits of additional precision for credits	CONN_ODD
IsoConnInit17	16 bits of additional precision for credits	CONN_ODD
IsoConnInit18	0:10 (11 bits) additional precision for credits (for a total of 256 bits of precision) 11:15 - RESERVED - MUST be 0	CONN_ODD
IsoConnInit19	RESERVED MUST be 0	CONN_ODD

TODO: Specify a deterministic subtraction and multiplication method for variable-precision credits.

Each node self-declares the number of words it MAY buffer (ideally 1 word) when transmitting the words across the required outgoing slot. The node MUST advertise this buffer size.

Reasoning for the existence of *IsoConnInit*: The credit value is a required part of every connection request and acts essentially as a credit-based hop-count limit.

7.1 Source Routing

The source provides a series of route instructions to be used at each hop ([Source Routing](#)). After an *IsoConnInit* sequence, the source node MUST stream a series of *IsoConnRoute* words that define the route, one switch at a time. The number of words consumed by each switch and their meaning MUST be defined and advertised by the switch.

A switch MAY advertise a requirement that the *IsoConnRoute* be some number of words longer than the *IsoConnInit* sequence used. This is to allow for low buffering of the *IsoConn*'s words. If the switch accepts an *IsoConnRoute* word(s) that is shorter than the *IsoConnInit* sequence, then the switch MUST implement a ring buffer to make up the difference and MUST queue all the future words belonging to that *IsoConn*, which adds to the advertised latency.

If an *IsoConnRoute* is recognized by a switch, the switch MUST allocate an output slot on the required link to stream the required *IsoConnInit* sequence. After the *IsoConnInit* sequence is streamed to the next node, the switch MUST mark the input slot as an *Active IsoConn*.

Let's say there's 1,000 words of *IsoConnRoute* (500 hops, 2 words each hop)

Effect of 1 word headers on latency at 1Mwps when travelling a 1,000 hop route:

$$(2 \text{ words} / 1,000,000 \text{ wps}) * 1,000 \text{ hops} = 2 \text{ w} / 1,000 \text{ wps} = 2\text{ms}$$

7.2 Active IsoConn

While an *IsoConn* is *Active*, a switch MUST copy each word on the input slot to the allocated output slot

Temporary physical link interference MUST NOT deactivate the *IsoConn*.

When the previously specified *IsoConnWordCount* is reached, the switch MUST send its own *IsoConnFooter* words, matching the exact number (but not the value) of *IsoConnHeader* words that were consumed by this switch in the *IsoConnRoute* sequence that activated this *IsoConn*.

7.3 In-band Signals

An *IsoConnSetInBandSignal* (SIG_EVEN) announces to all switches along a route the initial timing and frequency of the *InBandSignal*. The *IsoConnSetInBandSignal* word defines:

- The timing for the initial *InBandSignal* word
- The separation (in words) between each subsequent *InBandSignal* word

All IsoGrid switches MUST note these two values and decode future *InBandMessages* within the *InBandSignal*. The destination endpoint MUST note these two values such that it can distinguish the *IsoConn* stream words from the *InBandSignal* words. An *InBandMessage* within the *InBandSignal* COULD target a single switch, many switches, or all switches. It doesn't seem useful to have an *InBandMessage* target the destination endpoint; that's what the *IsoConn* stream is for. An *IsoConnSetInBandSignal* that arrives while an existing *InBandSignal* is active MUST be ignored: There can only be one *InBandSignal* within an *IsoConn*. Each *InBandMessage* MUST have a single IsoGrid Header. An *InBandMessagePreamble* word MUST be sent before each *InBandMessage*. *InBandMessage* words MUST have CONN_ODD parity. The *InBandMessagePreamble* word has SIG_EVEN parity. If a switch recognizes an *InBandMessagePreamble* word within the *InBandSignal*, it MUST read the next word of the *InBandSignal* as the first word of the next *InBandMessage*. A valid *InBandSignal* can only begin with a SIG_EVEN word: If the first word of the *InBandSignal* has CONN_ODD parity, the node MUST assume that the *InBandSignal* has ended (or never existed).

Name	Value
ISO_CONN_SET_IN_BAND_SIGNAL_SEWORD	<p>A floating point value describing the period of the <i>InBandSignal</i>:</p> <p>0:2 (3 bit) a 3 bit fraction (assumed initial 1 'hidden bit' if all zeros)</p> <p>3:6 (4 bit) exponent, biased by 3 (such that the minimum expressible value is exactly 1)</p> <p>7:11 (5 bits) exponent, no bias</p> <p>This value describes the number of <i>IsoConn</i> words between this word and the first <i>InBandSignal</i> word</p> <p>A value of zero means that the next word is the first <i>InBandSignal</i> word.</p> <p>12:15 (4 bits) - BITWISE_NOT([0:3] XOR [4:7] XOR [8:11])</p>
IN_BAND_MESSAGE_PREAMBLE_SEWORD	0xFFFF
IN_BAND_SIGNAL_NO_MESSAGE_COWORD	0x0000

All IsoGrid switches MUST support decoding the following *InBandMessages*:

- *SetNewIsoConnWordCount*
- *EndInBandSignalNow*
- *EndInBandSignalFuture*

7.3.1 SetNewIsoConnWordCount

This *InBandMessage* sets a new *IsoConnWordCount*. This message MUST be ignored if it's not immediately preceded by an *InBandMessagePreamble*. This message allows source nodes to start off with smaller *IsoConn* word counts, and then dynamically increase them if they continue to work well. A *SetNewIsoConnWordCount* with a word count shorter than the previous word count MUST be ignored.

All IsoGrid switches MUST recognize 0x9000 as a *ProtocolAlias* of {0x1000, 0xDC30, 0x4396, 0x11E5, 0xB970, 0x0800, 0x200C, 0x9A66} which refers to this *SetNewIsoConnWordCount InBandMessage*. The first word of the 2 word payload MUST be a normal *IsoConnWordCount*, but with an extra two bits for the exponent (total 8), leading to the maximum expressible word count being roughly the number of atoms in the observable universe. NOTE: Perhaps 7 bits for the exponent is enough, and using the extra bit for the fraction is a better use of the bit?

The second word MUST be identical to the first.

TODO: Can this use a better FEC?

7.3.2 EndInBandSignalNow

This *InBandMessage* allows source nodes to end an *InBandSignal*. This message MUST be ignored if it's not immediately preceded by an *InBandMessagePreamble*. The message has no payload. If a switch decodes this *InBandMessage*, it MUST stop decoding this particular *InBandSignal*.

All IsoGrid switches MUST recognize 0x8000 as a *ProtocolAlias* of {0x0000, 0x305F, 0xEFD7, 0x4598, 0x934E, 0x9A0E, 0x5544, 0xD954} which refers to this *EndInBandSignalNow InBandMessage*. This message is fast and efficient, but possibly risky on noisy links.

7.3.3 EndInBandSignalFuture

This *InBandMessage* allows source nodes to set a future end to the *InBandSignal*. This message MUST be ignored if it's not immediately preceded by an *InBandMessagePreamble*. The message has a two word payload. The first word MUST contain a simple 16-bit integer defining the number of words of *InBandSignal* that will follow the end of this *InBandMessage*. The second word MUST duplicate the first, and MUST be ignored otherwise. Using multiple *EndInBandSignalFuture* messages allows for redundancy to best ensure that both endpoints and every switch along the route agree on the precise word where the *InBandSignal* ends.

All IsoGrid switches MUST recognize 0x9001 as a *ProtocolAlias* of {0x8001, 0xEFBD, 0x26E2, 0x4AD7, 0x9C74, 0x19B7, 0xF9E3, 0xEA59} which refers to this *EndInBandSignalFuture InBandMessage*.

8 Packets

The global IsoGrid standard has a practically infinite packet size limit and Maximum Transmission Unit (MTU) of $\sim 2^{136}$ words.

Packets CAN be sent over a normal *IsoConn*: This is referred to as an *IsoPacket*. This method requires that:

1. The sender MUST have an Isochronous Link available over which it can send the *IsoPacket*
2. The sender MUST have available credits to pay for the *IsoPacket* to reach its destination
3. The sender MUST know the full isochronous path to the destination prior to sending the *IsoPacket*

IsoPackets have the same low latency as a normal *IsoConn*.

Packets MAY be sent from node to node using any Link-Layer protocol. Links MAY use Geo-Location routing, offer bounties for successful delivery, or other routing strategies. These methods have larger latency, but don't have the strict *IsoPacket* requirements and might be useful for edge nodes without access to dedicated Isochronous Links (like Internet-of-Things devices and sensors).

9 Economics and Growth

In the early stages of implementing the IsoGrid, there won't be any widely-deployed native-IsoGrid services. The only practical use of the IsoGrid during the early stages would be as a gateway to the existing Internet: The IsoGrid Protocol will need to act as a reasonably inexpensive alternative to traditional Internet Service Providers. Two of the biggest costs of traditional Internet Service Providers are: 1) infrastructure for connectivity over the last-mile to customers, and 2) customer acquisition costs. With a local IsoGrid, last-mile infrastructure is provided by the customer. Customer acquisition costs might be significantly reduced because the IsoGrid is likely to spread from neighbor to neighbor by word-of-mouth, precisely for the purpose of getting cheaper internet access. The early adopters are likely to be willing to start the IsoGrid before financial benefits are clear, due to being dissatisfied with their existing ISP options (or lack thereof). Once a small IsoGrid is started, the connected participants have a strong financial incentive to connect up more of their neighbors.

Once a significant portion of the population of developed countries start using the IsoGrid for Internet access, the hardware and software costs will get much lower. Having a grid topology allows for implementations with very simple initial setup and maintenance. When these begin to appear, the IsoGrid will spread to developing countries, providing cheap, scalable, and dependable connectivity to its users.

9.1 Network Management

Initially, the early adopters of the IsoGrid will have to manage their own networking equipment and handle credit settlement between neighbors. However, over time, we might expect to see the emergence of companies offering "Network Management" services. These netMan services are likely to attempt various business models:

- Consumer leases equipment, netMan service provides flat-rate internet service
- Consumer owns equipment, netMan service handles software management and takes a cut on credit exchanges between neighbors
- Open source
- Etc.

Many of these models might rely on consumer brand loyalty: If a brand of netMan becomes known for good service for the value, it's likely to gain customers from competitors. A brand with security holes is likely to suffer. Flat rate might die out as policing a commons can get expensive.

10 Undefined Higher-Level Services and Protocols

There are a number of services and protocols defined at a higher level that IsoGrid nodes and switches would likely need in order to operate. These services and protocols are intentionally

left out of the global IsoGrid Protocol specification in order to allow the overall system more flexibility over time.

Link Advertisements

Endpoint nodes on the IsoGrid will need to be able to determine routes to destination nodes. The idea here is to have a node advertise to its neighbors (and perhaps to anyone who asks) things like the following:

- A Tag which refers to the below route attributes, and used as an *IsoConnRoute*
- Required number of *IsoConnRouteFiller* words to affix to the *IsoConnRoute* relative to the number of *IsoConnInit* words.
- Geo Location of switch
- Geo Location of switch on the other end of the advertised link
- Scale factor for credit exchange
- Cost (in credits) per word
- Type: On-Demand-Immediate vs. On-Demand-Wait vs. Reservation
 - On-Demand-Immediate: If the connection cannot be switched immediately, it is cancelled
 - Reservation: The connection requires a key procured via a prior exchange
 - On-Demand-Wait: The source is willing to wait for a lower-priority connection to be closed (requires large buffers and complex switches to support this type)
- Target Number of words buffered
- Maximum number of connections that can be started in a single frame
 - The node may be using its link to transmit packets as well as isochronous connections. Specifying this allows it to better optimize its frame buffer allocation
- Timing class { TCG, GPSDO, Quartz, others? }
- Level of IsoGrid standards support
- Latency above a direct point-to-point wireless link (speed of light in a vacuum)
- A maximum *IsoConnWordCount* that the link will allow
- A maximum number of credits the node is willing to transfer per word of *IsoConn*
- Perhaps, link failure insurance? (credits refunded on destination link failure)
- The availability time when this tag is first valid
- The expiration time when this tag is no longer valid

Forward Error Correction Protocols

The IsoGrid is designed to support sending portions of the data across the grid over separate routes. With this in mind, it's good to have more connections to allow for redundancy. But if you split the data evenly, that just increases the likelihood that a link failure will cause data loss. However, if the source and destination agree to use a Forward Error Correcting Code (FEC code), then with just a bit of overhead, the combined stream can be tolerant of link failures. The fact that corrupted words are sent as *LostWord* allows them to be counted as Erasures for the FEC algorithm, which provides more efficient redundancy.

Low Latency Game Streaming

With low latency access to a decentralized network, it's possible to implement game streaming services; where folks share (or rent) game console access from your neighbors. Generalizing, this may evolve into distributed compute.

Decentralized Storage

People will eventually be able to store data in decentralized storage networks, with connectivity provided by The IsoGrid. Data should be both more secure (encrypted at the source) and reliable (distributed very widely, with excellent erasure FEC codes).

Alarm Systems

The ability of neighbors (and perhaps even neighborhoods) to link up their alarm systems can reduce the cost and/or increase the effectiveness of the systems. Additionally, it isn't necessary that the system be centralized, and smaller systems are less of a target for hackers.

IsoGrid Internet Service Providers

The success or failure of IsoGrid basically hinges on whether it's cheaper or less of a hassle to have an IsoGrid-based ISP instead of a Traditional ISP, like Comcast or Verizon.

IsoGrid-based ISPs are referred to as Minimal ISPs (or minISPs) and are very similar to Traditional ISPs except that they don't run links all the way to the customer, instead they rely on a local *IsoGrid* to provide the last-mile connectivity to and from their customers. Only customers that have at least one or (hopefully) more connections to an *IsoGrid* can use a minISP. A minISP must follow all the rules and regulations that apply to ISPs.

Link Tunnels

Tunneling an IsoGrid link across the IsoGrid between distant nodes should be efficient, with low overhead. This could make for quicker, easier, and more reliable long distance *IsoConn* streams. This effectively reduces the hop-count on an *IsoConn* stream that traverses the tunnel.

IP Transit

Tunneling other network protocols on top of the *IsoGrid* should be efficient, with low overhead.

Large Scale, High-Precision Timing

There are a number of large scale projects that aren't practical with the IP Internet, but could be undertaken if the IsoGrid were massively successful. One property of a full-scale IsoGrid is a very precise synchronized time source. With excellent, synchronized clocks it's potentially possible to build:

1. Cheap differential GPS, everywhere
2. Good-signal, Terrestrial GPS in urban areas
3. Indoor GPS
4. Distributed deep-space antenna arrays