

The IsoGrid/IsoMesh Requirements

Version 1.002

The Internet Has:

- High and Unbounded Latency
- Wasteful, Underused Links
- Limited Node/Switch/Hop Counts (IoT)
- Low Redundancy
- A Tendency to Centralize Power
- Choke-point Surveillance and Censorship
- Disaster Vulnerabilities
- Tragedy of the Commons

What follows is a free and open proposal for a solution to these problems: A new open network protocol with a mesh topology.

1 Protocol Requirements

The following are the goals and requirements for a new network protocol with a mesh topology called the IsoGrid.

1.1 Socioeconomic Vision:

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

1.2 Primary Tech Requirements

- Very-low maximum-latency bounds
 - No overcommit, no oversubscribe
 - Always QOS
- Efficiently scales to arbitrarily high bandwidth links
- Efficiently scales to arbitrarily high node/switch count
- Mesh Topology
 - Multi-path redundancy
 - Disaster Resistant
- Seamless connectivity for mobile and "Internet of Things" nodes
- Avoid global protocol mandates that limit economic freedom

1.3 Secondary Tech Requirements:

- Differential GPS everywhere, even indoors?
- Great multi-cast
 - Any switch is allowed to be a multi-caster
- Support for asymmetric links
- Enable high-quality crowd-sourced deep-space antenna arrays
- Enable efficient use of long-haul space-based wireless laser meshes
- IP Tunnels over The IsoGrid should be better than existing non-LAN networks with respect to:
 - Reliability
 - Latency
 - Cost
 - Speed
 - Security

1.4 Non-Goals:

- Does not need to be limited to wireless networks
- Does not need to be simple, or easy
- Does not need to fit into existing hardware
- Does not need to work easily with existing infrastructure
- Does not need to preserve or promote existing power structures
- Does not need to conform to any 'model' of networking

2 Justification: What's wrong with Internet Protocol (IP)?

The goal of IP was simple: Create an interop-protocol to connect the world's networks. In this regard, IP has seen fantastic success. However, it wasn't designed with socioeconomic goals in mind. In this day and age, when much of global commerce seems to rely on IP, it's tempting to think that "Anything can run on IP, why not run the IsoGrid on top of IP?" But this would be akin to asking "Why not build the grid of roads only on top of the hub-and-spoke railway network?" When you think about the goals of the IsoGrid, you'll see that this is a ridiculous proposal. It's reasonable to try to rely on the existing IP infrastructure in the early phases of building the IsoGrid: Like the railways connected the cities prior to interstate highway systems.

High and Unbounded Latency

IP switches have to decode the entire header of a packet and then lookup routing tables before the packet can be routed to the next switch. But it's worse than that, with IP, there are no

guarantees that a given packet will be routed within a certain amount of time, or even serviced at all. If 10 packets arrive on 10 different links at the same time, and all 10 have the same destination link, then some of the packets needs to wait for their turn. If the switch in this case only has buffers for 8 packets, then the last 2 are simply dropped.

The IsoGrid must provide bounded latency, and must provide low latency even as the network scales up.

Wasteful Underused Links

The majority of the links that comprise the Internet run at less than 50% utilization. This is related to the same latency/buffering problem above: In order to have even reasonable latency and low levels of packet-loss, links must typically be at less than 50% utilization.

The IsoGrid must allow for 100% link utilization without any increase in latency on existing connections and without suffering congestive collapse.

Limited Node/Switch/Hop Counts

IP has limits on the number of switches, nodes, and hops that make it ill-suited to an "Internet-of-Everything".

The IsoGrid must have no limits on the number of participating nodes or switches, and routes must be able to have an arbitrarily large number of hops.

Low Redundancy

Typically, most consumers and businesses have only a single link to the IP Internet. This is often because there is only one high-speed provider at a given location. But also, Internet links are mostly paid for by link bandwidth, rather than the actual bandwidth used. It becomes cost-prohibitive to pay for multiple under-utilized links. Finally, IP itself doesn't have good support for multi-path.

The IsoGrid must promote a mesh topology, where it actually makes sense to have more than just one link.

Centralization of Power, and thus Wealth

With the IP Internet, Economies-of-Scale make massive centralized services cheaper than distributed services (even if similarly massive). These centralized services seem to leave very little room for a healthy middle class.

With the IsoGrid, distributed services should be cheaper to provide than centralized services. Distributed services can spread the benefits of a growing economy more widely.

Choke-point Surveillance and Censorship

Because the Internet and the services running on it are so centralized, powerful governmental systems have the clear capability to surveil and/or censor it.

The IsoGrid must scale up to have no Choke-points or Check-Points: You should be able talk with your neighbors without permission from a central authority.

Vulnerable to Disaster

Major damage to a critical building in most major cities is likely to bring down IP Internet service in the region for weeks or months. War, terrorism, earthquakes, or coronal mass ejections could all completely bring down the Internet, knocking out both local and long distance communications, hindering recovery efforts.

The IsoGrid must not rely on central hubs.

Tragedy of the Commons

The Internet is a Commons, where everyone is expected to behave themselves or face removal from the network by network admins. This is expensive to police. The biggest example of this how it costs practically zero for spammers to send comment and email spam.

The IsoGrid must not suffer from Tragedy of the Commons, it should rely on micro-payments in exchange for accepting requests.

2.1 Details

It should be immediately clear to the reader that communication protocols, like TCP/IP, have dramatically changed the world. What is less clear, however, is that the specifics of protocol design can have wide-ranging social and economic effects, some positive, some negative.

Rail Networks are to Road Networks, as the TCP/IP Internet is to the IsoGrid

2.1.1 Economies of Scale

Many networks have a design that reinforces economies of scale. One can see this clearly with the train railroad network: The bigger, more interconnected systems beat out the smaller, less interconnected systems. This also creates huge barriers for new entrants, making it practically impossible for them to compete against the established providers. We submit that economies of scale with the present Internet are creating fewer and fewer providers, and concentrating control in a few individuals in the same way that the rail system of the late 1800s did. This doesn't have to be the case though. The grid topology of our roadways doesn't seem to have these same effects. The barrier-to-entry for personal use of the roadways is much smaller compared to that of the railway providers. How small can we make the barriers to entry in the telecommunications market?

The Railroads lead to Railroad Tycoons

The Internet lead to Internet Tycoons

Where are the Roadway Tycoons?

In the same way, the IsoGrid is designed to be less effective than the Internet at centralizing wealth and power.

2.1.2 Isochronous Grid/Mesh

The "Iso" in IsoGrid is short for Isochronous, meaning 'same time'. Isochronous means that bits are sent (and then arrive) at a specific well-defined frequency. A theoretical isochronous stream running at a frequency of 1 MHz would transmit one word of data every millionth of a second. Implementations of Isochronous protocols can operate with statically sized buffers, and bounded latency guarantees. VoIP, Internet video, and Internet radio are best sent over an Isochronous connection. Over 50% of peak Internet traffic is actually perfectly suited to Isochronous streams. The early POTS telephone network operated with an analog audio stream, and so early digital communications protocols that evolved from this network could be considered Isochronous. However, over time, the world has instead settled on a packetized asynchronous solution, which is now called the Internet.

Being based on packets, the Internet has terrible support for isochronous connections. This is why your YouTube movies always need to 'buffer' before playing: It's sending a few seconds (or more) of the video to the recipient before it starts playing so that it can compensate for the randomly-timed delivery of packets. If the net had support for true isochronous connections, then it would be possible to watch YouTube and Netflix videos without having to wait for the 'buffering' to complete.

2.1.3 Topology

At its core, IP allows a maximum of 255 hops for any packet. This inherently restricts the topology of the Internet: As it stands, it can never be a world-wide mesh. Instead, you end up with large hubs and choke-points.

So far, all attempts to create interop standards for nodes that can hop between networks have been unsuccessful. The IP addressing scheme also makes it very difficult to have a mesh: IP addresses are assigned hierarchically.

The name "Inter-Net" describes the problem directly: The Internet isn't a global network that just anyone can contribute or connect to; instead, the Internet is just a protocol for *inter-*connecting the world's centrally owned and operated *networks*.