AN ENGINEER'S GUIDE TO TMoIP

Richard W. Hoffman III
GDP Space Systems

ABSTRACT

As telemetry transport systems move inexorably closer to a unified telemetry-over-IP approach, the operators and engineers who have traditionally deferred to a separate communications group can benefit from a more comprehensive understanding of the intricacies of the transport medium and protocol. Ethernet, and more specifically IP network hardware, has gained increased robustness, as well as much of the reliability enhancing functionality of more venerable transport solutions, but with these increasingly integrated feature sets comes an emphasized demand on the telemetry systems operator to be able to configure the telemetry transport network devices in more dynamic environments. This paper will seek to serve as a handbook for the telemetry community, guiding discussions of the strengths, weaknesses, legacy, and future outlook of this transport methodology both within and without the groups involved in most range telemetry transport environments.

KEY WORDS

GUI, IP Networking, TMoIP

WHAT IS TMoIP?

Telemetry-over-IP, hereafter referred to as TMoIP, is a method of transporting telemetry data over a network at OSI layer 3, or the network layer, though in practical terms, this could be done at layer 2 as well. Data is acquired, packetized, transmitted, and re-serialized as part of the process of transporting the TM stream.

THE FUTURE OF TMoIP

As IP networks continue to become more ubiquitous in the range environments, the expansion path for the TM transport systems becomes increasingly clear. Regardless of the carrier, the ease-of-use and implementation of IP-based solutions make them attractive transport options. With the continuing implementation of IPv6 networks, there is a trade-off between the learned, ease-of-use of IPv4 and the new benefits of IPv6 such as true point-to-point data streaming, globally routable multicast, etc.

IP NETWORK TOPOLOGY

In describing an IP network, the nodes, interconnections, routes, and data taps are referred to collectively as the network topology. A network's topology, viewed graphically, could also be considered to be represented by a constellation diagram, assuming numerous many-to-many, one-to-many, and one-to-one data paths.

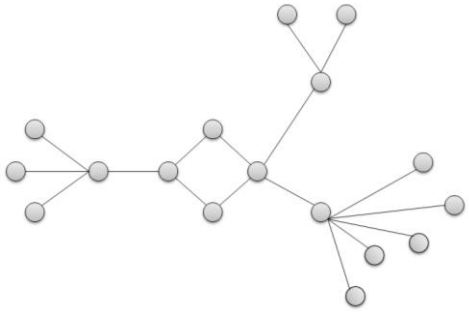

Figure 1- Network Constellation Diagram

SIMPLE TOPOLOGY

In its most simplistic form, an IP network could consist of two nodes, or devices, connected directly. This would be considered a pure point-to-point connection. While this connection would simplify the need to understand some of the more esoteric complications that can arise when using IP networks to transport telemetry data, this is an impractical means of implementing a transport solution.

A simplified, though still functionally practical network would involve the TMoIP devices, separated by one or more ethernet switches, passing data over some indeterminate network, represented quite appropriately ambiguously as a cloud.

Figure 2 - Simple Network Topology

In all practical terms, this type of connection would typically be implemented by using a means of encapsulating the transport layer telemetry stream within a transport protocol for ease of transmission over legacy infrastructure before breaking it back out into the transport layer streams.

CONVENTIONAL TOPOLOGY

A typical IP network will consist of many different interconnected nodes, utilizing a number of intermediate networking routes. These connections are themselves potentially portions of larger networks. Very complex networks can be simplified down to models

that appear to be the same as in the above graphic. Intervening network hardware, which the user may have no direct control over, can be considered to be part of the "cloud".

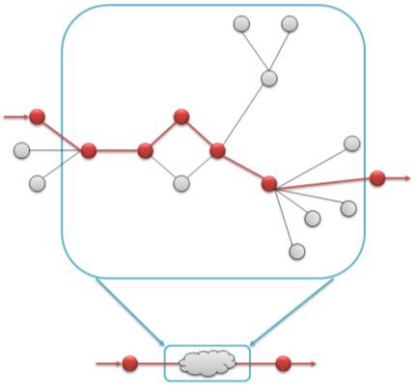

Figure 3 - Network Data Path and Simplification

PROTOCOLS

There are two primary transport layer protocols which could be utilized in order to provide the level of information required by the network to route the telemetry data stream.

TCP

Transmission Control Protocol, or TCP, is a connection based protocol that leverages acknowledgment, sequencing, and error-checking capabilities to ensure maximum data integrity is maintained throughout the transport process.

Delivery acknowledgment is a means of ensuring that previously sent packets have arrived at their destination and works hand-in-hand with the sequencing portion of the protocol, using incrementing values based on the initial sequence number for subsequent acknowledgment packets. The error checking, in the form of a per-packet checksum field in each TCP header, provides the last layer of integrity assurance.

The latency incurred by TCP data transmissions is largely a result of the acknowledgment process and can vary depending on relative endpoint distance, network congestion, routing and prioritization rules, and the imposition of the synchronization mechanism.

In understanding transport overhead encountered while using TCP, it's important to consider that there is a per-packet cost incurred, as well as an impact on the total bandwidth utilization of the transmission. The packet header for TCP consists of 24 bytes with a varying number of bytes allocated to the data payload. In addition to the data overhead, there is an additional bandwidth requirement when there is a need for a packet to be re-transmitted. Further overhead is added in order to facilitate the acknowledgement process.

While the robustness of the protocol may seem appealing, it's important to note that use of TCP is not specified in IRIG218 and as such, will not be discussed beyond this cursory explanation of the protocol.

UDP

User Datagram Protocol (UDP) is a connectionless protocol that trades the advantages of TCP, as enumerated above, for more desirable transmission characteristics in the real-time applications that range telemetry users face. UDP is the de facto protocol for delivery of time sensitive data like voice and telemetry and is the basis upon which higher level transport protocols, like RTP, are constructed.

CHALLENGES

Many of the challenges facing operators of TMoIP systems are inherent to the UDP protocol itself. Where these issues are considered to have a potentially significant impact on the desired functionality of a TMoIP system, a brief discussion of the challenge and a high-level view of a solution are presented below.

NON-ASSURED DELIVERY

Unlike in TCP, there is no assurance that the delivery of transmitted packets was successful when utilizing UDP. While efforts have been made to make UDP more "robust", these efforts are at least somewhat contradictory to the intent and purpose of UDP and do much to re-introduce the time and data overhead of TCP back into the transport process.

In most cases, it is more favorable that a device report lost packets than for it to attempt to initiate a re-transmission of the lost packet. There are numerous reasons that this is undesirable, including an additive delay for each failed packet and an increasingly heavy requirement for buffering and checking at high bitrates. The TMoIP device should implement an informative function by utilizing a proprietary sequence number, inserted into a proprietary TMoIP header at the acquisition site by the source device to allow the receiving device to notify the user when packets fail to arrive, as indicated by an out-of-sequence packet. Utilizing the knowledge that this state has been observed, a user can take steps to attempt to address the network condition that caused it, or mark the relevant records indicating this failure.

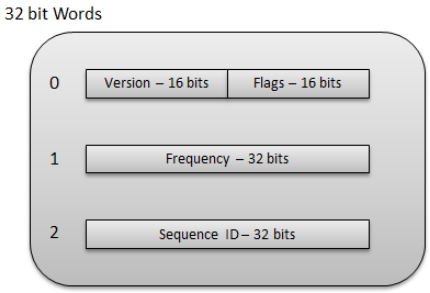

Figure 4 - Example TM Packet Header with Sequence ID

PACKET SEQUENCE

In an IP network, the path which a stream of packets takes on its way to its destination device is non-deterministic, changing to accommodate traffic shaping rules, congested network conditions, and failed network hardware, among other perturbations. In practical terms, this may cause packets to arrive via different routes to the destination, and at different times in relation to previously received packets. This condition presents the opportunity for the same failure behaviors displayed as in the non-assured delivery condition described above. As in the above failure, the appropriate solution is likely not to attempt to re-sequence packets, but to report to the user that this condition has been observed and allow the user to take appropriate action to address this condition. As above, this can be achieved by utilizing a sequence number in a proprietary, low overhead packet header.

LATENCY

The UDP packet has a packet size ranging between the minimal packet header size of 4 bytes and the maximum length field size of 65,507 bytes. This wide range of values for the packet size presents the potential for the packet length to become a factor with regard to latency. This condition can arise when a data rate is low enough in relation to the packet size that it takes a relatively long time for the device to acquire, packetize, and transmit the data. A telemetry data rate of 28 kbps would require almost half a second to fill the payload portion of a 1500 byte packet. This challenge makes the case for an adjustable buffering scheme, whereby the user can specify a packet payload size which will best balance their mission requirements for latency with their network requirements for bandwidth usage, trading time for efficiency.

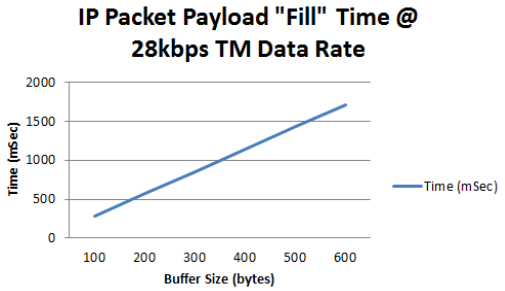

Figure 5 - Graph of time required to fill packet payload

PACKET JITTER

Due to the non-deterministic nature of IP networks, the arrival time of packets, created and transmitted within a relatively consistent period of time, can vary wildly. This variability in the reception of packets of telemetry data requires the receiving device to be able to buffer received data in order to "smooth" any delays in the inter-packet arrival time and eliminate stutters or pauses in the reconstructed data stream. Telemetry hardware typically encountered in the range environments, like bit-syncs, have been known to experience issues with many of the more simplistic clock and data smoothing algorithms implemented by TMoIP device manufacturers. Ideally, the TMoIP device must have an adequately sophisticated algorithm so that the data does not skew outside of the acquisition or tracking range of the bit-sync.

Implementing this buffering necessarily introduces further latency into the system, and should be user-configurable in order to allow for tuning to meet mission requirements.

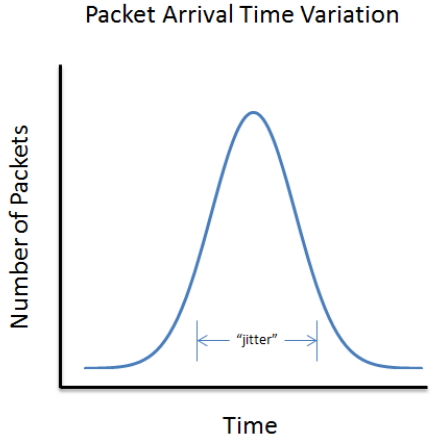

Figure 6 - Example Arrival Time Variation of IP Packets

PACKET SIZE

As mentioned above, a discussion on the implications of packet size is obviously linked to a discussion on the mission requirements for latency, but there are additional

considerations that are tied to the capabilities of the network itself. In developing a TMoIP solution, consideration must be given to the maximum transmission unit, or MTU, of the network hardware. This MTU, in bytes, can potentially impose a tighter restriction on packet size than the user would desire when consideration is focused purely on the efficiency of the IP link. It is absolutely critical that the user fully understands the limitations and configuration of the switches and routers in the network. Additionally, parameters must be provided in the TMoIP device to allow the user the level of control needed to accommodate these limitations.

COLLISION/CONGESTION

Packet collision or congestion in an IP network can occur as a result of inadequate line speed or network device queuing, as well as prioritization rules imposed by network administrators. Congestion can occur over the whole span of the network, but is most commonly encountered in single segments of the network where there are unexpected, unlooked for, and even unintended rules governing network traffic.

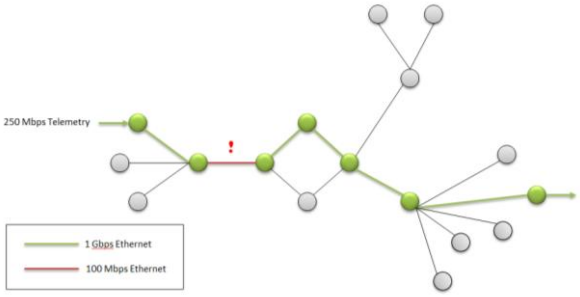

Figure 7 - Network Segment Data Rate Bottleneck

There is little that can be done by the TMoIP device vendor to account for this, but an awareness of the symptoms can assist the user in determining that this condition exists. Most often, this failure manifests itself as a link that "flickers" in and out of sync when tested with a BERT. Additional diagnostic information can be obtained by running a packet capture utility on a workstation and looking for gaps in the packet arrivals that are significantly longer than the "normal" delay between packets. This is likely a result of too many packets being buffered, or a prioritization that is superseding the TM stream. In the latter case, a discussion on DSCP/TOS, detailed below, can assist the operator in mitigating these disruptions.

PROTOCOL OVERHEAD

Protocol overhead is a valid concern for the TMoIP user and becomes a larger concern as the data rates get lower while the latency demands for the mission become tighter. Because UDP allows for an extremely small packet payload, the acquisition and packetization latency can be minimal. In a case where this is required, the payload portion of an IP packet will necessarily be quite small. This configuration presents a

condition where protocol overhead could constitute a very high percentage of the total data transmitted over the network.

Handling this challenge is, oftentimes, going to be a matter of tuning the packet sizes in order to most effectively meet the mission requirements, usually trading latency for efficiency. Understanding the limitations of the network is critical to making this determination, but the crucial considerations are those of line speed and router buffering and will require attention to the network hardware itself. The TMoIP device vendor can, as detailed in the 'Packet Size' section, provide mitigating parameters.

LEVERAGING NON-TM NETWORK STRENGTHS

While the number of complications that could potentially arise from utilizing a TMoIP solution might create the impression that there are other, better options for the telemetry operator, there are equally compelling reasons and solutions that make the case for the IP network based solution to some of the aforementioned challenges.

DSCP/TOS

The differentiated services code point (DSCP) field in the IP packet header allows a network user to "mark" telemetry traffic on a per-device, per-channel, or even per-packet, basis. The values used to mark these streams are then used to create prioritization rules in network hardware that will ensure telemetry traffic is delivered without incurring delays or disruptions due to a high traffic condition on the network segment.

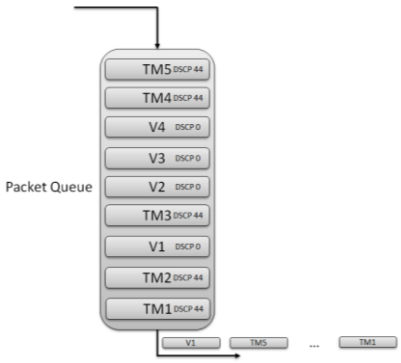

Figure 8 - Example DSCP Priority Queuing

The implementation of the DSCP queuing is specific to the network hardware, but ranges from requiring all packets carrying a specific DSCP value to be sent from the queue before moving to the next DSCP value, to time allotments being given to each DSCP value.

SPANNING/REDUNDANCY

Rapid spanning tree protocol (RSTP) is a method of ensuring that data traversing a network does not have looping paths where broadcast and multicast data could potentially loop perpetually and establishes a single path between any two network nodes. RSTP also allows for the user to provision redundant data paths to ensure that a loss of a single branch or node of the tree does not cause a loss of data on the receiving end.

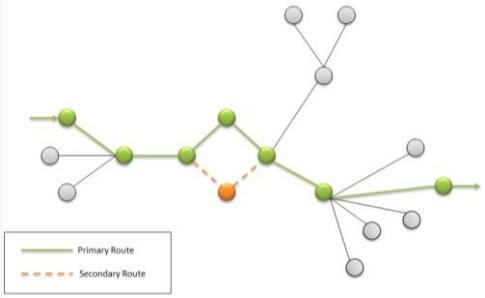

Figure 9 - Spanning Tree Showing Redundant Path

MULTICAST

Multicast is an advancement of the broadcast transmission of packets. It cuts down on the network load by introducing the concept that groups of devices on which the user is interested in receiving the transmitted packets will inform the network that they should receive a "copy" of that data stream. The receiving TMoIP device must support IGMP multicast "group join" requests, but the transmitting device need not do anything other than address the appropriate address range. Multicast is typically encountered in a one-to-many configuration where the TMoIP device is used to perform a similar function to what would traditionally be left to a matrix switch or distribution device.

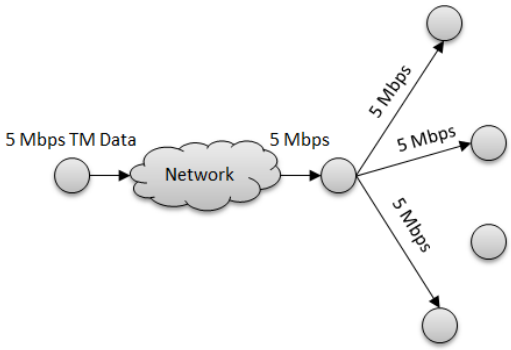

Figure 10 - Multicast One-to-Many TM Path

SECURITY/IPSEC

Internet protocol security, or IPSec, is a collection of protocols that allow a user to perform packet-level encryption. This encryption can be carried out between two TMoIP devices on the same network, a TMoIP device and a routing device, or two routing

devices. The benefit to this method of encrypting the data is that the device does not necessarily need to be burdened with the process, which would add latency and processing overhead. Additionally, the need for expensive additional hardware is removed; depending upon the implementation, the encryption suites can add very little cost to the solution.

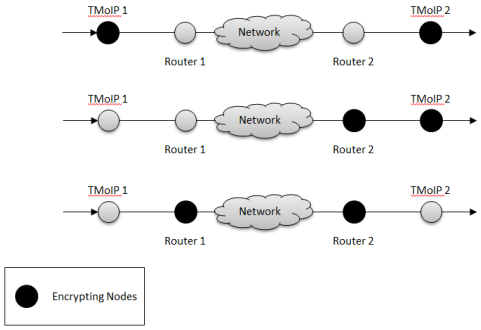

Figure 11- IPSec Encryption at Alternative Nodes

CONCLUSION

Implementing a TMoIP solution to a telemetry transport requirement presents opportunities to incrementally migrate to a more modern telemetry backbone in existing range environments. While the proper user of a TMoIP device will invariably require a deeper understanding of concepts traditionally outside the domain of the telemetry operator, this paper strives to present these principles in an understandable manner. While greater collaboration between telemetry operators, network administrators, information assurance teams, and information security groups will almost certainly aid in a smoother implementation of a TMoIP solution, an understanding of the concepts as presented herein will aid the user in more simply navigating the challenges they will encounter.