Mapping and Transport Standard for OTU4

Mark Loyd Jones, Ph.D.
Tellabs
mark.jones@tellabs.com

Abstract: Recently revised ITU-T G.709 defines a container for 100 Gbps signals called an OTU4. This paper highlights some of the reasons the OTU4 bit rate, frame structure, and mapping scheme were chosen.

1. Introduction

How did the OTU4 (Optical Channel Transport Unit 4) transport standard come to be? It was triggered by the IEEE project to develop a standard Ethernet signal at 100 Gbps (100G for short). Network operators with optical networks wanted to integrate those 100G Ethernet signals into their OTN (Optical Transport Network) dense wavelength division multiplexing (DWDM) networks carrying 2.5G, 10G, and 40G channels already. This paper explains some of the choices made during the development of the OTU4 standard, now included in ITU-T G.709[1].

2. OTN Background

One must first understand a little bit about the OTN to understand the decisions required for defining an OTU4 signal. The OTN is an ITU-T standardized multiplexing and transport hierarchy for optical networks. The OTN Operations, Administration, and Maintenance (OAM) functionality, for optical wavelengths as well as the digital signals they contain, enables network operators to efficiently manage their optical networks. DWDM based optical networks have optical channels that contain digital signals, such as the OTN OTUk. The OTN OTUk is the complete digital signal (level k) on a wavelength, including OTUk Overhead, space for an Optical channel Data Unit (ODUk) of the same level, and a forward error correction (FEC) block. The OTUk frame consists of four rows of 4080 bytes in each row, independent of the level k. Each ODUk contains an Optical channel Payload Unit (OPUk), which carries one or more lower level ODUks or a client signal (see Fig. 1). (For simplicity, this paper will not differentiate between the ODUk and OPUk going forward, given they are always one for one. Payload signals, including lower level ODUks and other clients, are always mapped into the OPUk, which is mapped into the ODUk.)

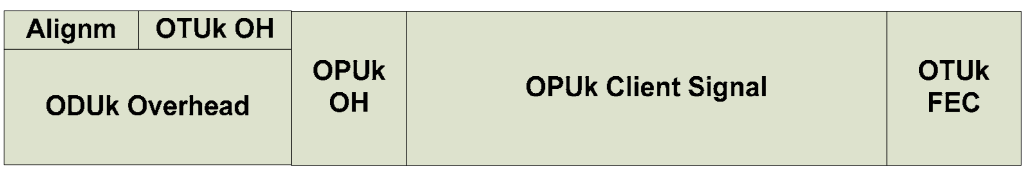

Fig. 1. OTN OTUk Frame and Contents (not to scale)

One could presume that development of an OTU4 standard was a straight-forward effort to simply apply the same tools from the lower OTU levels to a higher rate signal. In a sense that was true. One required a precise bit rate, a structure to carry the client signals, and a mapping for adapting the client signals into the container. However, the details behind those three basic aspects were not easy to come by.

3. OTU4 Bit Rate

First off, an OTU4 required a bit rate. Each higher rate of the OTN hierarchy was nominally four times the bit rate of the next lower signal (OTU1 at ~2.5G, OTU2 at ~10G, OTU3 at ~40G), sized originally to carry the SONET/SDH signals as clients. To maintain that growth trend, one would expect an OTU4 bit rate nominally at 160G. However, the client signal driving the new rate was only at 100G and the technology to support a signal closer to 100G was far more available and economical than technology to support a signal at 160G. The likelihood of supporting a similar transmission span distance as 10G and 40G DWDM was also far better at 100G than at 160G. 120G was also considered as a possibility, given that it would support three ODU3s. Even at 120G, the cost, complexity, and impact to the span distance for carrying the additional 20G was considered more than it was worth, and agreement was reached to select a rate closer to 100G. One could still support two 40G signals or ten 10G signals in it. At the lower rates, the transparent transport of the 10GBASE-R signal had triggered industry deployment of over-clocked ODU2 and ODU3 signals (often named ODU2e and ODU3e), since the 10GBASE-R signal rate was higher than the payload size of a standard ODU2. There was a great desire among the ITU-T developers to define the OTU4 in such a way that there would be no necessity for an over-clocked ODU4. To address that objective, the OTU4 had to be sized sufficiently large enough to carry ten ODU2e signals. One additional requirement for the OTU4 bit rate was to have it easily derivable from the SONET/SDH clock, which was already used as the basis for all of the other OTUk bit rates. To address all of these requirements, an OTU4 bit rate of approximately 111.8G was chosen (10 times 255/227 times the bit rate of the 10G SONET/SDH signals, STS-192/STM-64).

4. OTU4 Frame Structure

Secondly, the OTU4 signal required a structure to contain client signals. As part of the OTN hierarchy, the OTU4 was defined with the same frame structure as existing OTN containers.

However, given the higher bit rate signal at nearly 112G, it was agreed that the OTU4 FEC would be mandatory instead of optional (like it is at the lower rates). Proposals for many enhanced FEC codes were received and debated, but in the end it was agreed that the FEC used at lower rates was sufficient for the inter-domain interface requirements and the added complexity and latency of the higher gain codes was not sufficient to justify using one of them. So, the same FEC code that is used at lower rates was specified for OTU4.

In addition to the frame and FEC definition, a structure for containing the different payload signals was required. Though the OTUk frame is the same size at every level, those frames are grouped together into multi-frame structures, containing a larger number of frames at each progressively higher OTU rate. Tributary Slots (TSs) are defined within that structureand are used to contain the payloads. (Time slots in SONET/SDH are a similar construct.) At the same time that the OTU4 frame structure was being developed, the ITU-T experts agreed to a new lower speed ODUk named ODU0 (nominally at 1.25G), sized to carry 1 Gigabit Ethernet (and capable of carrying other low speed clients signals as well). To maintain simplicity in mapping ODU0 and other clients into the higher rate ODUs, it was decided that all ODUk TS structures would be revised to support 1.25G tributary slots instead of only the traditional 2.5G TSs. (Two of the 1.25G TSs fit within the same space as the original 2.5G TSs.) As a result, the OTU4 was divided into 80 1.25G TSs.

The OTN TS structure is similar at each ODU rate, though the TS sizes are graduated so that the higher speed ODUs have sufficient room for the lower speed ODUs to nest inside of them along with any additional overhead required for mapping one into the other. This was also true at the OTU4 rate, making it possible to fit an ODU2e into only 8 TSs at OTU4, even though an ODU2e requires 9 TSs at the OTU3 rate.

5. OTU4 Clients and Mappings

Thirdly, one must have a way to map the desired clients into the OTU4 structure. Known clients for OTU4 included the 100G Ethernet signal and the lower rate ODUk signals (see Fig. 2). With the continued adoption of Ethernet signals into transport networks and the desire to accept the "odd rate" ODU2e into the OTN hierarchy, there was sufficient motivation to come up with a flexible mapping scheme for mapping any client signal into any larger container, instead of continuing the trend of developing similar but unique mappings for each adaptation. Just such a mapping scheme had been sitting on the G.709 Living List for years, and experts began debating mapping schemes of that sort to select a Generic Mapping Procedure (GMP). The debate split the experts into two camps for months, but finally a single GMP mapping scheme was chosen. The GMP mapping selected is an algorithmic scheme that distributes data and stuff bytes evenly through the payload area of the container.

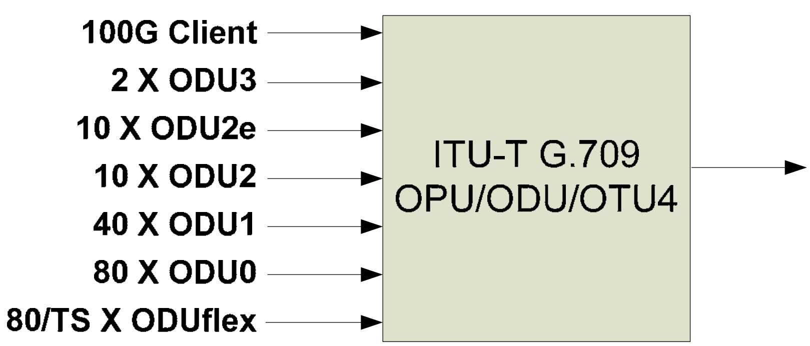

Fig. 2. OTU4 Client List

With such a flexible mapping scheme, it was easy to agree on proposals to also define an ODUk with variable bit rate, given the name ODUflex. The ODU2e, with its size that did not fit the original ODUk hierarchy rates could be treated as the first ODUflex, in the way it is multiplexed into higher ODUks. The ODUflex structure enabled the user to set its bit rate to fit the desired client signal and then only manage a single ODUflex when transporting it. This should be much simpler to manage than the timeslot concatenation scheme used in SONET/SDH, where one must manage the individual members (which could be many) as the concatenated group is transported. The definition of an ODUflex also meant that new client signals that would come up in the future could be efficiently mapped into the OTN without requiring the definition of a new ODUk type (at least until higher bit rate signals are required).

All clients, including the lower level ODUks and other client signals are mapped into ODU4 using the newly developed GMP.

To maximize the backwards compatibility and minimize the need for more interfaces (among other things), it was agreed that the ODU0, ODU2e, and ODUflex must be multiplexed into a higher level ODUk to be transported over an OTUk and an optical channel. The only standard OTN OTUks are the long-standing OTU1, OTU2, OTU3, and the new OTU4.

6. OTU4 Standard Development

The OTU4 standard was not developed as a stand-alone activity, but as part of the work on a major ITU-T G.709 revision. As such, it was interrelated to the other OTN enhancements made, such as the definition of new ODU0 and ODUflex and the development of GMP. This overview explains a few of the many decisions that were reached as a result of the hard work of many people in the ITU-T Question 11 – "Signal Structures, Interfaces and Interworking for Transport Networks" (technical subcommittee of ITU-T Study Group 15 - Optical Transport Networks and Access Network Infrastructures"). The revised ITU-T G.709, including the full specification for OTU4 will be available for free download, after it is approved, from http://www.itu.int/rec/T-REC-G.709/en.

7. References

[1] Draft Revised ITU-T Recommendation G.709/Y.1331 "Interfaces for the Optical Transport Network (OTN)," submitted to ITU-T approval process from the October 2009 SG15 meeting.