TICTOC S. Rodrigues
Internet-Draft Zarlink Semiconductor
Intended status: Informational K. Lindqvist
Expires: January 6, 2010 Netnod
July 5, 2009

TICTOC Requirement
draft-ietf-tictoc-requirements-00.txt

Status of this Memo

Copyright Notice

Abstract

Distribution of high precision time and frequency over the Internet and special purpose IP networks is becoming more and more needed as

IP networks replace legacy networks and as new applications with need for frequency and time are developed on the Internet. The IETF formed the TICTOC working group to address the problem and perform an analysis on existing solutions and the needs. This document summarizes application needs, as described and agreed on at an TICTOC interim meeting held in Paris from June 16 to 18, 2008.

Table of Contents

1. Introduction

There is an emerging need to distribute highly accurate time and frequency information over IP and over MPLS packet switched networks (PSNs). In this draft, the requirements for transporting accurate time and/or frequency are addressed.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119]

3. Applications Requirements

There are many applications that need synchronization. Some applications only need frequency; for others a combination of frequency and time of day or phase may be required. At the TICTOC interim meeting, it was agreed that these applications be grouped based on what was believed to be common requirements, and where the requirements where distinct from each other. This section describes these applications (or groups of applications) that was agreed on at the TICTOC interim meeting.

3.1. Cellular Backhauling

Within Cellular backhauling, there are several applications that need to be considered. Some of these applications only require frequency information, others require time-of-day, and others require phase. The cellular backhauling applications to be considered are:

- GSM
- Mobile Wimax
- LTE
- UMTS FDD
- UMTS TDD
- CDMA2000
- TD-SCDMA

Conventionally GSM and UMTS FDD base stations obtain this reference

frequency by locking on to the E1/T1 that links them to the base station controller. With the replacement of TDM links with Packet Switched Networks (PSNs) such as Ethernet, IP or MPLS, this simple method of providing a frequency reference is lost, and frequency information must be made available in some other way.

The synchronization requirement is derived from the need for the radio frequencies to be accurate. Radio spectrum is a limited and valuable commodity that needs to be used as efficiently as possible. In GSM, transmission frequencies are allocated to a given cellular base station and its neighbours in such fashion as to ensure that they do not interfere with each other. If the radio network designer cannot rely on the accuracy of these frequencies, the spacing between the frequencies used by neighbouring sites must be increased, with significant economic consequences.

There is an additional requirement derived from the need for smooth handover when a mobile station crosses from one cell to another. If the radio system designer can not guarantee that the preparations required for handover occur in a few milliseconds, then they must allow the mobile station to consume frequency resources simultaneously in both cells in order to avoid service disruption. The preparations required involve agreement between the mobile and base stations on the new frequencies and time offsets; these agreements can be accomplished quickly when the two base stations' frequency references are the same to a high degree of accuracy.

3.1.1. Cellular Backhauling

The requirements for the "Cellular Backhauling" are depicted in the following sections.

3.1.1.1. GSM/UMTS FDD

The requirements for the GSM/UMTS FDD are as follows:

Synchronization Type (e.g. time, frequency or phase): frequency
Frequency stability: 50-250 ppb (1)
Frequency accuracy : 50-250 ppb (1)
Uncalibrated time/time stability:
Uncalibrated time/time accuracy: NA
Stabilization Time : As soon as possible
Jitter on recovered timing signal: Depends on the oscillator stability
Wander on recovered timing signal: Depends on the oscillator stability
What expected network characteristics
(WAN, LAN, MAN, private, public, etc)?:
Does the application require security? (if so, which one:
authentication, encryption, traceability, others): No (2)
Reliability requirements (e.g. fault tolerance):
Traceability to a specific clock, clock quality, path, time:
Holdover requirement:
Cost (consumer, enterprise, carrier):
Auto-configuration (plug and play): No
Manageability (how much effort the operator needs to put in to manage
this application?) - In-band or out-of-band of protocol (MIBs?):
Scale and scalability:

Note (1) This is requirement in the air interface. In practice more accurate frequency is required at the input. For example OBSAI RP1 defines 16 ppb

Note (2) assumes a private network

3.1.1.2. UMTS TDD

The requirements for the UMTS TDD are as follows:

Synchronization Type (e.g. time, frequency or phase): phase alignment
Frequency stability: 50-250 ppb (1)
Frequency accuracy : 50-250 ppb (1)
Uncalibrated time/time stability:
Uncalibrated time/time accuracy: The phase alignment of neighbouring base stations shall be within 2.5us
Stabilization Time : As soon as possible
Jitter on recovered timing signal: Depends on the oscillator stability
Wander on recovered timing signal: Depends on the oscillator stability
What expected network characteristics (WAN, LAN, MAN, private, public, etc)?:
Does the application require security? (if so, which one: authentication, encryption, traceability, others): No (2)
Reliability requirements (e.g. fault tolerance):
Traceability to a specific clock, clock quality, path, time:
Holdover requirement:
Cost (consumer, enterprise, carrier):
Auto-configuration (plug and play): No
Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?):
Scale and scalability:

Note (1) This is requirement in the air interface. In practice more accurate frequency is required at the input. For example OBSAI RP1 defines 16 ppb

Note (2) assumes a private network

3.1.1.3. Mobile Wimax

The requirements for the Mobile Wimax (1) are as follows:

Synchronization Type (e.g. time, frequency or phase): phase alignment
Frequency stability: 15 ppb
Frequency accuracy : 15 ppb
Uncalibrated time/time stability:
Uncalibrated time/time accuracy: The phase alignment of neighbouring base stations shall be within 1us
Stabilization Time : As soon as possible
Jitter on recovered timing signal:
Wander on recovered timing signal:
What expected network characteristics (WAN, LAN, MAN, private, public, etc)?:
Does the application require security? (if so, which one: authentication, encryption, traceability, others): No (2)
Reliability requirements (e.g. fault tolerance):
Traceability to a specific clock, clock quality, path, time:
Holdover requirement:
Cost (consumer, enterprise, carrier):
Auto-configuration (plug and play): No
Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?):
Scale and scalability:

Note (1) 1024 OFDM carriers, BW 10 MHz, Cyclic prefix ratio 1:8, RF carrier 3.5 GHz

Note (2) assumes a private network

3.1.1.4. LTE

The requirements for the LTE are as follows:

Synchronization Type (e.g. time, frequency or phase): phase alignment
Frequency stability: 50-250 ppb (1)
Frequency accuracy : 50-250 ppb (1)
Uncalibrated time/time stability:
Uncalibrated time/time accuracy: From 1us to 50us (2, 3)
Stabilization Time : As soon as possible
Jitter on recovered timing signal: Depends on the oscillator stability
Wander on recovered timing signal: Depends on the oscillator stability
What expected network characteristics (WAN, LAN, MAN, private, public, etc)?:
Does the application require security? (if so, which one: authentication, encryption, traceability, others): No (4)
Reliability requirements (e.g. fault tolerance):
Traceability to a specific clock, clock quality, path, time:
Holdover requirement:
Cost (consumer, enterprise, carrier):
Auto-configuration (plug and play): No
Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?):
Scale and scalability:

Note (1) This is requirement in the air interface. In practice more accurate frequency is required at the input. For example OBSAI RP1 defines 16 ppb

Note (2) : no precise phase accuracy requirements defined in standard. The actual requirement will depend on implementation and network scenario.

Note (3) : In general LTE TDD systems may be defined to operate with 10-50 microseconds phase accuracy by making some limitations on the deployment (e.g. cell range), and radio frame configuration, however further investigations are required. When no assumption possible, microsecond or sub-microsecond requirement would apply.

Note (4) assumes a private network

3.1.1.5. CDMA2000

The requirements for the CDMA2000 are as follows:

Synchronization Type (e.g. time, frequency or phase): phase alignment
Frequency stability: 50-250 ppb (1)
Frequency accuracy : 50-250 ppb (1)
Uncalibrated time/time stability:
Uncalibrated time/time accuracy: The pilot time alignment error should be less than 3us and shall be less than 10us(compared to system time)
Stabilization Time : As soon as possible
Jitter on recovered timing signal: Depends on the oscillator stability
Wander on recovered timing signal: Depends on the oscillator stability
What expected network characteristics (WAN, LAN, MAN, private, public, etc)?:
Does the application require security? (if so, which one: authentication, encryption, traceability, others): No (2)
Reliability requirements (e.g. fault tolerance):
Traceability to a specific clock, clock quality, path, time: System Time, synchronous to UTC time (except for leap seconds) and uses the same time origin as GPS time. (3)
Holdover requirement:
Cost (consumer, enterprise, carrier):
Auto-configuration (plug and play): No
Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?):
Scale and scalability:

Note (1) This is requirement in the air interface. In practice more accurate frequency is required at the input. For example OBSAI RP1 defines 16 ppb

Note (2) assumes a private network

Note (3) 3GPP2, C.S0010-B version 2.0, 2004

3.1.1.6. TD-SCDMA

The requirements for the TD-SCDMA are as follows:

Synchronization Type (e.g. time, frequency or phase): phase alignment
Frequency stability: 50-250 ppb (1)
Frequency accuracy : 50-250 ppb (1)
Uncalibrated time/time stability:
Uncalibrated time/time accuracy: The phase alignment of neighbouring base stations shall be within 3us
Stabilization Time : As soon as possible
Jitter on recovered timing signal: Depends on the oscillator stability
Wander on recovered timing signal: Depends on the oscillator stability
What expected network characteristics (WAN, LAN, MAN, private, public, etc)?:
Does the application require security? (if so, which one: authentication, encryption, traceability, others): No (2)
Reliability requirements (e.g. fault tolerance):
Traceability to a specific clock, clock quality, path, time:
Holdover requirement:
Cost (consumer, enterprise, carrier):
Auto-configuration (plug and play): No
Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?):
Scale and scalability:

Note (1) This is requirement in the air interface. In practice more accurate frequency is required at the input. For example OBSAI RP1 defines 16 ppb

Note (2) assumes a private network

3.1.2. Cellular Backhaul Requirements Summary

Based on the sections above, the following can be summarized for the cellular backhaul applications. Two families of technologies can be identified:

- o Those only requiring the recovery of an accurate and stable frequency synchronization signal as a reference for the radio signal (e.g. GSM, UMTS FDD, LTE FDD). In this case the requirement ranges between 15 ppb (Wimax) and 250 ppb (UMTS Home Base Stations). This requirement is applicable on the air interface. In practice more accurate frequency on the long term is required at the input of the Base Stations (e.g. 16 ppb might be required for applications operating with 50 ppb)

- o Mobile technologies that in addition to frequency synchronization, also need phase synchronization. This is the case for the TDD technologies such as UMTS TDD. The requirement in this case ranges between 2.5 microseconds (phase error between Base Stations) to several tens of microseconds that could be sufficient

for some LTE TDD configurations. There is also a case (CDMA2000) that in addition to phase synchronization also requires the distribution of accurate time of day (3 microseconds max error during normal operation).

3.2. Circuit Emulation

The PWE3 WG has produced three techniques for emulating traditional low-rate (E1, T1, E3, T3) TDM services over PSNs, namely SAToP [RFC4553], CESoPSN [RFC5086], and TDMoIP [RFC5087]. The Network Synchronization reference model and deployment scenarios for emulation of TDM services have been described in [RFC4197], Section 4.3. The major technical challenge for TDM pseudowires is the accuracy of its clock recovery.

TDM network standards for timing accuracy and stability are extremely demanding. These requirements are not capriciously dictated by standards bodies, rather they are critical to the proper functioning of a high-speed TDM network. Consider a TDM receiver utilizing its own clock when converting the physical signal back into a bit-stream. If the receive clock runs at precisely the same rate as the source clock, then the receiver need only determine the optimal sampling phase. However, with any mismatch of clock rates, no matter how small, bit slips will eventually occur. For example, if the receive clock is slower than the source clock by one part per million (ppm), then the receiver will output 999,999 bits for every 1,000,000 bits sent, thus deleting one bit. Similarly, if the receive clock is faster than the source clock by one part per billion (ppb), the receiver will insert a spurious bit every billion bits. One bit slip every million bits may seem acceptable at first glance, but translates to a catastrophic two errors per second for a 2 Mb/s E1 signal. ITU-T recommendations permit a few bit slips per day for a low-rate 64 kb/s channel, but strive to prohibit bit slips entirely for higher-rate TDM signals.

3.2.1. Circuit Emulation Requirements

The requirements for the Circuit Emulation are as follows:

Synchronization Type (e.g. time, frequency or phase): frequency
Frequency stability: NA
Frequency accuracy : NA
Uncalibrated time/time stability: NA
Uncalibrated time/time accuracy: NA
Stabilization Time :
Jitter on recovered timing signal: G.8261/G.823/G.824
Wander on recovered timing signal: G.8261/G.823/G.824
What expected network characteristics (WAN, LAN, MAN, private, public, etc)?:
Does the application require security? (if so, which one: authentication, encryption, traceability, others): No
Reliability requirements (e.g. fault tolerance): NA
Traceability to a specific clock, clock quality, path, time:
Holdover requirement: Yes
Cost (consumer, enterprise, carrier):
Auto-configuration (plug and play): No
Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?):
Scale and scalability:

3.2.2. Circuit Emulation Requirements Summary

The Circuit Emulation application requires the receiver to recover the same long term frequency accuracy as the original TDM signal. The phase noise (jitter and wander) of the recovered signal has to be limited according to the relevant ITU-T recommendation (e.g. G.823). There are no requirements on phase or time synchronization in this case.

3.3. Test and Measurement

Note: The application information and the requirements for this section was provided by the LXI Consortium Technical Committee.

In the test and measurement sector there is a desire to move from special purpose communications infrastructure with calibrated wiring run back to a centralize controller, to a distributed system, in which instructions are distributed in advance to be executed at a predetermined time, and in which measurements are taken remotely and communicated back to a common point for later correlation and analysis.

Test and Measurement (T&M) is a very diverse industry and as would be expected, requirements vary widely with the application. However the vast majority of the newer instruments and systems make use of LAN technology and many have a connection to the local enterprise network for data transfer, or monitoring and control.

Because of the increasingly heavy use of LAN technology in T&M instruments and systems, we are dependent on the availability of network infrastructure, e.g. bridges, and low level silicon, e.g. PHYs and PHY/MAC, that supports not only T&M connectivity (data transport) but increasingly timing and frequency transfer support as well.

Furthermore T&M is going to require this support not only for the existing 10/100/1000 BaseT technology but on the newer high throughput LAN technology under development. While most instruments produce data at modest rates, many can source or sink data at rates well in excess of 40Gsamples/s. In addition, the time and phase coherence requirements on the data transport, e.g. LAN, typically are tighter on the high data rate instruments.

The other major headache in the use of LAN in T&M is latency and jitter because it compromises the determinism needed for some applications. One of the promises of LAN-based precise time is that in many circumstances precise time can be used to overcome latency issues. For example, for many data acquisition applications the ability to precisely and accurately timestamp data at the collection point makes LAN latency and jitter a non-issue.

Many T&M applications are localized, often to a bench or rack of equipment. The LAN will be local and private although there is often a connection to the local enterprise network. It is not uncommon in such applications to include a rubidium oscillator to provide a phase-coherent stable frequency source to critical instrumentation such as counters, scopes, signal generators and analyzers. In many cases the LAN, in principle, could fill the frequency distribution role if the LAN technology supported it. In these systems time transfer is becoming important first for timestamping data to facilitate data management and post acquisition processing, and in some cases as part of the control structure. The precise time specifications vary from milliseconds for general applications to nanoseconds for the most critical.

There is an important class of applications where time, and sometimes frequency traceable to international standards is required, generally due to regulatory issues, e.g. testing of medical, safety critical or military devices. The ability to deliver traceable time and frequency over the network to the enterprise would be a big help in these applications.

There are also T&M applications that are widely distributed due to the nature of the device or system being measured. Environmental measurement systems, surveillance, SCADA systems, and the telecommunication system itself are examples. Timestamping data is

an essential requirement to overcome the communication latency and jitter issues. The specific timing requirements clearly cover a wide range. Environmental and SCADA is typically a ms. However to really instrument a telecom system will require timing at least on the order of a packet length or better. Even more extreme are timing for RF test ranges (which can cover several miles), long-baseline interferometry, and RF surveillance where the time accuracy must be on the order of ns. In some cases public networks will be used if the time distribution is adequate.

3.3.1. Test and Measurement Requirements

The requirements for the Test and Measurement are summarized in this section. Where appropriate both the low and high end of the requirements spectrum are given to illustrate the breadth of requirements for the application areas discussed. Note that typically the applications with the most demanding requirements are also the high dollar value applications and in many cases the most critical in terms of the cost of failure, e.g. failure of a surveillance system, monitoring of telecommunications, military test systems where either the operational cost of downtime or the cost of the device being tested is high.

The requirements for the Test and Measurement are as follows:

Synchronization Type (e.g. time, frequency or phase): Time: ms to ns - high value applications will open up as this spec improves. Frequency: part in 10^9 minimum, 10^11 desirable and with the lowest phase noise obtainable for critical applications.

Frequency stability: When applicable (high end RF) the lowest phase noise possible in the short term, long term consistent with accuracy and calibration intervals- better than 1 ppm/year desirable.

Frequency accuracy : Generally consistency across the system is more important than absolute accuracy. For calibration applications at least 1ppm.

Uncalibrated time/time stability: Short term from fractional ms to ns or better. Long term comparable to GPS distributed time .

Uncalibrated time/time accuracy: Usually self-consistency requirements are tighter: ms to ns system wide. Absolute accuracy (traceable) is probably ms to 100 ns.

Stabilization Time : Not usually important. Many times critical instruments themselves need minutes to hours to stabilize. However stabilization times greater than a few minutes will reduce the number of practical wide-area applications.

Jitter on recovered timing signal: In the most critical applications,

the lowest phase noise achievable, in terms of TIE less than the stability requirement.

Wander on recovered timing signal: Modest for most measurements. For surveillance, long baseline, and similar less than the required stability over the duration of the test.

What expected network characteristics (WAN, LAN, MAN, private, public, etc)?: Most are private or enterprise LAN. Large scale applications will benefit from using the public telecommunications networks.

Does the application require security? (if so, which one: authentication, encryption, traceability, others): To date timing security requirements have been rare with the possible exception of measurement systems with legal requirements. Data security is more important when the public networks are involved.

Reliability requirements (e.g. fault tolerance): Has not been an issue to date in most systems.

Traceability to a specific clock, clock quality, path, time: Traceability to a path means that if there is on-path support we want to trace the path. Can also help to avoid time loops. Traceability is needed to establish NIST traceability. T&M will expect that public networks solve the timing loop problem. T&M end systems are typically strictly hierarchical networks without multiple paths.

Holdover requirement: Has not been an issue to date- but as T&M increasingly is integrated into operational systems it will become more important. Telecom requirements are probably sufficient.

Cost (consumer, enterprise, carrier): In most T&M systems component cost is very important. In many, operational cost is important.

Auto-configuration (plug and play): Very important. T&M customers to date would prefer to avoid any network related configuration.

Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?): As little as possible operator interaction. However visibility into system performance, including timing, is very important both operationally and during debug and commissioning.

Scale and scalability: T&M systems range from small systems with perhaps 2 or 3 instruments to large scale data acquisition with thousands of end devices. The physical scale of T&M systems varies widely from a few instruments on a bench to a few instruments separated by miles, and from several thousand instruments and sensors concentrated on a local device such as a jet engine to several thousand spread over many miles in environmental monitoring, or monitoring the telecommunications system. In all cases it is very common for these systems to grow as

additional test requirements are imposed so scalability is important.

3.4. Industrial Automation

In the industrial sector there is a desire to move from special purpose communications infrastructure with calibrated wiring run back to a centralize controller, to a distributed system. One example of this tendency is described below.

In the printing industry there is a need to control operations in multi-stand printing machines. The paper travels through these machines at a speed of nearly 100 km/h. At these speeds, co-ordination error of 1 microsecond between operations taking place at different positions in the machine produces a 0.03mm color offset, which is visible to the naked eye and results in an unacceptable degradation in quality.

3.4.1. Industrial Automation Requirements

The requirements for the Industrial Automation are summarized as follows.

Synchronization Type (e.g. time, frequency or phase):
Frequency stability:
Frequency accuracy :
Uncalibrated time/time stability:
Uncalibrated time/time accuracy:
Stabilization Time :
Jitter on recovered timing signal:
Wander on recovered timing signal:
What expected network characteristics (WAN, LAN, MAN, private, public, etc)?:
Does the application require security? (if so, which one: authentication, encryption, traceability, others):
Reliability requirements (e.g. fault tolerance):
Traceability to a specific clock, clock quality, path, time:
Holdover requirement:
Cost (consumer, enterprise, carrier):
Auto-configuration (plug and play):
Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?):
Scale and scalability:

3.5. ToD/ Internet

General time distribution over the Internet or IP networks, is often called Time of Day or Wall-clock. Most existing use cases are using NTP over the Internet with low precision requirements. However, new applications are arising that require higher precision rates than what is currently available.

Internet TOD is important to the maintenance of IT infrastructure in an organization. Generally the larger an organization becomes, the more important time synchronization is. Time synchronization is critical for the following: 1. Server and router log file entry time tags 2. "Date modified" attributes for files 3. Chron job scheduling 4. Security protocol with limited time windows for key exchange.

Server and Router log file time tag accuracy is essential to network diagnostic tools. Such tools are used to determine the root cause of a network failure or security breach. Often it is important to determine the order in which certain events occur amongst a number of network devices. The "Date modified" fields of files may also be part of this type of analysis.

Often Chron jobs perform operations on files depending on the times in the "Date modified" attributes files. These files might reside on more than one computer or server.

Many security protocols, such as Kerberos, depend on authentication "tickets" which expire after a short time. This means that an authenticating server gives a ticket to a client, which the client can send to another server for some service which requires authentication. The time limit is intended to reduce the threat of the "Man in the middle attack." To work the two servers need to have clocks synchronized to a time error which is smaller than the ticket time out period. To increase security there is a desire to reduce the ticket time interval. As the time interval becomes shorter the need for server clock agreement is increased. The trend over time is to reduce the ticket time out period.

3.5.1. ToD/Internet Requirements

The requirements for the ToD/Internet are summarized as follows:

Synchronization Type (e.g. time, frequency or phase): time
Frequency stability: no requirement
Frequency accuracy : no requirement
Uncalibrated time/time stability: no requirement
Uncalibrated time/time accuracy: 10 ms
Stabilization Time : 1 hour
Jitter on recovered timing signal: 100 ms
Wander on recovered timing signal: 10 ms
What expected network characteristics (WAN, LAN, MAN, private, public, etc)?: All network types
Does the application require security? (if so, which one: authentication, encryption, traceability, others): Authentication sometimes used
Reliability requirements (e.g. fault tolerance): high availability. Clients must see multiple servers
Traceability to a specific clock, clock quality, path, time: Not important
Holdover requirement: 1 hour to 1 year. Depends on server redundancy architecture
Cost (consumer, enterprise, carrier): 0 - $10,000 USD (1)
Auto-configuration (plug and play):
Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?): 30-90 minutes to configure a new server, 5 minutes to configure a new client. Almost no management after initial deployment.
Scale and scalability: system must cover entire IT infrastructure of organization. Any 1 server will cover 1 building or campus.

(1) The free option implies pointing all clients at ntp servers available on the public internet.

3.6. Networking

Editor's note: need more info on this application.

3.6.1. Networking SLA Requirements

The requirements for the Networking SLA are summarized as follows:

Synchronization Type (e.g. time, frequency or phase):
Frequency stability:
Frequency accuracy :
Uncalibrated time/time stability:
Uncalibrated time/time accuracy:
Stabilization Time :
Jitter on recovered timing signal:
Wander on recovered timing signal:
What expected network characteristics (WAN, LAN, MAN, private, public, etc)?:
Does the application require security? (if so, which one: authentication, encryption, traceability, others):
Reliability requirements (e.g. fault tolerance):
Traceability to a specific clock, clock quality, path, time:
Holdover requirement:
Cost (consumer, enterprise, carrier):
Auto-configuration (plug and play):
Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?):
Scale and scalability:

3.6.2. Networking CDR Requirements

The requirements for the Network CDR are summarized as follows:

Synchronization Type (e.g. time, frequency or phase):
Frequency stability:
Frequency accuracy :
Uncalibrated time/time stability:
Uncalibrated time/time accuracy:
Stabilization Time :
Jitter on recovered timing signal:
Wander on recovered timing signal:
What expected network characteristics (WAN, LAN, MAN, private, public, etc)?:
Does the application require security? (if so, which one: authentication, encryption, traceability, others):
Reliability requirements (e.g. fault tolerance):
Traceability to a specific clock, clock quality, path, time:
Holdover requirement:
Cost (consumer, enterprise, carrier):
Auto-configuration (plug and play):
Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?):
Scale and scalability:

3.7. Legal Uses of Time

With legal uses of time is meant the cases where high precision wall-clock is needed, just as in the ToD case, but with where the time source is traceable to UTC in a secure manner, i.e. through a certificate chain. It's also important for the legal-time case that the certificate chain is set-up so that it provides for an audit trail, where the ToD provided at any given moment can be traced to a known source or standard (i.e. a national timescale or time laboratory). One typical application that would benefit from high accuracy legal time is event correlation in computer systems logs, and similar applications.

3.7.1. Legal Uses of Time Requirements

There are timing applications for which accuracy and other characteristics are legally mandated, such as:

- Pay-by-time services (e.g., parking meters, taxicab meters, coin-operated laundries),

- First-arrival succeeds applications (e.g., races, stock-market exchanges),

- Devices that require accurate frequency for calibration (e.g., police radar, RF broadcast).

For such applications the legal requirements usually dictate both precision and accuracy, and frequently also traceability and security considerations. There also may be requirements for keeping of logs for some amount of time, for certification of correct operation by qualified personnel, and specification of the national timing standard to be used.

Due to the large number of disparate applications covered by legal uses of time, it is not useful to attempt to codify all the possible requirements in a table, however, the following is a typical subset of requirements.

Synchronization Type (e.g. time, frequency or phase): usually time, some frequency
Frequency stability:
Frequency accuracy : 100 ppm
Uncalibrated time/time stability:
Uncalibrated time/time accuracy: from 10s of ms to better than 1 second
Stabilization Time :
Jitter on recovered timing signal:
Wander on recovered timing signal:
What expected network characteristics (WAN, LAN, MAN, private, public, etc)?: private well-engineered IP networks, public Internet
Does the application require security? (if so, which one: definitely authentication, encryption, traceability, others):
Reliability requirements (e.g. fault tolerance): Yes, but not always specified
Traceability to a specific clock, clock quality, path, time: to national standard
Holdover requirement:
Cost (consumer, enterprise, carrier):
Auto-configuration (plug and play):
Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?):
Scale and scalability: usually not an issue

3.8. Metrology

Metrology for time and frequency is today mostly using tailored equipment and cabling for time/frequency transfer when doing laboratory work. However, in the future, using IP over existing networks in the laboratories would allow for greater flexibility and reuse of existing infrastructure rather than building out more special purpose infrastructure.

3.8.1. Metrology Requirements

We should distinguish between "primary" metrology of time and frequency performed in national metrology laboratories and some other timing centers (which deals with highly accurate and stable frequency standards, typically cesium standards and hydrogen masers and which ensures synchronization with a nanosecond accuracy) and applied metrology that mainly calibrates oscillators and clocks used as "secondary" standards in research organizations and industry. The use of time and frequency transfer in packet networks is limited in "primary" metrology, as it operates with frequency accuracy and stability in the order of 1e-14 and better and time accuracy in nanoseconds (1 ns represents 1 foot of the light path in vacuum). In turn, time and frequency transfer through packet networks is quite challenging for applied metrology - it can profit from any

improvement of transfer accuracy, therefore the values in table 8 should be considered as minimum target values. Whenever possible, accuracy of distributed time should be better then time accuracy provided by a GPS receiver. Short distance application (over LAN) usually require better accuracy than long distance application using WAN.

Note: some applications might belong into both metrology and measurement application groups.

The requirements for the Metrology are summarized as follows:

Synchronization Type (e.g. time, frequency or phase): Time and frequency
Frequency stability: 1 ppb, lowest possible phase noise
Frequency accuracy : 1 ppb
Uncalibrated time/time stability: Lowest possible phase noise
Uncalibrated time/time accuracy: 1 us
Stabilization Time : Not important, 1 hour is acceptable
Jitter on recovered timing signal: 1 us
Wander on recovered timing signal: 1 us
What expected network characteristics (WAN, LAN, MAN, private, public, etc)?: Any that can offer required parameters
Does the application require security? (if so, which one: authentication, encryption, traceability, others): Authentication is required when public networks are used
Reliability requirements (e.g. fault tolerance): Low fault tolerance, user should know whether expected parameters were assured or not
Traceability to a specific clock, clock quality, path, time: Very important
Holdover requirement:
Cost (consumer, enterprise, carrier): Cost should correspond with provided parameters
Auto-configuration (plug and play): Important at client side
Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?): Both provider and customer should accept network related configuration, long distribution path might require calibration
Scale and scalability:

3.9. Sensor Networks

More generally, there is growing interest in clock synchronization in massively parallel sensor networks. Advances in wireless communications have enabled the development of low power miniature

sensors that collect and disseminate data from their immediate environment. Although each sensor has limited processing power, through distributed processing the network becomes capable of performing various tasks of data fusion, but only assuming a common time base can be established.

3.9.1. Sensors Networks Requirements

The requirements for the Sensor are summarized as follows.

Synchronization Type (e.g. time, frequency or phase): time
Frequency stability: 1%
Frequency accuracy : 1000ppm
Uncalibrated time/time stability: 1 part per hundred
Uncalibrated time/time accuracy: 1 second
Stabilization Time : intermediate
Jitter on recovered timing signal: several seconds up to 1 milliHz
Wander on recovered timing signal: less than one second above 1 milliHz
What expected network characteristics (WAN, LAN, MAN, private, public, etc)?: distributed hop-to-hop network
Does the application require security? (if so, which one: authentication, encryption, traceability, others): authentication, encryption
Reliability requirements (e.g. fault tolerance): intermediate
Traceability to a specific clock, clock quality, path, time: one master
Holdover requirement: 1% wander per 10 days
Cost (consumer, enterprise, carrier): depends on application, from very low to intermediate
Auto-configuration (plug and play): strong requirement as there are many sensors
Manageability (how much effort the operator needs to put in to manage this application?) - In-band or out-of-band of protocol (MIBs?): in-band, self organizing, little to no storage on device
Scale and scalability: must scale to 1000s of intercommunicating sensors

4. Network Dependencies

When using packet networks to transfer timing, packet delay variation, propagation asymmetry, and maximum permissible packet rate all have a significant bearing on the accuracy with which the client is able to determine absolute time. Thus the network environment has a large bearing on the quality of time that can be delivered.

Timing distribution is highly sensitive to packet delay variation, and thus can deteriorate under congestion conditions. Furthermore the disciplining of the client's oscillator (the sole component of

frequency transfer, and a critical component of time transfer) is a function that should not be disrupted. When the service is disrupted the client needs to go into "holdover" mode, and its accuracy will consequently be degraded. Depending on the relative quality of the client's clock and the required quality after disciplining, a relatively high packet rate may be required.

Packet delay variation can to some extent be addressed by traffic engineering, thus time transfer within a constrained network environment might reasonably be expected to deliver a higher quality time service than can be achieved between two arbitrary hosts connected to the Internet. Greater gains can probably be obtained by deploying equipment that incorporates IEEE 1588 style on-the-fly packet timestamp correction (or any other form of on-path support), or follow-up message mechanisms that report the packet storage and forward delays to the client. However one can only be sure that such techniques are available along the entire path in a well-controlled environment. Therefore, time transfer protocols should not assume the availability of on path support, but utilizes it where available.

The packet rate between the time-server and its client also has a bearing on the quality of the time transfer, because at a higher rate the smart filter has a better chance of extracting the "good" packets. How the packet rate relates to the accuracy is dependent on the filter algorithm in use. In a controlled environment it is possible to ensure that there is adequate bandwidth, and that the server is not overloaded. In such an environment the onus moves from protecting the server from overload, to ensuring that the server can satisfy the needs of all of the clients.

Congested and overloaded paths might influence the quality of timing transfer. In a constrained network environment, it's assumed that a service provider will ensure that packet delivery is done in according to the timing transfer needs of the network operator.

5. Network Topology

Editor's note: This section needs to be discussed.

6. Security Considerations

Time and frequency services are a significant element of network infrastructure, and are critical for certain emerging applications. Hence time and frequency transfer services MUST be protected from being compromised, and for some of the applications described above such as legal time, the ability to provide and audit trail to the

timing source. One possible threat is a false time or frequency server being accepted instead of a true one, thus enabling an attacker to alter the time and frequency service provided. Other possible scenarios are to be able to distinguish between trusted clients and non-trusted clients when providing service.

Any protection mechanism must be designed in such a way that it does not degrade the quality of the time transfer. Such a mechanism SHOULD also be relatively lightweight, as client restrictions often dictate a low processing and memory footprint, and because the server may have extensive fan-out.

The following authentication mechanisms need to be considered:

1. of server by client (depending on the application)

2. of client by server (depending on the application)

3. transactions (depending on the application)

7. IANA Considerations

No IANA actions are required as a result of the publication of this document.

8. Acknowledgements

The authors wish to thank Stewart Bryant, Yaakov Stein, Karen O'Donoghue, Laurent Montini, Antti Pietilainen, Stefano Ruffini, Vladimir Smotlacha, Greg Dowd, Doug Arnold and the LXI Consortium Technical Committee for input on this draft.

9. Informative References

[1588] IEEE, "1588-2002 Standard for A Precision Clock Synchronization Protocol for Networked Measurement and Control Systems".

[G8261] ITU-T, "Recommendation G.8261/Y.1361 - Timing and synchronization aspects in packet networks", April 2008.

[G8262] ITU-T, "Recommendation G.8262/Y.1362 - Timing Characteristics of Synchronous Ethernet Equipment Slave Clock (EEC)", August 2007.

[G8264] ITU-T, "Recommendation G.8264/Y.1364 - Distribution of timing through packet networks", 2008.

[RFC1305] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation", RFC 1305, March 1992.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4197] Riegel, M., "Requirements for Edge-to-Edge Emulation of Time Division Multiplexed (TDM) Circuits over Packet Switching Networks", RFC 4197, October 2005.

[RFC4553] Vainshtein, A. and YJ. Stein, "Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)", RFC 4553, June 2006.

[RFC5086] Vainshtein, A., Sasson, I., Metz, E., Frost, T., and P. Pate, "Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)", RFC 5086, December 2007.

[RFC5087] Stein, Y(J)., Shashoua, R., Insler, R., and M. Anavi, "Time Division Multiplexing over IP (TDMoIP)", RFC 5087, December 2007.

Appendix A. Existing Time and Frequency Transfer Mechanisms

In this section we will discuss existing mechanisms for transfer of time information, frequency information, or both. It should be noted that a sufficiently accurate time transfer service may be used to derive an accurate frequency transfer. Indeed, this is exactly what happens in a GPS disciplined frequency standard. On the other hand, an accurate frequency transfer service, while itself unable to transfer absolute time, is usually used to support and improve the performance of the time transfer service. Indeed, implementations of NTP or IEEE 1588 clients can be considered to consist of two phases. First, a local oscillator is locked to the server's frequency using incoming information from the incoming packets, and then the local time set based on the server's time and the propagation latency. By maintaining a local frequency source, the client requires relatively infrequent updates, and can continue functioning during short periods of network outage. Moreover, it can be shown that this method results in significantly better time transfer accuracy than methods that do not discipline a local clock.

Time transfer mechanisms can be divided into three classes. The

first class consists of mechanisms that use radio frequency transport, while the second mechanism uses dedicated "wires" (which for our purposes include optical fibers). The third, which will be our main focus, exploits a Packet Switched Network for transfer of timing information. Advantages and disadvantages of these three methods are discussed in the following subsections.

A.1. Radio-based Timing Transfer Methods

The transfer of time by radio transmission is one of the oldest methods available, and is still the most accurate wide area method. In particular, there are two navigation systems in wide use that can be used for time transfer, The LOng RAnge Navigation (LORAN) terrestrial radio system, and the Global Navigation Satellite System (GNSS). In both cases the user needs to be able to receive the transmitted signal, requiring access to a suitable antenna. In certain situations, e.g. basement communications rooms and urban canyons, the required signal may not be receivable.

Radio systems have high accuracy, far better than what we will later see can be achieved by existing PSN technologies. However coverage is limited; eLORAN for example only covers North America, and GPS does not have good coverage near the poles.

Although civilian use is sanctioned, the GPS was developed and is operated by the U.S. Department of Defense as a military system. For this reason there are political concerns that rules out its use in certain countries. The European Union is working on an alternative system called Galileo, which will be run as a commercial enterprise. In addition, GPS has some well-documented multi-hour outages, and is considered vulnerable to jamming. One major PTT also reports that they see a 2% per year failure rate for the antenna/receiver/ clock-out chain.

While a radio-based timing service may be acceptable for some sites, it is frequently impractical to use on a per equipment basis. Hence, some form of local timing distribution is usually also required.

A.2. Dedicated Wire-based Timing Transfer Methods

The use of dedicated networks in the wide area does not scale well. Such services were available in the past, but for reasons of cost and accuracy have been superseded by GPS based solutions.

In the local area, one new technique is emerging as a mechanism for time transport, namely DOCSIS Timing Interface(DTI). DTI was designed by DOCSIS for the distribution of time in a cable head-end in support of media access control. Time transfer is packet-based

over a multi-stage hub and spoke dedicated network. It uses a single twisted-pair in half-duplex to eliminate inaccuracies due to the length differences between the pairs in a multi-pair cable.

The DTI approach is applicable for special applications, but the need for a dedicated network imposes significant drawbacks for the general time transfer case.

Synchronous Ethernet is a technique that has recently been approved by ITU-T, it provides frequency distribution over Ethernet links. Modern dedicated-media full-duplex Ethernet, in both copper and optical physical layer variants, transmits continuously. One can thus elect to derive the physical layer transmitter clock from a high quality frequency reference, instead of the conventional 100 ppm crystal-derived transmitter rate. The receiver at the other end of the link automatically locks onto the physical layer clock of the received signal, and thus itself gain access to a highly accurate and stable frequency reference. Then, in TDM fashion, this receiver could lock the transmission clock of its other ports to this frequency reference. Apart from some necessary higher layer packet based configuration and OAM operations to transport synchronization status messaging, the solution is entirely physical layer, and has no impact on higher layers.

At first sight it would seem that the only application of Synchronous Ethernet was in frequency transfer (it has no intrinsic time transfer mechanism). However, the quality of packet-based time transfer mechanism can be considerably enhanced if used in conjunction with Synchronous Ethernet as a frequency reference.

A.3. Transfer Using Packet Networks

When using a PSN to transfer timing, a server sends timing information in the form of packets to one or multiple clients. When there are multiple clients, the timing packets may be multicast. Software/hardware in the client recovers the frequency and/or time of the server based on the packet arrival time and the packet contents.

There are two well-known protocols capable of running over a general-purpose packet network, NTP [RFC1305], and IEEE 1588 [1588]. NTP is the product of the IETF, and is currently undergoing revision to version 4. PTP (a product of the IEEE Test and Measurement community) is specified in a limited first version (1588-2002), and the second version (1588-2008)was approved recently.

It is important that NTP, IEEE-1588 or any other future packet based time transfer mechanism do not break each other if they run in the same network.

A.3.1. NTP summary description

NTP is widely deployed, but existing implementations deliver accuracy on the order of 10 milliseconds. This accuracy is not adequate for the applications described above. Current NTP suffers from the fact that it was designed to operate over the Internet, and the routers and switches make no special concessions to NTP for preservation of time transfer accuracy. Furthermore, typical update rates are low and can not be significantly increased due to scalability issues in the server. In addition most NTP time servers and time receivers have a relatively unsophisticated implementation that further degrades the final time quality. However, proprietary NTP implementations that use other algorithms and update-rates have proved that NTP packet formats can be used for higher accuracy.

A.3.2. IEEE1588 summary description

The information exchange component of IEEE 1588 is a protocol known as Precision Time Protocol (PTP). PTP version 1 (1588-2002) was a time transfer protocol that exclusively used multicast technique and it was primarily developed for Industrial Automation and Test and Measurement applications. It is widely anticipated that wide scale deployment of PTP will be based on PTP version 2 (1588-2008).

IEEE Std 1588-2008 can be considered to consist of several components:

1. A configuration and control protocol

2. A time transfer protocol

3. A time correction protocol

4. Physical mapping

The configuration and control protocol is based on the multicast approach of IEEE Std 1588-2002 (multicast IP with recommended TTL=1, UDP, PTP payload with equipment identifier in the payload). The rationale for this approach was that the equipment needed to be "plug and play" (no configuration), was required to map to physical media other than Ethernet, and had to have a very low memory and processor footprint. IEEE Std 1588-2008 includes Unicast messages.

The time transfer protocol is a standard two-way time transfer approach used in other packet-based approaches. Like all such approaches it is subject to inaccuracies due to variable store and forward delays in the packet switches, and due to the assumption of symmetric propagation delays. For IEEE Std 1588-2008, the time

transfer packets (in both directions) may be operated in a multicast or unicast mode.

The time correction protocol is used to correct for propagation, store and forward delays in the packet switches. This again may be operated multicast or unicast. This mechanism requires some level of hop-by-hop hardware support. This mechanism may also be considered a concept in its own right and may be adapted to enhance other packet time transfer protocols such as NTP.

The IEEE Std 1588-2008 specification describes how the PTP operates over the Ethernet/IP/UDP protocol stack. It includes annexes that describe PTP operation over pure layer 2 Ethernet, and over a number of specialist media.

The mappings of interest for telecommunications are PTP over UDP/IP, PTP over MPLS , and perhaps PTP over Ethernet. They may operate in unicast or multicast. Issues of a suitable control management and OAM environment for these applications are largely in abeyance, as are considerations about the exact nature of the network environment.

It is also worth noting the existence of a second IEEE effort, IEEE 802.1AS. This group is specifying the protocol and procedures to ensure synchronization across Bridged and Virtual Bridged Local Area Networks for time sensitive applications such as audio and video. For these LAN media the transmission delays are assumed to be fixed and symmetrical. IEEE 802.1AS specifies the use of IEEE 1588 specifications where applicable in the context of IEEE Standards 802.1D and 802.1Q. Synchronization to an externally provided timing signal (e.g., a recognized timing standard such as UTC or TAI) is not part of this standard but is not precluded. IEEE 802.1AS will specify how stations attached to bridged LANs to meet the respective jitter, wander, and time synchronization requirements for time-sensitive applications.

Appendix B. Other Forums Working in this Problem Space

The NTP WG is the IETF group working on time distribution, but is presently only documenting NTPv4 and is not working on new algorithms or protocols. It is expected that many participants of the NTP WG will participate in the TICTOC effort.

The PWE3 WG has discussed frequency distribution for the TDM PW application, however it is not chartered to develop protocols for this purpose. It is expected that participants of the PWE3 WG who were active in the TDM PW discussions will participate in the TICTOC effort.

The IEEE approved the version 2 of the IEEE 1588 protocol (IEEE Std 1588- 2008) that will run over more types of PSNs. The protocol to be specified contains elements that will be of use in an IETF environment, but is unlikely to be regarded as being a complete, robust solution in such an environment. If the IEEE 1588 structure is deemed to be a suitable platform, then the IETF could contribute an Internet profile, including a complete distributed systems environment suitable for our purposes. Alternatively, the IETF could perhaps borrow some of the delay correction mechanisms and incorporate them into a development of a new version of NTP.

In addition, IEEE 802.1AS is working on Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks, basing itself on the IEEE 1588 standard.

ITU-T SG15 Question 13 has produced Recommendation G.8261 "Timing and synchronization aspects in packet networks" [G8261]. This Recommendation defines requirements for various scenarios, outlines the functionality of frequency distribution elements, and provides measurement guidelines. It does not specify algorithms to be used for attaining the performance needed. ITU-T has also consented G.8262 "Timing Characteristics of Synchronous Ethernet Equipment Slave Clock (EEC)" [G8262], and G.8264 "Distribution of timing through packet networks" [G8264]. G.8262 specifies the requirements for Synchronous Ethernet clocks and G.8264 defines the protocol for Synchronization Status Message (SSM) for Synchronous Ethernet. To date the ITU-T has focused on Ethernet infrastructure, but this is likely to extend to an MPLS environment. Two new work items, G.paclock.bis and G.pacmod.bis extend the work, and in particular, G.pacmod.bis intends to introduce time transfer. The scope for G.paclock.bis is to define the requirements for packet-based clocks. This is an area where the IETF, with its expertise in IP and MPLS networks, may co-operate with the ITU.

Authors' Addresses

Silvana Rodrigues
Zarlink Semiconductor
400 March Road
Ottawa K2K 3H4
Canada

Phone: +1 613 2707258
Email: silvana.rodrigues@zarlink.com

Kurti Erik Lindqvist
Netnod
Bellmansgatan 30
Stockholm S-118 47
Sweden

Phone: +46 708 30 60 01
Email: kurtis@kurtis.pp.se