17IND06 – FutureGrid II

Good practice guide on reviewing accuracy and security of digital time synchronization protocols

Grant Agreement number	**17IND06**
Project short name	**Future Grid II**
Project full title	**Metrology for the next-generation digital substatior instrumentation**
Version 1.0 drafted	August 12th 2021

Deliverable D5

Due date: May 2021	Actual submission date: November 2021

Partners Involved

PTB

Lead Partner

PTB

Authors:

Kristof Teichel

Contributors:

Kristof Teichel: kristof.teichel@ptb.de

This project has received funding from the EMPIR programme co-financed by the Participating States and from the European Union's Horizon 2020 research and innovation programme.

The EMPIR initiative is co-funded by the European Union's Horizon 2020 research and innovation programme and the EMPIR Participating States

Report Status: CO Confidential, only for members of the consortium (including EURAMET and the European Commission Services)

Table of Contents

1 PURPOSE AND SCOPE

This guide presents good practice concerning the assessment and review of time synchronization setups, for their properties regarding

1. Accuracy,
2. Security and reliability, and
3. Effort and (in)convenience.

In particular, this guide is supposed to assist users who are in the process of evaluating which technologies to use for their synchronization needs, and/or what results to expect from a given technology.

We differentiate between such assessment or review that is done pre-deployment of a synchronization setup (by analysis and prediction) versus such assessment or review that is done post-deployment (by experiment and measurement).

This guide focuses on pre-deployment techniques, for the following reasons. Firstly, post-deployment assessment of accuracy is relatively well-established, and treated in other documents such as Deliverable D4 of this same project, 17IND06 *FutureGrid II*. Secondly, post-deployment, experimental assessment of security can be impractical, because the desirable case of not finding vulnerabilities is not conclusive. Thirdly and most importantly, post-deployment assessment is often not helpful given our scope and target group of users who are still in the process of selecting technologies.

2 RESPONSIBILITIES

Responsibility on correctness and documentation of measurement and assessment results is on the person performing the review. Responsibility on the correctness of this document is on the author.

3 EXISTING TECHNOLOGIES

In this section, we list synchronization methods and technologies that are already available to users and provide security, as a starting point for evaluation. This list is not exhaustive, and a technology being listed does not necessarily constitute a recommendation from the author to use that technology, neither in a specific use case nor in general. It does, however, constitute a recommendation to be aware of the technology's existence, and to keep our comments on it in mind in a potential evaluation of it.

3.1 Internet-based technologies

We first list technologies designed for use via the internet. Thus, even those recommended for security properties should only be used if internet-level accuracy (low-millisecond to mid-microsecond level) suffices.

Table 3.1: Secured internet-based synchronization designs.

Name	Year	Type	Implementation	Ref.	Author Comments
NTP-MD5	1992	IETF Standard	Public Source	[1]	Outdated, use at your own peril
NTP-Autokey	2010	IETF Experimental	Public Source	[2, 3]	Deprecated, do NOT use!
ANTP	2016	Paper	None Available	[4]	Cannot comment
STS	2018	Paper	Proprietary	[5]	Looks promising
NTS	2020	IETF Standard	Multiple (Public)	[6]	Recommended (Careful: microsecond level at best!)

3.2 Radio-based and LAN-based technologies

We now list technologies not intended for internet use. They may come with their own limitations

Table 3.2: Secured internet-independent synchronization.

Name	Year	Type	Implementation	Ref.	Author Comments
TinySeRSync	2006	Standard	Public Source	[7]	For wireless sensors, locally
ASTS	2007	Standard	None Available	[8]	Seems unsecure, best NOT use!
Galileo OS-NMA	2018	ESA Standard	Unfamiliar	[9]	Looks promising (But: GNSS reception required)
Galileo PRS	2018	ESA Standard	Unfamiliar	[10]	Access restriction level for critical infrastructure users unclear
GPS Chimera	2017	Standard	Unfamiliar	[11]	Looks promising (But: GNSS reception required)
PTP (v2.1)	2020	IEEE Standard	In development	[12]	Recommended (for local use, long-range is effortful)

4 PRE-DEPLOYMENT ASSESSMENT

We present a collected overview on how to assess both the accuracy and reliability levels and relate them to the required effort, for different digital methods of synchronizing clocks. The presented process is intended for end users who require time synchronization but are not certain about how to judge at least one of the aspects. It can not only be used on existing technologies but should also be transferable to many future approaches. We further relate this approach to several examples. We discuss approaches such as medium-range White Rabbit connections over dedicated fibres, a method that occupies an extreme corner in the evaluation, where the effort is exceedingly high, but also yields excellent accuracy and significant reliability.

The presented was directly motivated by this EMPIR project 17IND06 regarding the change from analog to digital instrumentation in the European power grid. Improving and cataloguing the availability of security, accuracy, and convenience of time transfer techniques is a stated goal in this. The topic is, however, not just applicable in energy grid contexts, but has come up in other areas as well. Specifically, these areas have been the financial market, in particular the EU guideline MiFID II [3], as well as the telecommunication area and data centre applications. There have also been recent efforts to classify, assess and improve different synchronization techniques, especially those that require satellite support.

One result we present is an informed but simplified procedure for numerical score-based assessment of time transfer technologies for the three categories accuracy, reliability, and effort. Another second result we present concerns example evaluations of setups such as medium-range and long-range White Rabbit links, which are very accurate (so much so that they are visibly sensitive to changes in outdoor temperature in proximity to the fibre connection) and quite reliable but require large effort, and NTP connections secured with NTS, which are low effort and very reliable, but relatively inaccurate.

An adapted version of this section has been published at EFTF 2021 [13].

4.1 The Time Transfer Trade-off Triangle

In our experience with the time transfer technologies existing to date, there is no single technology that will achieve even near the best scores in all three categories at once. In other words, the search for an “optimal” technology always involves a trade-off. A given technology (of the ones available, this does not seem to be an inherent immutable problem) reliably is either not very accurate, or not very reliable, or not very convenient in the sense that it requires high effort. This can be expressed in a triangular graph, which is why we dub this phenomenon the Time Transfer Trade-off Triangle. A sketch of this can be seen in Fig. 4.1, which also outlines the corner cases and gives examples for each of them:

Fig. 4.1: The Time Transfer Trade-off Triangle.
A = accurate, C = convenient, R = Reliable

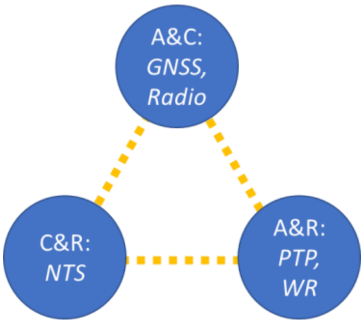

- Both GNSS and radio methods offer good accuracy and convenience, but they are inherently not very reliable.
- An NTS-Secured NTP connection is both very convenient and very reliable, but not very accurate.
- Protocols such as PTP and White Rabbit give very accurate time transfer with at least decent reliability (which can be improved with cryptography), but what they do not offer is good convenience.

4.2 Questionnaire

With the Time Transfer Trade-off Triangle, we have a problem statement of sorts. Now what is needed is an approach for how to evaluate a set of given technologies regarding their position in that triangle. For this, we present a questionnaire that enables users to perform an informed but simplified procedure for numerical score-based assessment of time transfer technologies for the three categories accuracy, reliability, and effort.

Our informed but simplified procedure for numerical score-based assessment of technologies consists of a tabular questionnaire and scoring system. The questionnaire is divided into three parts. The first part concerns the transportation method used in the time transfer method. This part deals with both the question of the transportation medium and the message flow model. The second part concerns cryptographic methods used to protect the time transfer data, treating both the question if any cryptography is used and, as importantly, how it is transported. The third part concerns the question of dedicated hardware required and used by the time transfer method.

In this questionnaire, accuracy is scored by our estimation of the attainable offset level, and reliability and effort are scored via arbitrarily chosen additive scores.

For accuracy rating, the score is calculated by determining the minimum value for both worst case (left value) and best case (right value). The two values then represent the worst and best case of the overall approach, respectively. The scores represent (very roughly) the order of magnitude of the resulting accuracy when used as a negative exponent to the power of 10 (e.g., a final score of 3-6 means: worst case 10^-3 seconds, so millisecond level, best case 10^-6 seconds, so microsecond level); thus, higher scores are better.

For reliability rating, the final score is calculated simply by addition of all relevant values in the questionnaire. Higher scores are better.

For the (in)convenience rating, the final score is also calculated by addition of all relevant values in the questionnaire. Higher scores mean more effort, so lower scores are more desirable.

Table 4.1: Questionnaire for numerical assessment.
AR = Accuracy Rating
RR = Reliability Rating
ER = Effort Rating

Question	**Option**	**AR**	**RR**	**ER**	**Example Technology**
Which transportation method is used?	Public Internet	1-4	1	0	NTP
	Closed Network	3-7	2	1	PTP
	Wireless Radio	4-8	3	0	GPS
	Dedicated Fibre (local)	8-12	10	3	WR
	Dedicated fibre (long-range)	7-11	5	10	WR, PTP
… and what is the message flow model?	One-way	-	-4	0	GPS
	Two-way	-	0	0	NTP
Is cryptographic protection involved?	No	-	0	0	NTP, PTP
	Yes, weak source authentication	-	4	0	PTP with group key
	Yes, strong source authentication	-	10	0	NTS, Galileo OSNMA
If YES, how is it communicated?	Same message as time data, without extra design effort	1-3	0	0	Roughtime
	Same message as time data, with deliberate design	2-6	0	0	NTS
	Separate message from the time data	-	-1	1	Secure PTP
To what extent is dedicated hardware used?	None, other than very common multi-purpose hardware (PCs)	1-6	0	0	NTP
	Dedicated hardware for all end devices	3-8	2	1	GPS
	Dedicated hardware as both end devices and middleware	-	4	3	PTP

4.3 Exemplaric Evaluations

For the sake of understandability, we give four examples of technologies and how their assessment turns out, with our evaluation method as presented above.

Long-range White Rabbit

One interesting corner case that presents itself is that of long-range White Rabbit connections. These offer great accuracy (often in the single nanosecond range), and very solid reliability, since the whole transportation network (both active electrical devices and fibre) are necessarily closely controlled. The effort of organizing a dedicated fibre connection plus the necessary White Rabbit hardware, however, is enormous. The evaluation of this corner case according to our assessment method as presented above can be seen in Tab. 4.2, with final scores of 7-11 for accuracy, 9 for reliability, but 13 for effort.

Table 4.2: Evaluation results for long-range White Rabbit

	Accuracy	**Reliability**	**Effort**
Dedicated fibre (long-range)	7-11	5	10
Two-way	-	0	0
No cryptography	-	0	0
Dedicated hardware both end and middle	-	4	3
Final Score	7-11	9	13

NTS-Secured NTP

Another potentially interesting corner case is that of a simple NTP connection [2] secured with measures according to the relatively new Network Time Security specification [5]. This offers about the highest reliability we could currently envision, and the effort is no more than having some kind of computer with an internet connection. The offered guaranteed accuracy is only in the millisecond range, however. This is visible in Tab. 4.3, where this corner case is evaluated according to our approach.

Table 4.3: Evaluation results for NTS-secured NTP

	Accuracy	**Reliability**	**Effort**
Public internet	1-4	1	0
Two-way	-	0	0
Strong source authentication	-	10	0
Same message as time data, with deliberate design	2-6	0	0
No dedicated hardware other than multi-purpose (PCs)	1-6	0	0
Final Score	1-4	11	0

Unsecured GNSS Synchronization

It is also interesting to look into GNSS synchronization from the angle of our assessment method. A typical GNSS link without any extra security (such as given by a standard GPS or Galileo receiver) is a method that is often used when users need a higher accuracy than NTP can offer, but shy away from the cost of dedicated fibres, opting instead for relatively cheaper GNSS hardware.

Table 4.4: Evaluation results for unsecured GNSS synchronization

	Accuracy	**Reliability**	**Effort**
Wireless radio	4-8	3	0
One-way	-	-4	0
No cryptography	-	0	0
Dedicated hardware for end devices	3-8	2	1
Final Score	3-8	1	1

Cryptographically secured GNSS Synchronization

It is also interesting to look into GNSS synchronization from the angle of our assessment method. A typical GNSS link without any extra security (such as given by a standard GPS or Galileo receiver) is a method that is often used when users need a higher accuracy than NTP can offer, but shy away from the cost of dedicated fibres, opting instead for relatively cheaper GNSS hardware.

Table 4.5: Evaluation results for unsecured GNSS synchronization

	Accuracy	**Reliability**	**Effort**
Wireless radio	4-8	3	0
One-way	-	-4	0
Cryptography with strong source authentication	-	10	0
Same message as time data, with deliberate design	2-6	0	0
Dedicated hardware for end devices	3-8	2	1
Final Score	2-6	11	1

5 HINTS ON POST-DEPLOYMENT ASSESSMENT

In this section, we present what advice we have to offer on post-deployment assessment techniques, for accuracy, security, and effort.

5.1 Accuracy

The post-deployment assessment of accuracy is one topic that is relatively well covered by the scientific and metrological community already. In essence, the most important techniques are the following:

- Logging: Especially in digital systems performing any kind of clock synchronization, keeping records of what synchronization operations are performed is immensely helpful in both proving that all possible effort was taken (in case of a failure or problem of some sort) and understanding the extent to which synchronization might have gone wrong.
- Back-monitoring: In the case of two-way synchronization (which is to say, with message flow going in both directions between time source and time receiver), it is recommended to pursue any options that include the time source also observing the time receiver. For example, if a system is getting (rough) time from a national metrology institute (NMI) via NTP, that NMI's NTP infrastructure can monitor that system's NTP device as if it were a server itself. That way, the NMI can later give informed statements about the synchronization status of that system.
- External comparison: The most important technique for post-deployment accuracy assessment remains the simple comparison of time receiver and time source via channels external to the primary synchronization channel. This can take a number of forms. For example, a system that gets its time via GNSS reception could additionally be monitored via NTP, to ensure that at least its synchronization status is not off by NTP-level accuracy. Or a mobile atomic clock, calibrated at an NMI, can be used as an additional channel for time and frequency offset measurements (see also Deliverable D4 of this same EMPIR project 17IND06).

5.2 Security and Reliability

Assessing some aspects of reliability post-deployment is easy. For example, this is the case for availability of service: one simply observes the system, and monitors when a service (such as GNSS or radio reception) is effectively available in practice.

However, all aspects regarding security (i.e., reliability and robustness in the face of deliberate attacks) are much more difficult to evaluate in practice. Granted, it is possible to perform penetration tests or hacking sessions, and to monitor systems for security breaches. And we recommend taking all those measures as far as possible. However, the fact remains that absence of evidence (of possible attacks) is not evidence of their absence. The most harmful attacks have exactly the property that the attacked user or system is not aware of an attack happening.

The best recommendation we can give is to make an effort to stay up-to-date on the security community's results regarding any technology that one uses. For example, if a user decides to employ Galileo's relatively new OS-NMA service, they should keep an eye open for potential research regarding potential attacks on that service. Such research is going to be actively conducted especially in the early years after a service goes live.

5.3 Effort and (In)Convenience

In a sense, assessing effort and (in)convenience of a technology works especially well post-deployment. After all, once the system is in place and running, one can see exactly the resources spent and the efforts made. However, this is also when this knowledge is perhaps the least useful.

6 CONCLUSIONS

We have presented the problem of the Time Transfer Trade-off Triangle, which tells us that a user for the most part has to pick two out of the three desirable properties of good accuracy, good reliability, and good convenience. We have also given a reasonably concise way to navigate the Triangle before the deployment of any technology. We have shown four corner-cases in it that future users might find useful.

The next steps in refining our approach could consist of documenting a common metrological consensus on the accuracy score (which currently represents our own prognosis) and researching a more concrete quantifying approach for reliability and effort scores. It is hard to quantify reliability, though, and even though it might be tempting to measure effort in monetary values or units such as person months, we feel that prognoses in this area carry a greater inherent risk of error, and thereby of misleading users into decisions that later may turn out to be wrong.

We believe that the coming years will bring increased need for users from all kinds of fields to select dedicated time transfer technologies for their applications, and that there will be no one-size-fits-all solution that is in some way ideal for everyone. We hope that our work can help clarify for individual entities how they should approach the search for their own personal best solution and navigate the Time Transfer Trade-off Triangle without fear of getting lost.

In case of any questions or suggestions, please do not hesitate to contact the author (see title page for contact).

REFERENCES

[1] David L. Mills. Network Time Protocol (Version 3): Specification, Implementation and Analysis. RFC 1305, IETF Secretariat, 03 1992. https://tools.ietf.org/html/rfc1305.

[2] David L. Mills, J. Martin, J. Burbank, and W. Kasch. Network Time Protocol Version 4: Protocol and Algorithms Specification. RFC 5905. https://tools.ietf.org/html/rfc5905. IETF Secretariat, June 2010.

[3] David L. Mills and Brian Haberman. Network time protocol version 4: Autokey specification. RFC 5906. https://tools.ietf.org/html/rfc5906. IETF Secretariat, 2010.

[4] Benjamin Dowling, Douglas Stebila, and Greg Zaverucha. Authenticated Network Time Synchronization. Cryptology ePrint Archive, Report 2015/171. http://eprint.iacr.org/2015/171. 2015.

[5] Gorgy Timing – The Time/Frequency Expert. 2019. url: http://www.gorgy-timing.fr/FTP/COM/Company-Profile-Gorgy-Timing.pdf

[6] Daniel Fox Franke, Dieter Sibold, Kristof Teichel, Marcus Dansarie, and Ragnar Sundblad. Network Time Security specification for the Network Time Protocol. Internet-Draft draft-ietf-ntp-using-nts-for-ntp. Internet Engineering Task Force (IETF), Oct. 2020. url: https://datatracker.ietf.org/doc/draft-ietf-ntp-using-nts-for-ntp/

[7] Kun Sun, Peng Ning, and Cliff Wang. Tinysersync: secure and resilient time synchronization in wireless sensor networks. In Proceedings of the 13th ACM conference on Computer and communications security, pages 264–277. ACM, 2006.

[8] X. Yin, W. Qi, and F. Fu. “ASTS: An Agile Secure Time Synchronization Protocol for Wireless Sensor Networks”. In: Wireless Communications, Networking and Mobile Computing, 2007. WiCom 2007. International Conference on. IEEE, 2007, pp. 2808–2811. doi: 10.1109/WICOM.2007.697.

[9] I. Fernandez-Hernandez, V. Rijmen, G. Seco-Granados, J. Simón, I. Rodríguez, and J. David Calle. “A Navigation Message Authentication Proposal for the Galileo Open Service”. In: Navigation - Journal of The Institute of Navigation 63 (Mar. 2016), pp. 85–102. doi: 10.1002/navi.125

[10] PRS Homepage. url: https://www.gsa.europa.eu/security/prs

[11] J. Anderson, K. Carroll, N. DeVilbiss, J. Gillis, J. Hinks, B. O’Hanlon, J. Rushanan, and L. Scott. “Chips-Message Robust Authentication (Chimera) for GPS Civilian Signals”. In: Navigation - Journal of The Institute of Navigation (Sept. 2017), pp. 2388–2416. doi: 10.33012/2017.15206

[12] Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems. Standard. url: https://standards.ieee.org/develop/project/1588.html

[13] Kristof Teichel, Tapio Lehtonen and Anders Wallin. “Assessing Time Transfer Methods for Accuracy and Reliability”. 2021 Joint Conference of the IEEE International Frequency Control Symposium and European Frequency and Time Forum (IFCS-EFTF 2021).

APPENDIX

1 TIME SYNCHRONIZATION SECURITY: WHAT TO LOOK FOR IN DETAIL?

This section is intended to relay some of experience gathered in roughly seven years of designing and standardizing security measures for time synchronization protocols. It attempts to give an overview of the most important goals (both security-related and operational) that security designs must follow.

The intended use is by users who feel that the reliability and specifically security indication from the questionnaire is not sufficient for their understanding, or who simply seek to understand what it is that is important for security properties of time transfer technologies in general. This might be helpful for executives looking to improve their understanding of the problem, but even more so for engineers and operators who are tasked with designing, setting-up, or operating a system that requires reliable time synchronization.

First off, we establish a few very general and obvious statements about the overall interests of the participants of (network) time synchronization protocols, first looking at time clients (that is to say, time receivers such as substation timing systems), then at time servers (time sources, such as NMIs). After this, we present a (non-exhaustive) list of operational security goals for digital time synchronization.

1.1 Interests of Time Clients

Clients have the inherent interest of having a service available that allows them to synchronize their own clock to the clocks of others. First and foremost, there is a self-evident internal interest for this, since it is simply beneficial to agree with other network participants about something as fundamental as time and date. As a largely new development, there may additionally be requirements on participants' clocks (and therefore on their synchronization procedures), posed by other parties. The details of such (internal or external) requirements vary from use-case to use-case. Generally speaking, higher quality of the synchronization (i.e., an expected increase in the resulting accuracy and/or stability of the client's clock) is preferable. Also generally speaking, higher degrees of traceability are preferable as well. In any given scenario, there may well be concrete requirements posed on the quantifiable properties of participants' clocks and/or their synchronization procedures. Specifically, this often includes one or more of the following possibilities:

- Requirements on the resolution and/or the precision of a clock. This is (mostly) an aside in our context, as synchronization does not usually influence it.
- Requirements on the accuracy and/or stability of a given clock. Both quality and performance of the synchronization procedures can influence this.
- Requirements on the traceability of measurements made with a given clock. Both genuineness and performance of the time synchronization procedure influence this.

Remark: Traceability requirements are especially tricky, particularly if there is a meta requirement to have provable traceability. In time synchronization contexts, it is relatively easy to prove one's ability to have a clock of certain accuracy or stability, or to make traceable measurements. What is harder, however, is to prove that one actually had that level of accuracy or traceability at a given point in the past.

1.2 Interests of Time Servers

The interests of time servers are, generally speaking, much less potentially complex than those of time clients. Servers want to provide the services that clients wish for. The only other typical meta goal is that providing these services is not too taxing on the servers themselves, or on other, third-party infrastructure.

1.3 List of Security and Operational Goals

Here, we describe objectives that apply to actual time synchronization traffic, as opposed to the overall resulting synchronization process. These result directly from the participants' interests as listed above.

1.3.1 Genuineness of Synchronization Traffic

The first (and in our context most obvious) overall objective is that the time synchronization traffic that a given participant witnesses is genuine. Not only is this critical for traceability requirements, it also essential for guarantees on accuracy and stability of a local clock that is being adjusted according to that witnessed traffic.

Prevention of Wrongful Message Injection

Note that wrongful message injection includes the operation of altering existing messages, since the altered message is technically a new message. (The special case where one existing message *M1* is altered to be equal to another already existing message *M2* can be disregarded, since it is equivalent to replay of *M2*). We first introduce the concepts of confirming a participant's identity or a participant' authorization. These are the first goals that need to be established in order to enable meaningful reasoning about genuineness of traffic, and in particular to prevent wrongful message injection.

Participant Identity

This goal states that participants must be enabled to confirm who it is that they are communicating with. Statements about participant identity often have a semantic format such as *the owner of key K is entity X*.

Participant Authorization

This goal states that participants must be able to confirm that the entity they are communicating with has the right to assume the role that they are assuming. In our context, such a role might be that of a time server. Statements about participant authorization usually have a semantic format close to Entity *X* is authorized to assume role *R*.

Remark: Participant identity and authorization are obviously closely linked, since it is often an identity that is connotated with an authorization. However, it can be perfectly legitimate to skip the middleman and make a statement The owner of key *K* is authorized to be a time server, without mentioning that owner's identity at all.

Message Integrity / Authenticity

The most essential sub-objective to establish that a time synchronization message is genuine is message integrity, i.e., that the message was transmitted and received as it was sent. Any form of integrity confirmation is typically a statement of the format The owner of key K confirms message M.

However, note that such a statement carries little semantic relevance if nothing is known about the owner of key K in the first place. Integrity can therefore, generally speaking, not be meaningfully confirmed without some form of identity and / or authorization confirmation.

Confirmation of Correct Context

In addition to preventing the wrongful injection of new or altered messages, there is another whole aspect that needs confirmation: the context of a message. This is an issue that, while not unique, is particularly emphasized for messages that are part of time synchronization procedures.

Sender-Recipient Context

The first and most obvious part of context is: which participant intended the message to be transmitted to which other participant? It might happen that a message M is delivered to A, but was intended by B to be delivered from B to C. In this case, checks that A runs regarding integrity, authenticity and so on might all be successful (the message is real, after all), but A needs to be aware that Ms contents are probably not meant for him to use as part of the protocol.

Communication Instance Context (Replay Attacks)

Another relevant piece of context, which is a good bit more involved but critical in time synchronization, is this: has a given received message already been received before, and does that change its value? A critical attack vector with regards to this piece of context is given by replay attacks. A message M, intended by B to be delivered to A is delivered once, and then wrongfully delivered again. The message should pass most checks regarding genuineness, but it should still be flagged as problematic, because the time at which it is received is suboptimal (in general, time synchronization messages should always be delivered as fast as possible). A message is called fresh if and only if it is delivered for the first time.

Timing Context (Delay Attacks)

The last piece of context that we go into is the one most specific to time synchronization: has a received message been delivered in a timely fashion, or has it been wrongfully delayed? If an attacker has simply delayed delivery of a message, it will still pass all cryptographic checks, as well as freshness checks. But this operation can still have a negative influence on the accuracy of the clock whose synchronization the message is used for, due to the distortion of the timestamp that is logged for the message's reception.

1.3.2 *Privacy Concerns*

At the time of this writing, it seems like, there is the consensus that privacy goals do not have to be pursued at all. Nevertheless, debates over privacy concerns in recent years (mostly specific to NTP) have seen these goals raised, dropped and picked back up multiple times. We therefore wish to mention them and their existence in discussion, if only for the sake of completeness.

1.3.3 *Categorical Imperative*

There is another set of goals that can be seen as more abstract than the previous ones. The previous goals consider a given protocol that two or more given participants are actively involved in and then limit its scope to consequences of that specific protocol to those specific participants. However, the following facts are also worth recognizing:

- Any single protocol that a participant engages in is likely just one of many activities that said participants pursues. That participant may well have a global set of goals – meaning some of them might go beyond the scope of that specific protocol. As far as those global goals can be known or guessed, no single protocol should (unnecessarily) compromise them.
- Any single participant of network-based time synchronization is likely also a part of a larger distributed system. As such, it has additional responsibilities to keep said system functional and its other participants unharmed and able to pursue their desired activities.
- These facts can and should be accommodated for, by adhering to the goals below. This set of goals can be seen as additional, in the sense that these goals do not stand on their own but are requirements as to the way in which the security goals above are handled.

Avoidance of Excess Overhead

First, we look at goals concerning the avoidance of several kinds of overhead. They boil down to requiring that the methods used must not be overly taxing for the participants and their environment.

Memory Overhead

In our context, the desire to not use too much memory for the participation in a protocol mostly relates to how the protocol scales. That is to say that participants want to avoid significant memory requirements per association, especially if they predict that they will have great numbers of such associations – such as a time server will have if it is frequented by millions of clients.

Computational Overhead

Limiting computational costs for any protocol is always a good idea. In our context, there is also the additional challenge of having to accommodate for devices such as routers or embedded devices, which simply do not bring the computational power that might otherwise just be expected of a modern device.

Other Quality Factors of Time Synchronization

In the context of time synchronization, there are a few additional factors that should be considered regarding overhead. The most important one is that timestamps should be taken as close to their related events (typically reception or sending of a message) as possible. Therefore, it is important to avoid computationally expensive tasks (such as cryptographic operations) between the taking of a timestamp and its associated event as far as possible.

Traffic Overhead

Generally speaking, it is always desirable to send fewer and smaller messages if one can achieve the same. Avoiding excess traffic overhead relates both to Item 1 (all participants should be interested in using as little bandwidth as possible) and Item 2 (usage of bandwidth concerns entities other than just the protocol participants).