Institut für Automation
Abt. für Automatisierungssysteme

Technische Universität Wien

Projektbericht Nr. 183/1-77
March 1997

Orthogonal Accuracy Clock Synchronization

Ulrich Schmid

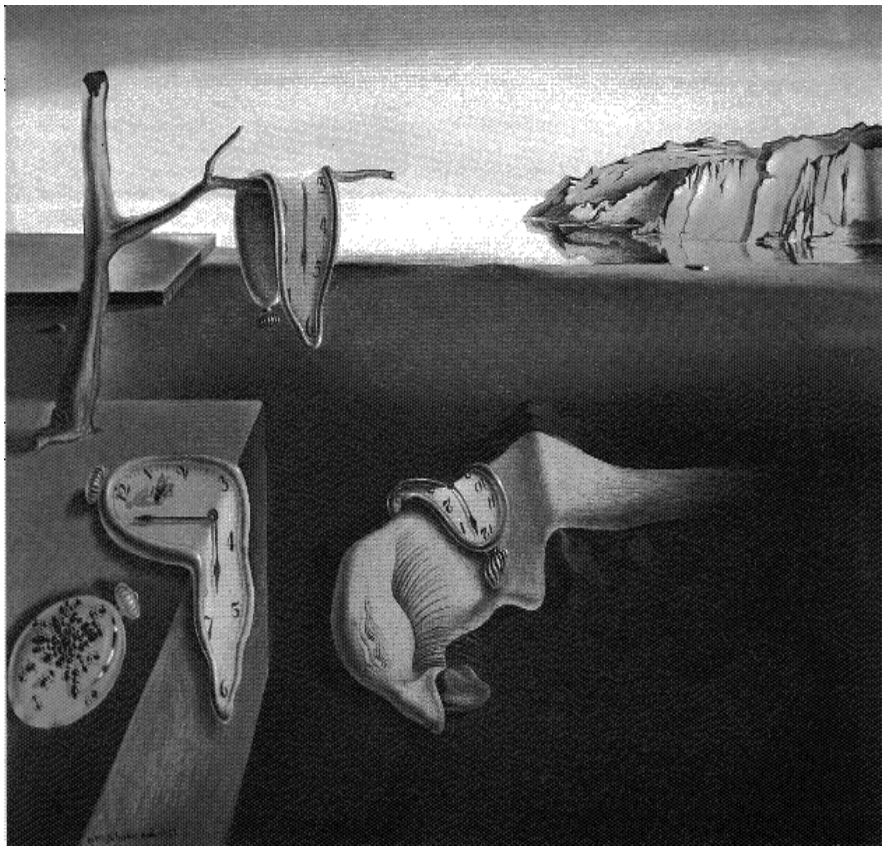

Salvador Dali, "Die Beständigkeit der Erinnerung"

Orthogonal Accuracy Clock Synchronization

Ulrich Schmid
Technische Universität Wien
Department of Automation
Treitlstraße 1, A-1040 Vienna
Email: s@auto.tuwien.ac.at

August 30, 1999

Abstract

We present description and analysis[1] of a novel *orthogonal accuracy* clock synchronization algorithm (OA), which takes care of both precision and accuracy with respect to external time. It is based upon the generic algorithm introduced in [SS97b] and utilizes a convergence function based on Marzullo's fault-tolerant intersection function. As far as precision is concerned, we show that OA has the same worst case performance as the well-known fault-tolerant midpoint algorithm of [LWL88]. However, relying upon a fairly realistic system model, our results are meaningful for very high precision as well: Impairments due to clock granularity and discrete rate adjustment cannot be ignored here anymore. Our accuracy analysis focusses on the nodes' local accuracy interval, which provides the application with an on-line bound on the current deviation from external time. We show that this bound could get larger than twice the necessary lower bound ("traditional accuracy"), hence OA is definitely suboptimal in this respect.

Keywords: interval-based clock synchronization, fault-tolerant distributed real-time systems, convergence functions, precision analysis, accuracy intervals, clock granularity, hybrid fault models, Marzullo function.

1 Introduction

Modern distributed systems usually run applications that rely on a global notion of time. In fact, algorithms for (fault-tolerant) distributed systems are considerably simplified and improved with respect to performance when mutually synchronized clocks are available, see [Lis93]. In addition, time rules daily life and hence most commercial computer applications, so system time should also have a well-defined relation to some external time like *Universal Time Coordinated* (UTC).

Excluding non-fault-tolerant solutions based on centralized clocks, we thus assume that each node p of a distributed system is equipped with a local clock $C_p(t)$ that continuously displays p's view of system time. Providing mutually synchronized ("precise")

[1]This research is part of our project SynUTC (http://www.auto.tuwien.ac.at/Projects/SynUTC/), which has been supported by the Austrian Science Foundation (FWF) under grant no. P10244-ÖMA and the Austrian START programme Y41-MAT.

local clocks is known as the internal synchronization problem, and numerous solutions have been worked out under the term fault-tolerant clock synchronization, see [YM93] for a bibliography. Providing mutually synchronized clocks that also relate to UTC (being "accurate" as well) is usually termed the external synchronization problem, due to the fact that UTC is provided externally to the system. A comprehensive collection of papers describing recent efforts in this area may be found in [Sch97b].

Among those is our research on *clock validation* introduced in [Sch95], which approaches the external synchronization problem by verifying whether the usually highly accurate (but possibly faulty) "authoritative time" provided by UTC time sources is consistent with the less accurate (but reliable) "validation time" formed by exchanging the information of all the local clocks in the system. If so, the authoritative time is accepted — if not, it is discarded and the nodes rely upon the validation time instead. The latter situation is encountered in case of failures[2] or unavailability of time information from UTC time sources, where clock validation obviously "degenerates" to pure internal synchronization. Therefore, the clock synchronization algorithm employed for computing the validation time must not only ensure precision but has to maintain high accuracy as well.

Traditional internal synchronization algorithms are ill-suited for coping with this requirement. In fact, although worst case accuracy bounds have been provided for most existing algorithms, it is nevertheless true that a static worst case bound is not representative for the "average" execution. A promising alternative are interval-based algorithms introduced in Marzullo's thesis [Mar84] and further exploited in [Lam87], [Mar90], [OSF92], [Sch95], [Mil95], [BI96], [SS97b], [Sch97c], where local time at external time t is expressed as an interval that contains t. Given a set of such intervals from different nodes, a usually smaller interval that contains t can be determined by means of a suitable interval-valued convergence function. Since accuracy bounds are hence maintained dynamically ("on-line") here, they are obviously representative for the "average" execution.

In [SS97b], we presented and analyzed a generic interval-based clock synchronization algorithm suitable for computing the validation time in our clock validation framework. According to the exposition above, it maintains bounded precision when started from an initially synchronized state and takes care of accuracy intervals as well. However, the algorithm is generic in the sense that the convergence function employed for computing the clock adjustments is left unspecified. As in [Sch86], all results (worst-case precision, accuracy, maximum adjustment, etc.) are hence expressed in terms of a few characteristic parameters of the convergence function. In order to determine the performance of a particular instance of the algorithm, all that needs to be done is to evaluate the characteristic parameters of the convergence function and to plug it into the generic results.

This paper provides description and detailed analysis of the *orthogonal accuracy algorithm* OA obtained by employing the *orthogonal accuracy convergence function* $\boldsymbol{\mathcal{OA}}$ in the generic algorithm of [SS97b]. It is organized as follows: Section 2 introduces our interval-based clock synchronization framework, including the generic algorithm (Subsection 2.1 and Appendix B), the generic precision analysis (Subsection 2.2 and Appendix C), and the particular fault model (Subsection 2.3). Section 3 is devoted to an in-depth investigation of Marzullo's function $\boldsymbol{\mathcal{M}}$, which plays a central role in the analysis of $\boldsymbol{\mathcal{OA}}$ in Section 4 and Appendix A. Our major results, namely, worst case bounds for accu-

[2] Our experimental evaluation [HS97] of the failures of six different GPS timing receivers revealed an average error probability of about 10^{-6}, with several different failure modes.

racy and precision of OA, are provided in Section 5. Some conclusions in Section 6, an appended graph showing the interdependency of the major parts of the analysis, and a glossary eventually round off the paper.

2 Interval-based Clock Synchronization

The core idea of the interval-based paradigm introduced in [Mar84] is to represent real-time (= external time) t not just by a time-dependent local clock value $C(t)$, but rather by a *local interval clock* $\boldsymbol{C}(t) = [C(t)-\alpha^-(t), C(t)+\alpha^+(t)]$. Any $\boldsymbol{C}(t)$ must be maintained appropriately to secure the *accurateness property* $t \in \boldsymbol{C}(t)$, which is of course increasingly meaningful if $\boldsymbol{\alpha}(t)$ becomes small, and may be violated at all if $\boldsymbol{C}(t)$ is faulty. Note carefully that an interval, i.e., a range of values where t can lie, is the best deterministic information one can get in practice, since the exact value of t is usually not known explicitly: Even the *1 pulse-per-second* (1pps) output of a GPS timing receiver, which indicates something like "now it is 10:00", actually means "the real-time when the 1pps signal actually occurred lies somewhere within 10:00 $\pm$ 150 ns", see [Dan97], [HS97].

Interval clock readings $\boldsymbol{A} = [T - \alpha^-, T + \alpha^+] = [T \pm \boldsymbol{\alpha}]$ taken at some fixed real-time t_0, i.e., $\boldsymbol{A} = \boldsymbol{C}(t_0)$, as well as intervals derived from those by means of the basic operations introduced in Subsection 2.1 are called *accuracy intervals*. Herein, T denotes $\boldsymbol{A}$'s *reference point* and $\boldsymbol{\alpha}$ its *interval of accuracies* taken relatively to the reference point. Note that we employ the usual notation of lower case letters like t for real-time values and upper case letters like T for logical time ones.

2.1 Generic Algorithm

The system model of [SS97b] assumes a distributed system consisting of n *nodes*, which communicate with each other by message passing over a fully connected point-to-point or broadcast network. Each node is equipped with a *processor* (with integer arithmetic only) for executing the clock synchronization algorithm, a *network interface*, and a local interval clock $\boldsymbol{C}_p(t)$ that continuously displays p's local accuracy interval. Consult [SKM+99] for details of an advanced prototype implementation based upon our *Network Time Interface* M-Module.

An interval-based clock synchronization algorithm is in charge of maintaining $\boldsymbol{C}_p(t)$ in a way that secures the following properties:

(P) *Precision requirement*: There is some fixed *precision* $\pi_{\max} \geq 0$ such that $|C_p(t) - C_q(t)| \leq \pi_{\max}$ for all nodes p, q that are non-faulty up to real-time t.

(A) *Accuracy requirement*: The interval of accuracies $\boldsymbol{\alpha}_p(t)$ is such that $-\alpha_p^+(t) \leq C_p(t) - t \leq \alpha_p^-(t)$ for all nodes p that are non-faulty up to real-time t.

Note that we restrict our attention to $\alpha_p^+(t), \alpha_p^-(t) = \mathcal{O}(t)$, with the implied constant[3] M being (much) smaller than 1. This forces $C_p(t)$ to be within a linear envelope of real-time, thereby excluding "degenerated" cases like $C_p(t) \equiv 0$, see [DHS86].

[3] Throughout this paper, we use the $\mathcal{O}(.)$-notation to characterize the order of magnitude of neglected terms. An expression like $\alpha_p^-(t) = \mathcal{O}(t)$ means that there is some (reasonably small) fixed constant $M > 0$ such that $|\alpha_p^-(t)| \leq M|t|$.

The generic interval-based clock synchronization algorithm of Definition 9 in Appendix B employs the usual round-based structure of traditional internal synchronization algorithms to ensure properties (P) and (A). Starting from an initially synchronized state, any node periodically executes the following steps whenever its local clock reads kP, $k \geq 1$:

1. Initiation of a *full message exchange* (FME) to provide each node with the accuracy intervals of all other nodes within the system.
2. *Preprocessing* of the set of received accuracy intervals to make them all "compatible", that is, represent the same real-time.
3. Application of a suitable interval-valued *convergence function* to the set of preprocessed intervals to compute and subsequently apply a clock correction upon $\boldsymbol{C}_p(t)$ for resynchronization.
4. Keeping track of real-time by means of $\boldsymbol{C}_p(t)$ up to the next resynchronization.

The algorithm relies upon two basic operations called drift compensation and delay compensation. *Drift compensation* —required in Step 2 and 4 of the algorithm— allows to shift ("drag") an accuracy interval in time by means of the local clock while maintaining accurateness of the resulting interval. Since the local clock can drift with respect to real-time, this requires sufficient enlargement ("deterioration") of positive and negative accuracy according to [SS97b, Def. 5]. The following Figure 1 shows an example of drift compensated intervals based upon some initial interval $[T_0 \pm \boldsymbol{\alpha}]$ representing t_0, which is dragged to real-times t_i characterized by $C(t_i) = T_i = T_0 + i\Delta T$, $i \geq 1$. We assume a fast but deaccelerating clock with a maximum drift $\in [-\rho, \rho]$ here and ignore advanced issues like clock granularity and discrete rate adjustment uncertainty for simplicity. For accurateness, deterioration must ensure that any resulting interval $\boldsymbol{I}'$ intersects with the line $T = t$.

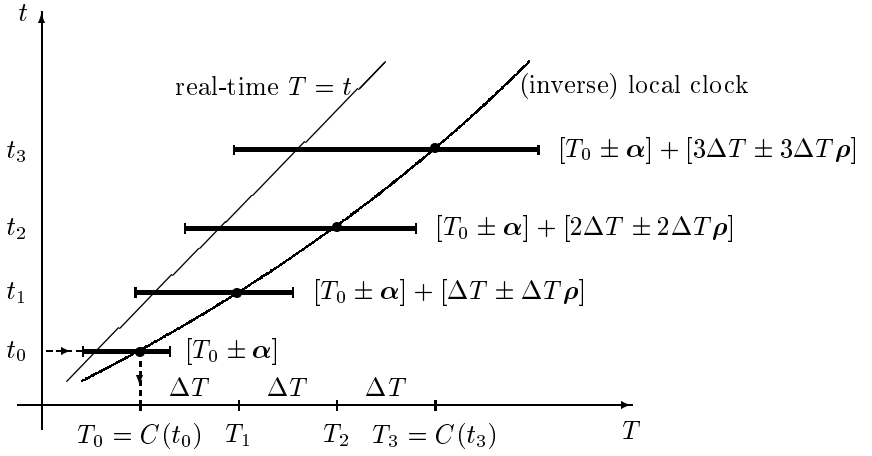

Figure 1: *Example of drift compensation in case of a fast but deaccelerating clock with max. drift* $\in [-\rho, \rho]$. *Proper deterioration ensures intersection with the line* $t = T$.

The *delay compensation* operation employed in Step 2 of the algorithm maintains accurateness of intervals that are transmitted over a network experiencing variable transmission delays $\delta' \in [\delta \pm \boldsymbol{\varepsilon}]$, with δ denoting the deterministic part of the delay. As before,

positive and negative accuracy must be enlarged according to [SS97b, Def. 6] to account for the maximum uncertainty in transmission delay $\boldsymbol{\varepsilon}$. This is illustrated in Figure 2, where it is assumed that the actual transmission delay is $\delta' > \delta$ and the sender node p's clock is drift-free, i.e., progresses as real-time does.

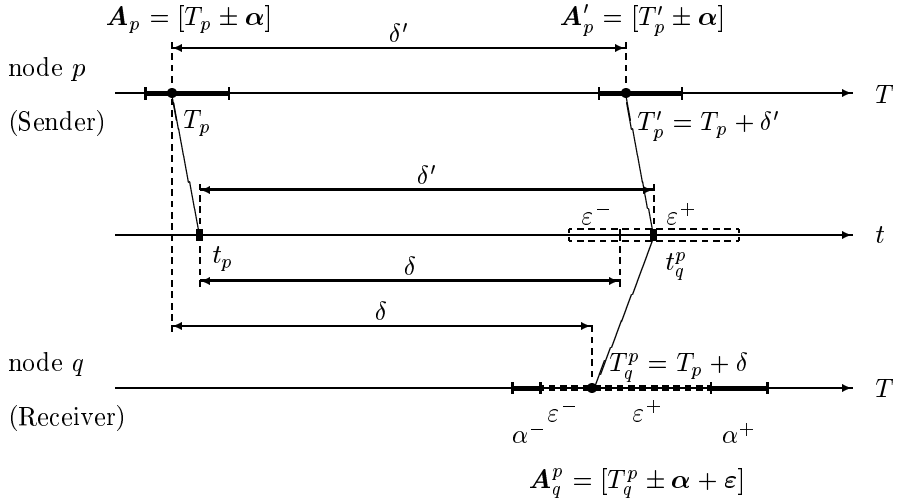

Figure 2: *Example of delay compensation when sending $\boldsymbol{I}$ from node p to q. The received interval must be enlarged by $\boldsymbol{\varepsilon}$ to secure accurateness.*

The middle time axis represents real-time, whereas the upper resp. lower one shows local time at node p resp. q. A line connecting two points at different axes, like T_p and t_p, indicates the "represents" relation, hence $t_p \in \boldsymbol{A}_p$ for accurateness. When interpreting Figure 2, one should consider the intervals $\boldsymbol{A}_p$ and $\boldsymbol{A}_q^p$ as "fixed" (since they do not depend upon the actual transmission delay δ'), whereas the reception time t_q^p and hence the interval $\boldsymbol{A}'_p$ vary with δ'. It is apparent that $t_q^p \in \boldsymbol{A}_q^p$ is always maintained if t_q^p remains within the dash-boxed region of the t-axis, that is, if $\delta' \in [\delta \pm \boldsymbol{\varepsilon}]$.

Together, drift compensation and delay compensation are employed to make all received intervals compatible, that is, represent a common real-time t_q^R: For an accuracy interval $\boldsymbol{A}_p$ sent by node $p \neq q$ at real-time t_p, delay compensation is applied to provide the receiver q with an interval $\boldsymbol{A}_q^p$ that covers the sender's current accuracy interval at the real-time of reception t_q^p. This interval $\boldsymbol{A}_q^p$ is then dragged locally by means of the receiver's clock to some (common) point in real-time t_q^R defined by $C_q(t_q^R) = T_q^R$, requiring an appropriate drift compensation to obtain the final accuracy interval $\boldsymbol{I}_q^p$. Provided that T_q^R is chosen large enough to ensure that the intervals of all non-faulty nodes can be received and processed, it follows by construction that $\boldsymbol{I}_q^p$ is accurate if (1) $\boldsymbol{A}_p$ was accurate, (2) transmission delay was not excessive, and (3) the receiver q is not faulty.

Finally, a suitable *convergence function* is applied to the set $\boldsymbol{\mathcal{I}}_q$ of node q's intervals $\boldsymbol{I}_q^p$ in Step 3, which provides a (small) interval that both contains real-time t_q^R (i.e., is accurate) and enhances precision — despite of some possibly faulty $\boldsymbol{I}_q^p$'s. Fortunately, only a few properties of the convergence function are actually required for predicting the above algorithm's worst case performance. Hence, the appropriate analysis in [SS97b] was conducted for a generic convergence function $\boldsymbol{\mathcal{CV}}$, which can be any interval-valued function that satisfies certain properties stated in Definition 11 in Appendix C.

2.2 Generic Precision Analysis

Since all operations used in our interval-based clock synchronization algorithm, namely, drift compensation, delay compensation, and finally application of the convergence function, are explicitly designed to preserve accurateness of the intervals involved, it is clear that accurateness of the local interval clocks $\boldsymbol{C}_p(t)$ of all non-faulty nodes is maintained during all rounds. Apart from the requirement of being accurate, which is a property of any single accuracy interval, however, we also have the precision requirement that applies to (the reference points of) a set of accuracy intervals. Precision and accuracy are in fact almost[4] independent of each other, since the reference point can in principle be placed arbitrarily within the accuracy interval. In the remainder of this section, we will briefly sketch the interval-based precision analysis framework established in [SS97b].

To that end, we first introduce the basic notation used throughout the paper: Our elementary objects are real intervals $\boldsymbol{I} = [x, y]$, $x \leq y$, with *lower edge* $x = \text{left}(\boldsymbol{I})$ and *upper edge* $y = \text{right}(\boldsymbol{I})$; the *empty interval* $\emptyset$ satisfies $\nexists t : t \in \emptyset$. For an interval $\boldsymbol{I} = [x, y]$, $|\boldsymbol{I}| = y - x$ denotes its *length* and $\text{center}(\boldsymbol{I}) = (x + y)/2$ its *centerpoint*. The *sum* of two intervals is defined by $[x, y] + [u, v] = [x + u, y + v]$, the *scalar product* by $s \cdot [x, y] = [sx, sy]$ for $s \geq 0$, and the *translation* by $\boldsymbol{I} + a = \boldsymbol{I} + [a, a] = [x + a, y + a]$ for some arbitrary scalar a. For two intervals $[x, y]$, $[u, v]$, the *intersection* is $[x, y] \cap [u, v] = [\max\{x, u\}, \min\{y, v\}]$ if $u \leq y$, $v \geq x$, and $\emptyset$ otherwise; the *union* reads $[x, y] \cup [u, v] = [\min\{x, u\}, \max\{y, v\}]$. Note that the definition of the union is also valid for $[x, y] \cap [u, v] = \emptyset$, hence incorporates the closure of two disjoint intervals as well. Note that both intersection and union extend to a scalar operand in the obvious way, i.e., $[x, y] \cup u = [x, y] \cup [u, u]$.

Accuracy intervals, as introduced at the beginning of this section, are intervals extended by a distinguished *reference point* r, which partitions the interval into a *negative accuracy* α^- and a *positive accuracy* α^+ according to

$$\boldsymbol{A} = [r \pm \boldsymbol{\alpha}] = [r - \alpha^-, r + \alpha^+]. \tag{1}$$

Herein, $\text{ref}(\boldsymbol{A}) = r$ denotes $\boldsymbol{A}$'s reference point, $\boldsymbol{\alpha} = [-\alpha^-, \alpha^+]$ its *interval of accuracies*, and $\alpha = |\boldsymbol{\alpha}| = \alpha^+ + \alpha^-$ its length. Note that we use bold letters like $\boldsymbol{A}$ for both ordinary intervals and accuracy intervals, since its actual type is usually clear from the context. Similarly, calligraphic bold letters like $\boldsymbol{\mathcal{A}}$ are used to denote a *set* of intervals or accuracy intervals.

For accuracy intervals $\boldsymbol{I} = [r \pm \boldsymbol{\alpha}]$ and $\boldsymbol{J} = [s \pm \boldsymbol{\beta}]$, we have $\text{left}(\boldsymbol{I}) = r - \alpha^-$, $\text{right}(\boldsymbol{I}) = r + \alpha^+$, $|\boldsymbol{I}| = \alpha^+ + \alpha^- = \alpha$, $\text{center}(\boldsymbol{I}) = r + (\alpha^+ - \alpha^-)/2$, $\boldsymbol{I} + \boldsymbol{J} = [r + s \pm \boldsymbol{\gamma}]$ where $\boldsymbol{\gamma} = \boldsymbol{\alpha} + \boldsymbol{\beta} = [-(\alpha^- + \beta^-), \alpha^+ + \beta^+]$, $\boldsymbol{I} + a = [r + a \pm \boldsymbol{\alpha}]$ for an arbitrary scalar a, and $s\boldsymbol{I} = [sr \pm \boldsymbol{\mu}]$ with $\boldsymbol{\mu} = s\boldsymbol{\alpha} = [-s\alpha^-, s\alpha^+]$ for any scalar $s \geq 0$. Finally, there is also a notation to express intervals obtained from (1) by *swapping* its positive and negative accuracy, namely

$$\overline{\boldsymbol{I}} = \overline{[r \pm \boldsymbol{\alpha}]} = [r \mp \boldsymbol{\alpha}] = [r - \alpha^+, r + \alpha^-] = [r \pm \overline{\boldsymbol{\alpha}}] \tag{2}$$

where $\overline{\boldsymbol{\alpha}} = [-\alpha^+, \alpha^-] = -\boldsymbol{\alpha}$.

Definition 1 (Interval Relations [SS97b, Def. 1]) *Accuracy intervals are categorized as follows:*

[4] They are not totally independent, as the computed reference point could lie outside the accuracy interval, see Figure 6.

(1) Two accuracy intervals $\boldsymbol{I} = \boldsymbol{I}(t_1)$ *representing* t_1 *and* $\boldsymbol{J} = \boldsymbol{J}(t_2)$ *representing* t_2 *are compatible iff* $t_1 = t_2$.

(2) Two compatible accuracy intervals $\boldsymbol{I}$ *and* $\boldsymbol{J}$ *are consistent iff* $\boldsymbol{I} \cap \boldsymbol{J} \neq \emptyset$.

(3) An accuracy interval $\boldsymbol{I} = \boldsymbol{I}(t)$ *representing real-time* t *is accurate iff* $t \in \boldsymbol{I}$.

Note that compatibility of two accuracy intervals means that they are comparable, i.e., represent the same real-time. Bear in mind, however, that the "represents"-relation neither implies consistency nor accurateness in case of faulty accuracy intervals.

The following definition of $\boldsymbol{\pi}$-precision is a key for our interval-based precision analysis. The underlying idea is to capture precision π of two clocks $C_p(t)$ and $C_q(t)$, that is, $|C_p(t) - C_q(t)| \leq \pi$, by means of consistency of suitably constructed precision intervals.

Definition 2 (Precision Intervals [SS97b, Def. 2]) *Given some fixed* $\boldsymbol{\pi} = [-\pi^-, \pi^+]$ *with* $\pi^-, \pi^+ \geq 0$ *and* $\pi = |\boldsymbol{\pi}| = \pi^- + \pi^+$, *and a set of* $n \geq 2$ *compatible accuracy intervals* $\boldsymbol{\mathcal{I}} = \{\boldsymbol{I}_1, \ldots, \boldsymbol{I}_n\}$ *with* $\boldsymbol{I}_j = [r_j \pm \boldsymbol{\alpha}_j]$, *the* $\boldsymbol{\pi}$-precision interval $\hat{\boldsymbol{I}}_j$ *associated with* $\boldsymbol{I}_j$ *is defined as* $\hat{\boldsymbol{I}}_j = [r_j \pm \boldsymbol{\pi}]$. *The set* $\boldsymbol{\mathcal{I}}$ *is called* $\boldsymbol{\pi}$*-precise iff* $\bigcap_{j=1}^{n} \hat{\boldsymbol{I}}_j \neq \emptyset$.

Figure 3 shows two accuracy intervals along with their associated precision intervals. Note carefully, however, that associated $\boldsymbol{\pi}$-precision intervals are not maintained "online", as are accuracy intervals, but rather computed from the reference points by adding the "static" values of $-\pi^-$, π^+ provided by the analysis. That is, it is sufficient for the clock synchronization algorithm to dynamically maintain the interval of accuracies and its reference point only.

Item (2) of the following Lemma 1 reveals that $\boldsymbol{\pi}$-precision implies (traditional) precision $\pi = |\boldsymbol{\pi}|$. Its assertions are quite trivial, as can be seen from the associated precision intervals shown in Figure 3.

Lemma 1 ($\boldsymbol{\pi}$-Precision vs. Precision) *Given a* $\boldsymbol{\pi}$*-precise set* $\boldsymbol{\mathcal{I}} = \{\boldsymbol{I}_1, \ldots, \boldsymbol{I}_n\}$ *of* $n \geq 2$ *compatible accuracy intervals* $\boldsymbol{I}_j = [r_j \pm \boldsymbol{\alpha}_j]$, *then*

(1) $|\hat{\boldsymbol{I}}_i \cup \hat{\boldsymbol{I}}_j| \leq 2\pi$ *for any* $1 \leq i, j \leq n$,

(2) $|r_i - r_j| \leq \pi$ *for any* $1 \leq i, j \leq n$.

Proof: See [SS97b], Lemma 3. □

Although the definition of $\boldsymbol{\pi}$-precision is a key for our precision analysis, there are only a few occasions where $\boldsymbol{\pi}$-precise intervals are encountered explicitly. In most cases, the slightly stronger predicate of $\boldsymbol{\pi}$-*accurateness* (implying $\boldsymbol{\pi}$-precision) is used, which is based on a suitable notion of an *internal global time* $\tau = \tau(t)$. More specifically, it was shown in [SS97b] that it makes sense to stipulate an "artificial" internal global time $\tau^k = \tau^k(t) = \tau_0^k + (t - t_0^k)$ for round k that progresses as real-time does. Of course, like real-time, internal global time is not directly accessible, and usually $\tau^k(t) \neq t$. However, internal global time of any fixed round is equivalent to real-time for specifying durations. Moreover, we we can of course unambiguously write $\hat{\boldsymbol{I}} = \hat{\boldsymbol{I}}(t) = \hat{\boldsymbol{I}}(\tau^k)$ for $\tau^k = \tau^k(t)$, meaning that $\hat{\boldsymbol{I}}$ represents $\tau^k(t)$ iff $\boldsymbol{I}$ represents t.

The concept of internal global time makes it possible to define an analogue to accurateness as follows:

Definition 3 (π-correctness [SS97b, Def. 3]) *For $\boldsymbol{\pi} = [-\pi^-, \pi^+]$ with $\pi^-, \pi^+ \geq 0$,*

(1) an accuracy interval $\boldsymbol{I} = \boldsymbol{I}(t)$ is $\boldsymbol{\pi}$-accurate (w.r.t. internal global time $\tau^k = \tau^k(t)$ of round k) iff the $\boldsymbol{\pi}$-precision interval $\hat{\boldsymbol{I}} = \hat{\boldsymbol{I}}(\tau^k)$ associated with $\boldsymbol{I}$ satisfies $\tau^k \in \hat{\boldsymbol{I}}$,

(2) an accuracy interval $\boldsymbol{I}$ is $\boldsymbol{\pi}$-correct (w.r.t. real-time and internal global time of round k) iff $\boldsymbol{I}$ is both $\boldsymbol{\pi}$-accurate and accurate,

(3) a set $\boldsymbol{\mathcal{I}}$ of compatible accuracy intervals is $\boldsymbol{\pi}$-correct if all members are $\boldsymbol{\pi}$-correct.

The following Figure 3 shows an example of two $\boldsymbol{\pi}$-correct intervals.

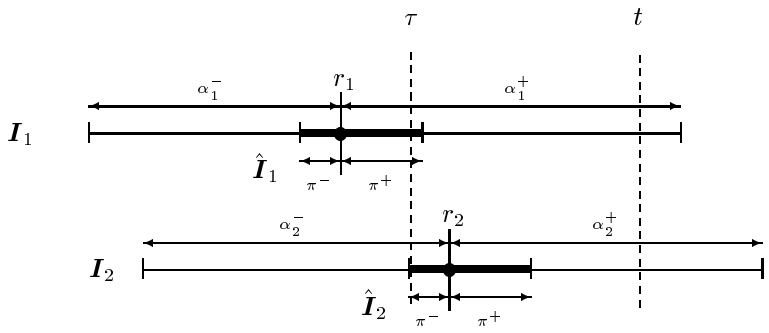

Figure 3: *Example of two $\boldsymbol{\pi}$-accurate intervals. Both accuracy intervals and (bold) $\boldsymbol{\pi}$-precision intervals are accurate with respect to t and τ, respectively.*

Now we are ready for explaining how the interval-based clock synchronization algorithm outlined at the beginning of this section maintains precision. The most important observation is that all steps of the algorithm, except the application of the convergence function, maintain precision automatically by maintaining accuracy. To understand why, assume that all members of the set $\boldsymbol{\mathcal{C}}(t)$ of non-faulty interval clocks are $\boldsymbol{\pi}_0$-correct at some real-time t^0 in round k, i.e., that their associated $\boldsymbol{\pi}_0$-precision intervals contain $\tau^k(t^0)$. In order to capture real-time t by $\boldsymbol{C}_p(t)$, it must be deteriorated (enlarged) appropriately to compensate for the drift of the local clock. However, if this is done properly to ensure $t' \in \boldsymbol{C}_p(t')$ for $t' > t^0$, then the associated $\boldsymbol{\pi}^H$-precision interval $\hat{\boldsymbol{C}}_p(t')$ captures internal global time $\tau^k(t') > \tau^k(t^0)$ as well, provided that $\boldsymbol{\pi}^H$ is the result of enlarging $\boldsymbol{\pi}_0$ by the maximum amount any $\boldsymbol{C}_q$ has been enlarged. This is a simple consequence of the fact that internal global time progresses as real-time does. Note anyway that enlargement of precision intervals is just a matter of analysis — the algorithm need not deal explicitly with precision intervals at all.

Whereas enlarging $\boldsymbol{\pi}_0$ to $\boldsymbol{\pi}^H$ guarantees $\boldsymbol{\pi}^H$-correctness of all non-faulty $\boldsymbol{C}_p(t')$, this cannot ensure bounded precision for $t \to \infty$. Periodic resynchronizations are required for this purpose, giving rise to our round-based algorithm. More specifically, at the end of the k-th round, the nodes' current $\boldsymbol{\pi}^H$-correct local interval clocks are set to newly computed accuracy intervals that are $\boldsymbol{\pi}_0$-precise for some $\boldsymbol{\pi}_0 \subset \boldsymbol{\pi}^H$ (= precision enhancement provided by the convergence function). Note that we cannot safely assume $\boldsymbol{\pi}_0$-correctness here, since it will usually happen that round k's internal global time τ^k does not lie in the intersection of the new $\boldsymbol{\pi}_0$-precision intervals. However, if a new "starting value" τ_0^{k+1} for internal global time $\tau^{k+1}(t)$ is defined for round $k+1$, $\boldsymbol{\pi}_0$-correctness

w.r.t. τ^{k+1} can of course be guaranteed. Consequently, resynchronization starts the next round $k+1$, during which initial precision $\boldsymbol{\pi}_0$ again deteriorates to $\boldsymbol{\pi}^H$.

Note carefully that only the interval clocks' associated $\boldsymbol{\pi}^H$-precision intervals experience (precision) enhancement during resynchronization. The local interval clocks $\boldsymbol{C}_p$ itself must continuously track real-time t, so that the accuracies could be monotonically increasing; the accuracy in round k can in fact be viewed as an accumulation of the $\boldsymbol{\pi}^H$-precision intervals during round $0, \ldots, k$. This eventually explains why t and τ will usually be apart, as mentioned earlier.

Generally speaking, the major advantage of the interval-based precision analysis conducted in [SS97b][5] is conceptual beauty and considerable flexibility with respect to incorporating non-standard features like clock granularity, broadcast latencies, etc. This is primarily a consequence of our notion of internal global time and $\boldsymbol{\pi}$-accurateness, which allows us to reason about precision by considering each local interval clock separately, i.e., without explicitly relating it to the other clocks in the system. Even more, there is no need to consider the "absolute position" of intervals, i.e., clock values, at all. In fact, any information required on some $\boldsymbol{I}(t) = [T \pm \boldsymbol{\alpha}]$ is provided by its interval of accuracies $\boldsymbol{\alpha}$ and the associated $\boldsymbol{\pi}$-precision interval: Since all (non-faulty) accuracy intervals must contain real-time t and internal global time τ by construction, the latter ones serve as a "common reference" for relating different intervals. Of course, the particular reference point $\mathrm{ref}(\boldsymbol{I})$ may lie anywhere in $[t-\alpha^+, t+\alpha^-]$ and $[\tau-\pi^+, \tau+\pi^-]$, according to the actually experienced clock drift, transmission delay, and initial accuracy, but there is no need for dealing with it explicitly.

2.3 Fault Model

In [SS97b], we argued that it does not make sense to stipulate a particular fault model for the generic algorithm, since it is primarily the convergence function that is concerned about faults. Now that we have to plug in a particular convergence function, however, we have to address this issue as well.

First of all, as in [SS97b], we assume that faulty nodes and network devices can take arbitrary steps and transmit (and "receive") any number of arbitrary messages. However, a faulty component must not be able to cause (serious) "global" disturbance of system operation, e.g. by impersonating other nodes or flooding/jamming the network.

Consequently, we require that a faulty sender node/network transmission may affect a receiving node only by means of its interval delivered during an FME. Note that faulty receiving nodes (which may behave arbitrarily anyway) can safely be ignored here, since the precision and accuracy properties (P) and (A) of Subsection 2.1 must be satisfied by non-faulty nodes only. Hence, it is perfectly reasonable to dispose of the usual system-global fault models, like the one that at most f nodes may crash during a round, in favor of one that just targets the intervals $\boldsymbol{\mathcal{I}}_q$ received at any non-faulty node q.

More specifically, our fault model rests upon pairs of "perceptions" $\boldsymbol{\mathcal{I}}_p$, $\boldsymbol{\mathcal{I}}_q$ of the intervals received during an FME at any two different nodes p and q. Any such pair of perceptions can be considered *in isolation*, that is, without taking into account the situation faced by other nodes: By the precision requirement (P), we only have to ensure

[5] Note that we provided a few extensions of the generic analysis [SS97b] in in Appendix C, which are used instead of the original version where required.

that any two non-faulty clocks satisfy $|C_p(t) - C_q(t)| \leq \pi$, independently of the other clocks in the system. Hence, system-global fault assumptions —like the one that all receivers perceive a certain fault consistently— are not at all required here. There is no need to restict the total number of faulty nodes during a round to, say, $\lfloor (n-1)/3 \rfloor$, but only the maximum number of faulty intervals in a pairwise perception.

The primary advantage of our fault model is proper handling of receive omissions: Receiver-originated faults like packet losses due to unrecognized packet headers or receiver overruns are quite likely in practice, but cannot "artificially" be attributed to sending nodes without sacrificing a reasonable maximum number f of tolerated failures. In fact, usual fault models enforce a maximum number f on the number of receive omissions within the whole system; the clever *a posteriori agreement* technique of [VRC97] even requires that there are broadcasts with $f = 0$. A more natural modeling of receive omissions, however, is to grant each receiving node a certain maximum number of those, independently of the situation at other nodes. However, allowing even a single receive omission at each node could eat up all n sending nodes, i.e., all n nodes were to be considered faulty in a conventional fault model. In our fault model, only two faulty pairs of intervals can show up due to this cause.

We start with the formal definition of faults of a single interval, which is primarily required for accuracy analysis.

Definition 4 (Single Faults) *An interval $\boldsymbol{I}$ representing t can suffer from the following faults:*

- Omission*: Missing interval, expressed by $\boldsymbol{I} = \emptyset$.*
- Non-accurate interval*: $t \notin \boldsymbol{I}$*
- Unbounded accuracy*: $t \in \boldsymbol{I}$ but $|\boldsymbol{I}|$ too large according to some condition (that need not be known explicitly).*

Remarks:

1. Non-accurate intervals can be caused by *timing faults* due to a faulty sending node/clock or excessive transmission delays, or by *accuracy faults* due to a faulty sending node/clock or a damaged message.

2. Masking or detecting —and thus ruling them out completely— unbounded accuracy faults is impossible in most circumstances. Indeed, although it is sometimes possible to determine the border between faulty and non-faulty accuracy values (see Theorem 4), it is nevertheless true that even limiting α^-, α^+ accordingly cannot prevent faulty nodes from considerably spoiling the "average" behavior.

3. Whereas it is usually impossible to decide locally whether an interval $\boldsymbol{I}$ is accurate or not, it is of course possible to detect omission faults. Hence, given a set $\boldsymbol{\mathcal{I}}$ of $n \geq 1$ compatible intervals with $f'_o \geq 0$ of them exhibiting omission faults, it is trivial to discard the f'_o omissive ones from $\boldsymbol{\mathcal{I}}$ and to proceed with the reduced set $\boldsymbol{\mathcal{J}}$ containing the $n' = n - f'_o$ non-empty intervals only.

For precision analysis, the single-interval faults of Definition 4 must be complemented by faults of *pairs of intervals* $\boldsymbol{I}_p^s$ resp. $\boldsymbol{I}_q^s$ obtained at nodes p resp. q in the broadcast from a single node s. Note that the classification of pairwise faults in Definition 5 is exhaustive.

Definition 5 (Pairwise Faults) *A pair of compatible accuracy intervals* $\{\boldsymbol{I}_p^s, \boldsymbol{I}_q^s\}$ *originating in a single sending node* s *and representing real-time* t *suffers from*

- *a* crash fault *iff* $\boldsymbol{I}_p^s = \boldsymbol{I}_q^s = \emptyset$,
- *a* symmetric fault *iff either*
  - *(1) both* $\boldsymbol{I}_p^s$ *and* $\boldsymbol{I}_q^s$ *are not accurate in the sense of* $t < \mathit{left}(\boldsymbol{I}_p^s)$ *and* $t < \mathit{left}(\boldsymbol{I}_q^s)$, *or else* $t > \mathit{right}(\boldsymbol{I}_p^s)$ *and* $t > \mathit{right}(\boldsymbol{I}_q^s)$,
  - *(2) without loss of generality,* $\boldsymbol{I}_p^s = \emptyset$ *and* $\boldsymbol{I}_q^s \neq \emptyset$ *does not suffer from an unbounded accuracy fault.*
- *an* asymmetric fault *iff either*
  - *(1) both* $\boldsymbol{I}_p^s$ *and* $\boldsymbol{I}_q^s$ *are not accurate in the sense of* $t > \mathit{right}(\boldsymbol{I}_p^s)$ *and* $t < \mathit{left}(\boldsymbol{I}_q^s)$ *or else* $t > \mathit{right}(\boldsymbol{I}_q^s)$ *and* $t < \mathit{left}(\boldsymbol{I}_p^s)$ *(true byzantine fault),*
  - *(2) without loss of generality,* $\boldsymbol{I}_p^s \neq \emptyset$ *is faulty and* $\boldsymbol{I}_q^s \neq \emptyset$ *is arbitrary (and none of the other faults applies).*

Remarks:

1. A crash fault causes an omission at both receiving nodes p and q. Note carefully, though, that it is impossible for either node to decide locally whether its omission is due to a crash fault or a more severe receive omission.
2. A symmetric fault is one that is perceived similarly at p and q (or, alternatively, involves a receive omission at one of them). This type of fault is usually produced by a sender clock that runs too slow or too fast.
3. The "classical" asymmetric fault is one that is perceived differently at p and q. Its distinguishing property is that node p arrives at the conclusion that the sender's clock is, say, too fast, whereas q thinks that it is too slow (or correct). This could occur, for example, when the transmission delay to p resp. q is excessively low resp. high or if the sending node exhibits byzantine behavior. However, an asymmetric fault also occurs in presence of an unbounded accuracy fault, which is due to the poor fault-tolerance degree of our convergence function in this respect, see Remark 2 on Lemma 2.
4. Note that Definition 5 does not cover the case where a more severe fault comes out as a less severe one. For example, it is reasonable to assume that an asymmetric fault could just be a symmetric or even a crash fault only. In this paper, we will typically use phrases like "asymmetric (or weaker) fault" to indicate such extensions.

Introducing different classes of faults as in Definition 5 is known as a *hybrid fault model* in literature, cf. [AK96]. It allows us to exploit the fact that masking f symmetric faults requires only $n \geq 2f+1$ nodes, whereas $n \geq 3f+1$ are needed if all faults are asymmetric ones. Since a large number of asymmetric faults is very unlikely in practice, cf. [Sch95], this effectively leads to a smaller n for tolerating a given number of faults, see (3).

Still, it is apparent that the above definitions are meaningful only for "pure" accuracy intervals and precision intervals, but not for the "combination" of both required

for $\mathcal{OA}$'s analysis in Section 4, recall Figure 3. For our final fault model, we have to take into account that faults may occur in both quite independently of each other. More specifically, we must distinguish faults affecting an accuracy interval and its associated precision interval either in the same way or differently: Let a single accuracy interval $\boldsymbol{I}$ that is faulty w.r.t. real-time t and/or internal global time τ be called t/τ-symmetrically faulty if $t < \text{left}(\boldsymbol{I})$ and $\tau < \text{left}(\hat{\boldsymbol{I}})$ or either $t > \text{right}(\boldsymbol{I})$ and $\tau > \text{right}(\hat{\boldsymbol{I}})$, and t/τ-asymmetrically faulty otherwise; a set $\boldsymbol{\mathcal{F}}$ of faulty accuracy intervals is *identically* t/τ-symmetrically faulty if t is either to the left or to the right for all members of $\boldsymbol{\mathcal{F}}$.

Assumption 1 (Hybrid Fault Model $\mathcal{F}$) *Let a pair of accuracy intervals $\{\boldsymbol{I}_p^s, \boldsymbol{I}_q^s\}$ originating in a single sending node s and representing real-time t, with the associated precision intervals $\{\hat{\boldsymbol{I}}_p^s, \hat{\boldsymbol{I}}_q^s\}$ representing τ, be called*

(1) simple faulty *if it suffers from a crash fault or a symmetric fault w.r.t. t and/or τ and all faulty intervals involved are identically t/τ-symmetrically faulty,*

(2) arbitrary faulty *if it suffers either from an asymmetric fault w.r.t. t and/or τ, or a symmetric fault involving at least one t/τ-asymmetrically faulty interval. Alternatively, an arbitrary fault could also be just a simple one.*

For all pairs of non-faulty nodes p and q, consider the sets of intervals $\boldsymbol{\mathcal{I}}_p = \{\boldsymbol{I}_p^1, \ldots, \boldsymbol{I}_p^n\}$ and $\boldsymbol{\mathcal{I}}_q = \{\boldsymbol{I}_q^1, \ldots, \boldsymbol{I}_q^n\}$ obtained after reception and preprocessing of the accuracy intervals disseminated in an FME. We assume that at most f_a resp. f_s of the n pairs of intervals $\{\boldsymbol{I}_p^s, \boldsymbol{I}_q^s\}$, $1 \leq s \leq n$, suffer from arbitrary resp. simple faults, where f_a and f_s are such that

$$n \geq 3f_a + 2f_s + 1. \tag{3}$$

3 Marzullo's Function

This section is devoted to an in-depth investigation of a fault-tolerant intersection function $\boldsymbol{\mathcal{M}}$, which plays a central role in our orthogonal accuracy convergence function. $\boldsymbol{\mathcal{M}}$ was introduced in Marzullo's thesis [Mar84] and termed *Marzullo function* in [Lam87]. Although its relevance was recognized in several papers (e.g. [Lam87], [Mar90], [OSF92], [Sch95], [Mil95]), it has been studied thoroughly in the context of replicated sensors only, see [Mar90] and [BI96]. For clock synchronization purposes, a number of additional properties are required, which will be established subsequently.

Definition 6 (Marzullo Function) *Given a set $\boldsymbol{\mathcal{I}} = \{\boldsymbol{I}_1, \ldots, \boldsymbol{I}_n\}$ of $n \geq 1$ (non-empty) compatible intervals with at least $n - f \geq 1$ of the intervals being accurate, $\boldsymbol{\mathcal{M}}_n^{n-f}(\boldsymbol{\mathcal{I}})$ is defined as the largest interval whose edges lie in the intersection of at least $n - f$ different $\boldsymbol{I}_j$'s.*

Therefore, to compute the left resp. right edge of $\boldsymbol{\mathcal{M}}_n^{n-f}(\boldsymbol{\mathcal{I}})$, one has to "sweep" over the set of intervals from left to right resp. right to left and stop when $n - f$ intervals intersect for the first time. $\boldsymbol{\mathcal{M}}$ can thus be computed in $\mathcal{O}(n \log n)$ time by sorting the intervals' edges, cf. [Mar90]. Figure 4 shows an example for $n = 4$ and $f = 1$. Note that the unknown t cannot lie in the region between $\text{right}(\boldsymbol{I}_3)$ and $\text{left}(\boldsymbol{I}_4)$ in this example.

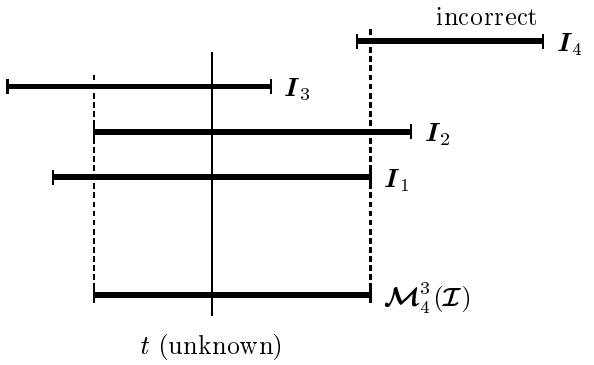

Figure 4: *Example of the Marzullo function* $\mathcal{M}$ *for* $n = 4$ *and* $f = 1$*. The edges of the result lie in* $n - f = 3$ *input intervals.*

However, since there is no way to decide whether t lies in the area left or right of this region, both areas must be covered by $\mathcal{M}_4^3$.

The most important feature of $\mathcal{M}$ is fault-tolerance with respect to faulty input intervals. For example, Figure 4 shows that $\mathcal{M}_4^3$ provides an accurate result despite of the fact that $\boldsymbol{I}_4$ was non-accurate. We will now embark formally on $\mathcal{M}$'s capabilities in this respect, according to the following roadmap:

1. We start with the "single node case", where $\mathcal{M}$'s result $\boldsymbol{M}_p$ computed at a particular non-faulty node p is considered in isolation. It suffices to account for single faults according to Definition 4 here, which affect $\mathcal{M}$'s input intervals at node p.

   - Lemma 2 gives the number of non-faulty input intervals required for tolerating a certain number of faults, and provides both a lower and an upper bound on $\boldsymbol{M}_p$.
   - Lemma 3 provides a few monotonicity properties of $\mathcal{M}$, that is, upper bound results like $\mathcal{M}_n^{n-f}(\mathcal{I}) \subseteq \mathcal{M}_n^{n-(f+k)}(\mathcal{I})$.

   Note that those results are primarily required for accuracy analysis.

2. For precision analysis, we also require pairwise properties, i.e., statements relating the results $\boldsymbol{M}_p$ and $\boldsymbol{M}_q$ of $\mathcal{M}$ computed at two different nodes p and q. Therefore, we have to consider pairwise faults according to Definition 5 here, which affect the pair $\{\boldsymbol{I}_p^s, \boldsymbol{I}_q^s\}$ received in the broadcast from node s.

   - Lemma 4 gives the number of non-faulty pairs of intervals required for tolerating a certain number of faults, and provides both a lower bound on $\boldsymbol{M}_p \cap \boldsymbol{M}_q$ and an upper bound on $\boldsymbol{M}_p \cup \boldsymbol{M}_q$.
   - Lemma 5 adopts Lemma 4 to a slightly extended fault model and, most importantly, shows what happens to $\boldsymbol{M}_p \cup \boldsymbol{M}_q$ when the failure assumptions are violated.

We should add here that all lemmas of this section deal with ordinary intervals only (without reference points). In fact, unlike in the analysis of $\mathcal{OA}$ in Section 4, there is

no need to deal with fully-fledged $\boldsymbol{\pi}$-accurate intervals (incorporating both precision and accuracy intervals) according to Figure 3.

The following Lemma 2 reveals how $\boldsymbol{\mathcal{M}}$ behaves in the presence of faults according to Definition 4. Extending the accomplishments of [Mar84] and [Mar90], it answers the question of how many non-faulty intervals are required for tolerating a certain number of non-accurate intervals (f_n) and unbounded accuracy faults (f_u). The most important property shown is that $\boldsymbol{\mathcal{M}}$'s result lies within the intersection of $n - 2f_n - 3f_u \geq 1$ non-faulty input intervals.

Lemma 2 (Accuracy $\boldsymbol{\mathcal{M}}$) *Let $\boldsymbol{\mathcal{J}} = \{\boldsymbol{J}_1, \ldots, \boldsymbol{J}_n\}$ be a set of $n \geq 1$ non-empty compatible accuracy intervals representing t, and define w^h to be the length of the largest intersection of $h \geq 1$ non-faulty intervals among them, i.e., $w^h = \max\{|\boldsymbol{W}| : \boldsymbol{W} \in \boldsymbol{\mathcal{W}}^h\}$ for*

$$\boldsymbol{\mathcal{W}}^h = \{\boldsymbol{W} : \boldsymbol{W} = \bigcap_{i=1}^{h} \boldsymbol{J}_{w_i} \quad \textit{with indices } w_i \neq w_j \textit{ for } i \neq j \textit{ and } \boldsymbol{J}_{w_i} \in \boldsymbol{\mathcal{J}} \textit{ being non-faulty}\}.$$

If $f'_u \geq 0$ of the $\boldsymbol{J}_j$ suffer from unbounded accuracy faults and $f'_n \geq 0$ are non-accurate, where $f'_u \leq f_u$ and $f'_n \leq f_n$ with $f'_u + f'_n = f' \leq f_u + f_n = f < n$ (so that $n - f' \geq n - f > 0$ of the n intervals are non-faulty), then

(1) $\boldsymbol{M} = \boldsymbol{\mathcal{M}}_n^{n-f}(\boldsymbol{\mathcal{J}})$ is accurate and contains any intersection $\boldsymbol{W} \in \boldsymbol{\mathcal{W}}^{n-f}$ of $n - f \geq 1$ different non-faulty input intervals $\boldsymbol{J}_{w_1}, \ldots, \boldsymbol{J}_{w_{n-f}}$, i.e.,

$$\boldsymbol{W} = \bigcap_{j=1}^{n-f} \boldsymbol{J}_{w_j} \subseteq \boldsymbol{M}, \tag{4}$$

so that $|\boldsymbol{M}| \geq w^{n-f}$ (minimal intersection property),

(2) there are at least $n - 2f - f'_u \geq n - 2f - f_u$ different non-faulty input intervals $\boldsymbol{J}_{b_1}, \ldots, \boldsymbol{J}_{b_{n-2f-f'_u}} \in \boldsymbol{\mathcal{J}}$ such that

$$\boldsymbol{M} \subseteq \bigcap_{j=1}^{n-2f-f'_u} \boldsymbol{J}_{b_j} \subseteq \bigcap_{j=1}^{n-2f-f_u} \boldsymbol{J}_{b'_j}, \tag{5}$$

where the set of indices $\{b'_j\}_{1 \leq j \leq n-2f-f_u}$ is obtained from $\{b_j\}_{1 \leq j \leq n-2f-f'_u}$ by discarding $f_u - f'_u$ elements. Hence, $|\boldsymbol{M}| \leq w^{n-2f-f'_u} \leq w^{n-2f-f_u}$.

(3) there are at least $f - f' + 1 \geq 1$ non-faulty intervals $\boldsymbol{J}_{\ell_k}$ resp. $\boldsymbol{J}_{r_k}$, $1 \leq k \leq f - f' + 1$, in $\boldsymbol{\mathcal{J}}$ satisfying left($\boldsymbol{M}$) $\leq$ left($\boldsymbol{J}_{\ell_k}$) resp. right($\boldsymbol{M}$) $\geq$ right($\boldsymbol{J}_{r_k}$).

Proof: Since $\boldsymbol{M} = \boldsymbol{\mathcal{M}}_n^{n-f}(\boldsymbol{\mathcal{J}})$ contains any intersection of at least $n - f$ input intervals by definition, it obviously contains any intersection of $n - f$ non-faulty intervals $\boldsymbol{W} = \bigcap_{j=1}^{n-f} \boldsymbol{J}_{w_j} \in \boldsymbol{\mathcal{W}}^{n-f}$. Note that w_j, $1 \leq j \leq n-f$, just denote the indices of the contributing intervals w.r.t. $\boldsymbol{\mathcal{J}}$ here. Therefore, it follows that $t \in \boldsymbol{M}$ and $|\boldsymbol{M}| \geq w^{n-f}$ as asserted in item (1) of the lemma.

Turning our attention to item (2), it is apparent that the total number of intersections of left and right edge of $\boldsymbol{M}$ with non-faulty input intervals is $g'_l + g'_r \geq 2(n-f) - 2f'_u - f'_n$,

because an interval $\boldsymbol{J}_j$ suffering from an unbounded accuracy fault (f'_u) could intersect with both edges of $\boldsymbol{M}$, whereas a non-accurate interval (f'_n) can only intersect with one edge of $\boldsymbol{M}$ due to $t \notin \boldsymbol{J}_j$ but $t \in \boldsymbol{M}$. However, since there are only $g' = n - f'$ different non-faulty intervals in $\boldsymbol{\mathcal{J}} = \{\boldsymbol{J}_1, \ldots, \boldsymbol{J}_n\}$, the pigeonhole principle reveals that

$$g'_l + g'_r - g' \geq 2n - 2f - 2f'_u - f'_n - n + f' = n - 2f - f'_u$$

of the intersected accurate intervals, say $\boldsymbol{J}_{b_1}, \ldots, \boldsymbol{J}_{b_{n-2f-f'_u}}$, must be the same. Therefore, $\boldsymbol{M}$ must lie in the intersection of those intervals and $|\boldsymbol{M}| \leq w^{n-2f-f'_u}$ as asserted. The upper bound in (5) follows immediately from $f'_u \leq f_u$.

Finally, to prove item (3), consider w.l.o.g. the left edge of $\boldsymbol{M}$. We show by contradiction that the left edge of at least $f - f' + 1$ non-faulty intervals lies at or right of left($\boldsymbol{M}$): If there were at most $f - f'$ such intervals, all $n - f' - (f - f') = n - f$ remaining non-faulty intervals would have their left edge strictly left of left($\boldsymbol{M}$). However, this would contradict the minimal intersection property established in item (1), completing the proof of our lemma. □

Remarks:

1. We excluded omission faults in our lemma, since $\boldsymbol{\mathcal{M}}$ as defined in Definition 6 cannot deal with empty intervals. However, as mentioned in Remark 3 following Definition 4, intervals with omission faults can of course be discarded before $\boldsymbol{\mathcal{M}}$ is applied. Therefore, if f'_o of presumed n intervals suffer from an omission fault, we just have to set $n := n - f'_o$ and $f := f - f'_o$ in Lemma 2 to obtain the results for this case as well. Note that it is feasible to let f depend on f'_o, see Lemma 4 below.

2. Interpreting item (2) of Lemma 2 and the previous remark in terms of the usual fault-tolerance degree notion, it follows that $n \geq f'_o + 2f + f'_u + 1$ nodes are required to guarantee that $\boldsymbol{M}$ remains bounded by the length of at least one non-faulty input interval. Hence, as many as

$$n \geq \begin{cases} f'_o + 1 & \text{for } f'_o \text{ omission faults,} \\ 2f_n + 1 & \text{for } f'_n \leq f_n \text{ non-accurate intervals,} \\ 2f_u + f'_u \leq 3f_u & \text{for } f'_u \leq f_u \text{ unbounded accuracy faults} \end{cases}$$

   nodes are required for tolerating faults of the given type. It is thus apparent that $\boldsymbol{\mathcal{M}}$ can tolerate $\lfloor (n-1)/2 \rfloor$ non-accurate intervals but only $\lfloor (n-1)/3 \rfloor$ intervals that suffer from unbounded accuracy faults, cf. [Mar90]. Note carefully that the numbers above do not solely depend on the *actual* number of faults (e.g., f'_u), but also on their maximum number (e.g., f_u); this is due to the fact that the latter is compiled into the superscript argument of $\boldsymbol{\mathcal{M}}$.

3. The lower bound on $|\boldsymbol{M}|$ in item (1) expresses the rather obvious fact that $\boldsymbol{\mathcal{M}}$ cannot improve the accuracy beyond the one "hidden" in the input intervals; the term *minimal intersection property* was coined in [Mar84]. Note that $\boldsymbol{M}$ contains *any* intersection of $n - f$ intervals, hence includes intersections involving unbounded accuracy faults as well.

4. Item (3) just says that $\boldsymbol{M}$ contains the left and right edge of at least one (not necessarily the same) non-faulty interval.

The following Lemma 3 establishes a few useful monotonicity relations with respect to both parameters and input arguments of $\mathcal{M}$.

Lemma 3 (Monotonicity $\mathcal{M}$) *Let $\mathcal{I} = \{\boldsymbol{I}_1, \ldots, \boldsymbol{I}_n\}$ be a set of $n > f \geq 0$ compatible non-empty accuracy intervals representing t, with f', $0 \leq f' \leq f$, faulty ones among them. Then, $\mathcal{M}_n^{n-f}(\mathcal{I})$ is accurate and satisfies the following monotonicity relations:*

(1) $\mathcal{M}_n^{n-f}(\mathcal{I}) \subseteq \mathcal{M}_n^{n-(f+k)}(\mathcal{I})$ for any integer k with $0 \leq k < n - f$,

(2) $\mathcal{M}_n^{n-f}(\mathcal{I}) \subseteq \mathcal{M}_n^{n-f}(\mathcal{J})$ for any $\mathcal{J} = \{\boldsymbol{J}_1, \ldots, \boldsymbol{J}_n\}$ with $\boldsymbol{I}_l \subseteq \boldsymbol{J}_l$ for $1 \leq l \leq n$,

(3) For $f \geq f' \geq 1$, if $\mathcal{L} = \mathcal{I} \backslash \{\boldsymbol{I}_j\}$ is obtained by discarding some faulty interval $\boldsymbol{I}_j$ from $\mathcal{I}$, $\mathcal{M}_{n-1}^{(n-1)-(f-1)}(\mathcal{L}) = \mathcal{M}_{n-1}^{n-f}(\mathcal{L})$ is accurate and satisfies

$$\mathcal{M}_{n-1}^{n-f}(\mathcal{L}) \subseteq \mathcal{M}_n^{n-f}(\mathcal{I}). \tag{6}$$

Proof: In the proof of item (1) of Lemma 2, we argued that $\boldsymbol{M} = \mathcal{M}_n^{n-f}(\mathcal{I})$ contains any intersection of at least $n - f$ input intervals by definition, hence must be accurate. Since the interval containing any intersection of $n - f - k$ input intervals obviously contains any intersection of $n - f$ input intervals, item (1) of the lemma follows.

Turning our attention to item (2), it is clear that any particular intersection of $n - f$ intervals in $\mathcal{J}$ contains the intersection of the corresponding intervals in $\mathcal{I}$ if it is non-empty at all. By the same token as before, it hence follows that $\mathcal{M}_n^{n-f}(\mathcal{J})$ must contain $\mathcal{M}_n^{n-f}(\mathcal{I})$ as well.

For proving item (3), the same argument is used again. First of all, $n > f \geq 1$ implies that $n-1 > f-1 \geq 0$, hence $\boldsymbol{M}_{n-1}^{(n-1)-(f-1)}(\mathcal{L})$ is accurate. Moreover, discarding a faulty $\boldsymbol{I}_j$ in $\mathcal{I}$ and simultaneously reducing f by 1 leaves the superscript argument $(n-1)-(f-1) = n - f$ in $\mathcal{M}_{n-1}^{n-f}(\mathcal{L})$ unchanged. Since the interval containing any intersection of $n - f$ input intervals in $\mathcal{I}$ must contain the interval containing any intersection of $n-f$ intervals in $\mathcal{L} \subseteq \mathcal{I}$, the statement given in item (3) of our lemma follows. □

Remarks:

1. It is immediately apparent from the definition of $\mathcal{M}$ that $\mathcal{M}_n^1(\mathcal{I}) = \bigcup_i \boldsymbol{I}_i$ and $\mathcal{M}_n^n(\mathcal{I}) = \bigcap_i \boldsymbol{I}_i$, hence $\mathcal{M}_n^{n-f}$ "changes" from union to intersection as $n - f$ goes from 1 to n.

2. It is not difficult to show that $\mathcal{M}$ is optimal with respect to worst case accuracy in presence of non-accurate intervals (but not for unbounded accuracy faults, see Remark 2 on Lemma 2) among all interval-valued functions of n interval arguments, as pointed out already in [Lam87]: Suppose there were a function $\mathcal{F}$ that provides an accurate interval satisfying $\mathcal{M}(\mathcal{I}) \not\subseteq \mathcal{F}(\mathcal{I})$, then $\boldsymbol{M}' = \mathcal{M}(\mathcal{I}) \backslash (\mathcal{M}(\mathcal{I}) \cap \mathcal{F}(\mathcal{I})) \neq \emptyset$ and there must be an intersection of $n - f$ accurate intervals $\boldsymbol{A} = \bigcap_{i=1}^{n-f} \boldsymbol{I}_{b_i}$ so that $\boldsymbol{A} \cap \boldsymbol{M}' \neq \emptyset$. However, the (valid) assumption that $t \in \boldsymbol{A} \cap \boldsymbol{M}'$ reveals that $\mathcal{F}(\mathcal{I})$ cannot be accurate, providing the required contradiction.

3. Regarding item (2) of Lemma 3, it is important to note that enlarging an input interval (even by a minor amount) can cause a discontinuous jump of an edge of $\mathcal{M}_n^{n-f}(\mathcal{J})$ if a "new" intersection of $n-f$ intervals comes up. Just consider shrinking or moving right the faulty interval $\boldsymbol{I}_4$ in Figure 4, which causes the right edge

of $\mathcal{M}_4^3(\mathcal{I})$ to shrink to right(I_3) as soon as left(I_4) > right(I_1). This implies that $\mathcal{M}$ does not satisfy a Lipschitz condition with respect to moving (edges of) input intervals, as already noted in [Lam87].

4. Item (3) of Lemma 3 implies that one should always try to detect and discard faulty intervals before $\mathcal{M}$ is applied, since this can only improve the result.

For establishing precision results, we also require certain "pairwise" properties of $\mathcal{M}$, i.e., statements relating the results M_p and M_q of $\mathcal{M}$ computed at different nodes p and q. This is provided by the following Lemma 4, which is an advanced version of a lemma introduced in [Sch95]. It gives the number of non-faulty pairs of intervals required for tolerating a certain number of

- crash faults ($f_c' \leq f_c$),
- symmetric faults ($f_s' \leq f_s$),
- asymmetric faults ($f_a' \leq f_a$).

The most important result of Lemma 4 is an upper bound on the union $M_p \cup M_q$, which must lie within at least $n - \min\{f_c' + f_s', 2f_c - f_c'\} - 2f_s - 3f_a \geq 1$ unions $I_p^s \cup I_q^s$ of non-faulty input intervals. Note that the union takes into account that two different nodes p and q usually receive slightly different intervals in the broadcast of a single node s, even if there is no fault.

Lemma 4 (Precision $\mathcal{M}$) *Let $\mathcal{I}_p = \{I_p^1, \ldots, I_p^n\}$ and $\mathcal{I}_q = \{I_q^1, \ldots, I_q^n\}$ be two ordered sets of $n > f_c + f_s + f_a$, $f_c, f_s, f_a \geq 0$, compatible (or empty) accuracy intervals representing t, where $f_a' \leq f_a$, $f_s' \leq f_s$, and $f_c' \leq f_c$ of the n pairs of intervals $\{I_p^i, I_q^i\}$ exhibit asymmetric, symmetric, and crash faults, respectively, and the remaining ones are non-faulty. Define u^h resp. v^h to be the length of the largest intersection of $h \geq 1$ unions resp. intersections of pairs of non-faulty intervals, formally $u^h = \max\{|U| : U \in \mathcal{U}_{pq}^h\}$ resp. $v^h = \max\{|V| : V \in \mathcal{V}_{pq}^h\}$ for*

$$\mathcal{U}_{pq}^h = \{U : U = \bigcap_{i=1}^{h} I_p^{u_i} \cup I_q^{u_i} \quad \text{with } u_i \neq u_j,\ i \neq j, \text{ and non-faulty } I_p^{u_i} \in \mathcal{I}_p,\ I_q^{u_i} \in \mathcal{I}_q\}$$

$$\mathcal{V}_{pq}^h = \{V : V = \bigcap_{i=1}^{h} I_p^{v_i} \cap I_q^{v_i} \quad \text{with } v_i \neq v_j,\ i \neq j, \text{ and non-faulty } I_p^{v_i} \in \mathcal{I}_p,\ I_q^{v_i} \in \mathcal{I}_q\}.$$

With d_p, $0 \leq d_p \leq f_s'$, denoting the number of empty intervals caused by symmetric faults at node p, let $\mathcal{J}_p = \{J_1, \ldots, J_{n_p}\}$ be the set of $n_p = n - o_p$ non-empty intervals obtained from $\mathcal{I}_p$ by discarding any of the present $o_p = f_c' + d_p \leq f_c + f_s$ empty intervals caused by crash and symmetric faults. Using the upper bound $f_p = f_s + f_a - \max\{0, o_p - f_c\}$ on the number of intervals in J_p that (still) may be faulty in presence of o_p omissions, define

$$M_p = \mathcal{M}_{n_p}^{n_p - f_p}(\mathcal{J}_p) \qquad \text{and analogously} \qquad M_q = \mathcal{M}_{n_q}^{n_q - f_q}(\mathcal{J}_q). \tag{7}$$

Then,

(1) both M_p and M_q are accurate and

$$M_p \cap M_q \supseteq \bigcap_{j=1}^{n - f_c' - f_s - f_a} I_p^{v_j} \cap I_q^{v_j} = V \tag{8}$$

for any possible subset $\boldsymbol{V} \in \boldsymbol{\mathcal{V}}_{pq}^{n-f'_c-f_s-f_a}$, *so that* $|\boldsymbol{M}_p \cap \boldsymbol{M}_q| \geq v^{n-f'_c-f_s-f_a}$ *(distributed minimal intersection property),*

(2) there are at least $n - \min\{f'_c + f'_s, 2f_c - f'_c\} - 2f_s - 2f_a - f'_a$ *pairs of non-faulty intervals* $\{\boldsymbol{I}_p^{u_k}, \boldsymbol{I}_q^{u_k}\}$ *with* $\boldsymbol{I}_p^{u_k} \in \boldsymbol{\mathcal{J}}_p$ *and* $\boldsymbol{I}_q^{u_k} \in \boldsymbol{\mathcal{J}}_q$ *such that*

$$\boldsymbol{M}_p \cup \boldsymbol{M}_q \subseteq \bigcap_{k=1}^{n-\min\{f'_c+f'_s,2f_c-f'_c\}-2f_s-2f_a-f'_a} \boldsymbol{I}_p^{u_k} \cup \boldsymbol{I}_q^{u_k} \tag{9}$$

and hence $|\boldsymbol{M}_p \cup \boldsymbol{M}_q| \leq u^{n-\min\{f'_c+f'_s,2f_c-f'_c\}-2f_s-2f_a-f'_a}$.

Proof: First of all, we note that f_p gives indeed an upper bound on the number of intervals in $\boldsymbol{\mathcal{J}}_p$ that still may be faulty in presence of $o_p = f'_c + d_p \leq f_c + f'_s \leq f_c + f_s$ omissions, since $f_p = f_s + f_a$ if $o_p \leq f_c$, and $f_p = f_s + f_a - (o_p - f_c)$ otherwise (accounting for $o_p - f_c > 0$ symmetric faults that must have caused omissions at node p), hence

$$f_p \leq f_s + f_a. \tag{10}$$

Evidently, at least $n_p - f_p$ of the intervals in $\boldsymbol{\mathcal{J}}_p$ must be non-faulty. Rewriting the definition

$$n_p - f_p = n - o_p - f_s - f_a + \max\{0, o_p - f_c\} = n - f_s - f_a + \max\{-o_p, -f_c\} \tag{11}$$

and applying $\max\{0, x\} \geq x$ for any x, and the simple fact that $\max\{-o_p, -f_c\} \leq -f'_c$ since obviously $o_p \geq f'_c$ and $f'_c \leq f_c$, it follows easily that

$$n - f_c - f_s - f_a \leq n_p - f_p \leq n - f'_c - f_s - f_a \leq n - f_s - f_a. \tag{12}$$

Of course, analogous bounds hold for $n_q - f_q$.

Lemma 2 is applicable, and it follows that $\boldsymbol{M}_p$ and $\boldsymbol{M}_q$ are both accurate and satisfy the (local) minimal intersection property. That is, $\boldsymbol{M}_p$ contains any intersection of $n_p - f_p \leq n - f'_c - f_s - f_a$ non-faulty intervals present in $\boldsymbol{\mathcal{J}}_p$. If $\{v_j\}_{1 \leq j \leq n-f'_c-f_s-f_a}$ denotes any set of different indices of non-faulty pairs of intervals $\boldsymbol{I}_p^{v_j} \in \boldsymbol{\mathcal{I}}_p$, $\boldsymbol{I}_q^{v_j} \in \boldsymbol{\mathcal{I}}_q$ (of course also present in $\boldsymbol{\mathcal{J}}_p$, $\boldsymbol{\mathcal{J}}_q$), we thus have

$$\boldsymbol{W}_p = \bigcap_{j=1}^{n-f'_c-f_s-f_a} \boldsymbol{I}_p^{v_j} \subseteq \bigcap_{j=1}^{n_p-f_p} \boldsymbol{I}_p^{v_j} \subseteq \boldsymbol{M}_p$$

and, for the same set $\{v_j\}$, $\boldsymbol{W}_q = \bigcap_{j=1}^{n-f'_c-f_s-f_a} \boldsymbol{I}_q^{v_j} \subseteq \boldsymbol{M}_q$. By elementary set algebra, it thus follows that $\boldsymbol{V} = \boldsymbol{W}_p \cap \boldsymbol{W}_q \in \boldsymbol{\mathcal{V}}^{n-f'_c-f_s-f_a}$ satisfies (8). Finally, $|\boldsymbol{M}_p \cap \boldsymbol{M}_q| \geq v^{n-f'_c-f_s-f_a}$ is a simple consequence of the definition of v^h as the maximum length of $\boldsymbol{V} \in \boldsymbol{\mathcal{V}}_{pq}^h$. This completes the proof of item (1).

For item (2), we distinguish two cases:

1. If —without loss of generality— $\boldsymbol{M}_p$ determines both left and right edge of $\boldsymbol{M}_p \cup \boldsymbol{M}_q$, Lemma 2 applies with $n := n_p$, $f := f_p$, and $f'_u \leq f'_a$ (as well as $f_u \leq f_a$). Hence,

by item (2) of this lemma, we know that there are at least $n_p - 2f_p - f'_u$ non-faulty intervals in $\boldsymbol{\mathcal{J}}_p$ the intersection of which majorizes $\boldsymbol{M}_p$. Using (11) and the definition of f_p, straightforward algebra yields

$$\begin{aligned} n_p - 2f_p - f'_u &\geq n + \max\{-o_p, -f_c\} + \max\{0, o_p - f_c\} - 2f_s - 2f_a - f'_a \\ &\geq n + \max\{-o_p, o_p - 2f_c\} - 2f_s - 2f_a - f'_a \\ &\geq n + \max\{-f'_c - f'_s, f'_c - 2f_c\} - 2f_s - 2f_a - f'_a. \end{aligned}$$

Abbreviating $\mu = \min\{f'_c + f'_s, 2f_c - f'_c\}$, we thus obtain

$$\boldsymbol{M}_p \cup \boldsymbol{M}_q \subseteq \boldsymbol{M}_p \subseteq \bigcap_{j=1}^{n-\mu-2f_s-2f_a-f'_a} \boldsymbol{I}_p^{b_j} \subseteq \bigcap_{j=1}^{n-\mu-2f_s-2f_a-f'_a} \boldsymbol{I}_p^{b_j} \cup \boldsymbol{I}_q^{b_j} \in \boldsymbol{\mathcal{U}}_{pq}^{n-\mu-2f_s-2f_a-f'_a}. \tag{13}$$

Note that any corresponding $\boldsymbol{I}_q^{b_i}$ must also be present in $\boldsymbol{\mathcal{J}}_q$, since our choice of f_p ensures that we got rid of any interval involved in a faulty pair. This eventually confirms (9) in this case.

2. If w.l.o.g. the left resp. right edge of $\boldsymbol{M}_p \cup \boldsymbol{M}_q$ is determined by $\boldsymbol{M}_p$ resp. $\boldsymbol{M}_q$, where the left resp. right edge of $\boldsymbol{M}_p$ resp. $\boldsymbol{M}_q$ intersects with $g_{p,l}$ resp. $g_{q,r}$ intervals belonging to a non-faulty pair of intervals in $\boldsymbol{\mathcal{I}}_p$, $\boldsymbol{\mathcal{I}}_q$, we must have

$$\begin{aligned} g_{p,l} &\geq n_p - f_p - f'_a - (s_{\text{left}} - d_{p,\text{left}}) \\ &\geq n - f_s - f_a + \max\{-o_p, -f_c\} - f'_a - s_{\text{left}} + d_{p,\text{left}} \\ g_{q,r} &\geq n_q - f_q - f'_a - (s_{\text{right}} - d_{q,\text{right}}) \\ &\geq n - f_s - f_a + \max\{-o_q, -f_c\} - f'_a - s_{\text{right}} + d_{q,\text{right}}. \end{aligned}$$

Herein, $s_{\text{left}} + s_{\text{right}} = f'_s \leq f_s$ are the number of symmetrically faulty pairs of intervals lying left resp. right of t, and $d_{p,\text{left}} + d_{p,\text{right}} = d_p$, $d_{q,\text{left}} + d_{q,\text{right}} = d_q$ denote the number of omissions among them at node p resp. q; the lower bounds follow immediately from (11).

However, we only have $g = n - f'_c - f'_s - f'_a$ different non-faulty pairs of intervals. Thus, the usual pigeonhole argument reveals that

$$\begin{aligned} g_{p,l} + g_{q,r} - g &\geq 2n + \max\{-o_p, -f_c\} + \max\{-o_q, -f_c\} - 2f_s - 2f_a - 2f'_a - f'_s \\ &\qquad + d_{p,\text{left}} + d_{q,\text{right}} - n + f'_c + f'_s + f'_a \\ &\geq n + \max\{-f'_c - d_{p,\text{right}}, -f_c + d_{p,\text{left}}\} \\ &\qquad + \max\{-f'_c - d_{q,\text{left}}, -f_c + d_{q,\text{right}}\} + f'_c - 2f_s - 2f_a - f'_a \\ &\geq n + \max\{-2f'_c - f'_s, -2f_c\} + f'_c - 2f_s - 2f_a - f'_a \\ &\geq n - \min\{f'_c + f'_s, 2f_c - f'_c\} - 2f_s - 2f_a - f'_a \end{aligned}$$

of them must be the same. Abbreviating $\mu = \min\{f'_c + f'_s, 2f_c - f'_c\}$, we can conclude that there are at least $n - \mu - 2f_s - 2f_a - f'_a$ pairs of accurate intervals, say, $\boldsymbol{I}_p^{b_1} \cup \boldsymbol{I}_q^{b_1}, \ldots, \boldsymbol{I}_p^{b_{n-\mu-2f_s-2f_a-f'_a}} \cup \boldsymbol{I}_q^{b_{n-\mu-2f_s-2f_a-f'_a}}$ with $\boldsymbol{I}_p^{b_i} \in \boldsymbol{\mathcal{J}}_p$ and $\boldsymbol{I}_q^{b_i} \in \boldsymbol{\mathcal{J}}_q$ satisfying

$$\boldsymbol{M}_p \cup \boldsymbol{M}_q \subseteq \bigcap_{j=1}^{n-\mu-2f_s-2f_a-f'_a} \boldsymbol{I}_p^{b_j} \cup \boldsymbol{I}_q^{b_j} \in \boldsymbol{\mathcal{U}}_{pq}^{n-\mu-2f_s-2f_a-f'_a}, \tag{14}$$

which proves (9) for this case as well.

To complete the proof of Lemma 4, it only remains to justify $|\boldsymbol{M}_p \cup \boldsymbol{M}_q| \leq u^{n-\mu-2f_s-2f_a-f'_a}$, which is a trivial consequence of (13) and (14). □

Remarks:

1. Note carefully that Lemma 2 could also be used to deduce a precision-related result: Since $\boldsymbol{M}_p$ and $\boldsymbol{M}_q$ are both accurate and hence contain t, it follows from item (2) that $|\boldsymbol{M}_p \cup \boldsymbol{M}_q| \leq 2w^{n-2f-f_u}$. However, comparison with item (2) of Lemma 4 reveals that this result is roughly twice as large and hence insufficient for precision enhancement.

2. Our crash faults are more severe than the (systemwide consistently perceived) *benign faults* of [AK96], since it cannot be decided locally whether an omissive interval belongs to a crash fault or to an (inconsistent) receive omission. However, it is of course possible to "merge" crash and symmetric faults, in the sense that the former are counted in f'_s resp. f_s and $f'_c = f_c = 0$ (note that $n_p - f_p = n - f_s - f_a$ in this case). After all, we already accounted for symmetric faults involving empty intervals in the proof of Lemma 4.

3. Interpreting the accomplishments of Lemma 4 and the previous remark in terms of the usual fault-tolerance degree notion, it turns out that $n \geq \min\{f'_c + f'_s, 2f_c - f'_c\} + 2f_s + 2f_a + f'_a$ nodes are required to guarantee that $\boldsymbol{M}_p \cup \boldsymbol{M}_q$ remains bounded by the length of the union of at least one pair of non-faulty input intervals. Hence, as much as

$$n \geq \begin{cases} \min\{f'_c + f'_s, 2f_c - f'_c\} + 1 & \text{for } f'_c \text{ crash faults,} \\ 2f_s + 1 & \text{for } f'_s \leq f_s \text{ symmetric faults,} \\ 2f_a + f'_a + 1 \leq 3f_a + 1 & \text{for } f'_a \leq f_a \text{ asymmetric faults} \end{cases}$$

   nodes are required for tolerating faults of the given type.

4. It should be clear from the proof of Lemma 4 that the property that really pins down symmetric faults is the following one: If a symmetrically faulty interval $\boldsymbol{I}^s_q$ intersects, say, with the right edge of $\boldsymbol{M}_q$ (correctly accounted for in s_{right}), then its corresponding $\boldsymbol{I}^s_p$ must not intersect with the left edge of $\boldsymbol{M}_p$ (since it is not accounted for in s_{left}). This is the reason why $\boldsymbol{I}^s_p \neq 0$ being faulty and $\boldsymbol{I}^s_q \neq 0$ being non-faulty must be counted as an asymmetric fault in item (2) of Definition 5.

The following lemma shows that the results of Lemma 4 remain valid if a more severe fault comes out as a less severe one, and shows what happens if certain fault assumptions are violated. Note that crash faults are counted as symmetric ones here.

Lemma 5 (Precision and Graceful Degradation $\mathcal{M}$) *Let $\boldsymbol{\mathcal{I}}_p = \{\boldsymbol{I}^1_p, \ldots, \boldsymbol{I}^n_p\}$ and $\boldsymbol{\mathcal{I}}_q = \{\boldsymbol{I}^1_q, \ldots, \boldsymbol{I}^n_q\}$ be two ordered sets of $n > f_s + f_a$, $f_s, f_a \geq 0$, compatible (or empty) accuracy intervals representing t, where $f'_s \leq f_s$ resp. $f'_a \leq f_a$ of the n pairs of intervals $\{\boldsymbol{I}^i_p, \boldsymbol{I}^i_q\}$ exhibit symmetric (or weaker) resp. asymmetric (or weaker) faults, and the remaining ones are non-faulty. As in Lemma 4, define u^h resp. v^h to be the length of the largest intersection of $h \geq 1$ unions ($\in \boldsymbol{\mathcal{U}}^h_{pq}$) resp. intersections ($\in \boldsymbol{\mathcal{V}}^h_{pq}$) of pairs of non-faulty intervals.*

Let $\boldsymbol{\mathcal{J}}_p = \{\boldsymbol{J}_1, \ldots, \boldsymbol{J}_{n_p}\}$ be the set of $n_p = n - o_p$ non-empty intervals obtained from $\boldsymbol{\mathcal{I}}_p$ by discarding any of the o_p empty intervals caused by omissions. Using the upper bound

$f_p = f_s + f_a - o_p$ on the number of intervals in $\boldsymbol{J}_p$ that (still) may be faulty in presence of o_p omissions, define

$$\boldsymbol{M}_p = \boldsymbol{\mathcal{M}}_{n_p}^{n_p - f_p}(\boldsymbol{\mathcal{J}}_p) = \boldsymbol{\mathcal{M}}_{n_p}^{n - f_s - f_a}(\boldsymbol{\mathcal{J}}_p) \quad \textit{and analogously} \quad \boldsymbol{M}_q = \boldsymbol{\mathcal{M}}_{n_q}^{n - f_s - f_a}(\boldsymbol{\mathcal{J}}_q).$$

Then,

(1) both $\boldsymbol{M}_p$ and $\boldsymbol{M}_q$ are accurate and

$$\boldsymbol{M}_p \cap \boldsymbol{M}_q \supseteq \bigcap_{j=1}^{n - f_s - f_a} \boldsymbol{I}_p^{v_j} \cap \boldsymbol{I}_q^{v_j} = \boldsymbol{V} \tag{15}$$

for any possible subset $\boldsymbol{V} \in \boldsymbol{\mathcal{V}}_{pq}^{n - f_s - f_a}$, so that $|\boldsymbol{M}_p \cap \boldsymbol{M}_q| \geq v^{n - f_s - f_a}$ (distributed minimal intersection property),

(2) there are at least $n - 2f_s - 2f_a - f'_a \geq n - 2f_s - 3f_a$ pairs of non-faulty intervals $\{\boldsymbol{I}_p^{u_k}, \boldsymbol{I}_q^{u_k}\}$ with $\boldsymbol{I}_p^{u_k} \in \boldsymbol{\mathcal{J}}_p$ and $\boldsymbol{I}_q^{u_k} \in \boldsymbol{\mathcal{J}}_q$ such that

$$\boldsymbol{M}_p \cup \boldsymbol{M}_q \subseteq \bigcap_{k=1}^{n - 2f_s - 2f_a - f'_a} \boldsymbol{I}_p^{u_k} \cup \boldsymbol{I}_q^{u_k} \subseteq \bigcap_{k=1}^{n - 2f_s - 3f_a} \boldsymbol{I}_p^{u'_k} \cup \boldsymbol{I}_q^{u'_k}, \tag{16}$$

where $\{u'_k\}_{1 \leq k \leq n - 2f_s - 3f_a}$ is obtained from $\{u_k\}_{1 \leq k \leq n - 2f_s - 2f_a - f'_a}$ by discarding $f_a - f'_a$ arbitrary elements. Hence, $|\boldsymbol{M}_p \cup \boldsymbol{M}_q| \leq u^{n - 2f_s - 2f_a - f'_a} \leq u^{n - 2f_s - 3f_a}$.

(3) Assume that the fault model is violated in the sense that $f' = f'_s + f'_a > f_s + f_a$ but still $n \geq 2f' + f'_u + 1$, where $f'_u \leq f'_a$ denotes the number of pairs of intervals that involve unbounded accuracy faults. If $\boldsymbol{M}_p$ and $\boldsymbol{M}_q$ exist, i.e., sufficiently many intersecting input intervals exist to compute $\boldsymbol{\mathcal{M}}$, then there are $n - 2f' - f'_u$ non-faulty intervals $\boldsymbol{I}_p^{p_1}, \ldots, \boldsymbol{I}_p^{p_{n - 2f' - f'_u}}$ in $\boldsymbol{\mathcal{J}}_p$ and $n - 2f' - f'_u$ non-faulty intervals $\boldsymbol{I}_q^{q_1}, \ldots, \boldsymbol{I}_q^{q_{n - 2f' - f'_u}}$ in $\boldsymbol{\mathcal{J}}_q$ such that

$$\boldsymbol{M}_p \cup \boldsymbol{M}_q \subseteq \left(\bigcap_{j=1}^{n - 2f' - f'_u} \boldsymbol{I}_p^{p_j} \right) \cup \left(\bigcap_{j=1}^{n - 2f' - f'_u} \boldsymbol{I}_q^{q_j} \right). \tag{17}$$

Hence, $|\boldsymbol{M}_p \cup \boldsymbol{M}_q| \leq w_p^{n - 2f' - f'_u} + w_q^{n - 2f' - f'_u}$, where w_p^h resp. w_q^h denote the length of the largest intersection of h accurate intervals in $\boldsymbol{\mathcal{I}}_p$ resp. $\boldsymbol{\mathcal{I}}_q$.

Nevertheless, $\boldsymbol{M}_p$ and $\boldsymbol{M}_q$ are not necessarily accurate and possibly not even consistent; accurateness is guaranteed, however, if $f' \leq f_s + f_a$ but all f' faults are asymmetric ones.

Proof: Since crash faults are now considered as symmetric ones and hence accounted for in f'_s and f_s, see Remark 2 on Lemma 4, items (1) and (2) follow directly from adopting the results of Lemma 4 to $f'_c = f_c = 0$. Note that $n_p - f_p = n - f_s - f_a$ here. To confirm the assertions for asymmetric faults appearing as weaker ones, just consider the expressions supplied by Lemma 4 when temporarily setting $f_a := f_a - 1$ and $f_s := f_s + 1$.

To show item (3), we first note that we only have to consider the case where $n_p = n - o_p \geq n - f_s - f_a$, since otherwise there would have been too many omissions to compute $\boldsymbol{M}_p$. Thus,

$$\boldsymbol{M}_p = \boldsymbol{\mathcal{M}}_{n_p}^{n - f_s - f_a}(\boldsymbol{\mathcal{J}}_p) \subseteq \boldsymbol{\mathcal{M}}_{n_p}^{n - f'}(\boldsymbol{\mathcal{J}}_p) \tag{18}$$

according to item (1) of Lemma 3. Lemma 2 is now applicable to the right-hand side of (18) and it follows by its item (2) that

$$\boldsymbol{M}_p \subseteq \bigcap_{j=1}^{n-2f'-f'_u} \boldsymbol{J}_p^{p_j}.$$

An analoguous result holds for $\boldsymbol{M}_q$. Of course, the majorizing intersections for $\boldsymbol{M}_p$ and $\boldsymbol{M}_q$ involve non-faulty intervals only, hence are accurate and thus consistent. This justifies (17) and thus the condition on $|\boldsymbol{M}_p \cup \boldsymbol{M}_q|$ given in the lemma. Note carefully, however, that this does not imply that $\boldsymbol{M}_p$ and $\boldsymbol{M}_q$ are accurate or even just consistent! On the other hand, if $f' \leq f_s + f_a$, it follows from item (1) of Lemma 2 applied to the left-hand side of (18) that $\boldsymbol{M}_p$ (and analogously $\boldsymbol{M}_q$) is accurate. □

Remarks:

1. It follows from item (3) of the above lemma that there are two possibilities in case of a violation of the failure assumptions: Either a node recognizes this fact because there are not enough accurate intervals to compute $\boldsymbol{\mathcal{M}}$, or the computed interval is not "too wrong". Obviously, this is some form of *graceful degradation* of the algorithm's performance.

2. Evidently, the worst situation with respect to the number of faults where one can hope to get a meaningful result is $n \geq 2f' + 1$. Item (3) of Lemma 5 can be used to deduce a result for this case as well: Setting $f'_u = 0$ and declaring any interval with an unbounded accuracy fault as being non-faulty, we get from (17) that $\boldsymbol{M}_p \cup \boldsymbol{M}_q$ lies in the union of the intersection of $n - 2f'$ "non-faulty" intervals in $\boldsymbol{\mathcal{J}}_p$ resp. $\boldsymbol{\mathcal{J}}_q$.

4 Orthogonal Accuracy Convergence Function

Whereas $\boldsymbol{\mathcal{M}}$ is optimal with respect to the length of the resulting accuracy interval, it can nevertheless exhibit large discontinuous jumps even for minor shifts of a faulty input interval, recall Remarks 2 and 3 on Lemma 3. For that reason, it has already been argued in [Lam87] that a "naive" clock synchronization algorithm, which places the reference point (= clock value) to the centerpoint of $\boldsymbol{\mathcal{M}}$'s result, is incapable of guaranteeing small precision if the input intervals are comparatively large. Our example run in Figure 5 justifies this claim: Consider a system of four nodes A, B, C, D with interval clocks having the following characteristics:

- A's interval clock is deteriorated by ± 2 units during the resynchronization period P, but actually runs at perfect rate,
- B's interval clock is deteriorated by ± 1 unit and is actually one unit ahead of real-time after one period P,
- C's interval clock is deteriorated by ± 1 unit and is actually one unit behind of real-time after one period P,
- D's clock exhibits byzantine faulty behaviour in the sense that D's accuracy interval received by node p mimics p's current interval clock.

Assuming fault-free, zero-delay communication and initially perfectly synchronized interval clocks $\boldsymbol{C}_p(0) = [0 \pm \mathbf{1}]$ for $p = A, B, C$, we obtain the scenario depicted in Figure 5.

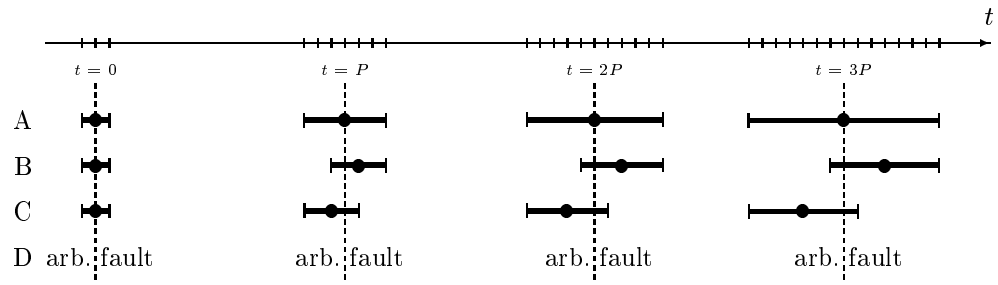

Figure 5: *Example showing the lacking precision enhancement property of Marzullo's function $\mathcal{M}$ with centerpoint setting. Since node D mimics $\boldsymbol{C}_p(t)$ when received at any node p, the interval computed by $\mathcal{M}_4^3$ at $t = P, 2P, \ldots$ reconfirms the current $\boldsymbol{C}_p(t)$*

Each node receives the interval clocks $\boldsymbol{C}_p$ of the non-faulty nodes A, B, C exactly as shown at $t = kP$, $k \geq 1$. One observes that the interval $\boldsymbol{M}_p$ obtained by applying $\mathcal{M}_4^3$ to the $\boldsymbol{I}_p^q$'s received by node p is always exactly $\boldsymbol{C}_p$ — no precision enhancement will ever take place here; the reference points of $\boldsymbol{C}_B$ and $\boldsymbol{C}_C$ will drift apart perpetually.

Our algorithm will fix this problem by applying $\mathcal{M}_4^3$ to the associated $\boldsymbol{\pi}^H$-precision intervals $\hat{\boldsymbol{I}}_p^q$ instead of $\boldsymbol{I}_p^q$, recall our exposition in Subsection 2.2. In the example of Figure 5, $\boldsymbol{\pi}^H$ must satisfy $|\boldsymbol{\pi}^H| = 2 + 4$ units because of the initial precision (= 2 units) plus twice the maximum deterioration during P (= $2 \cdot 2$ units). The intervals fed into $\mathcal{M}_4^3$ are hence trimmed to length π^H, and considering the resulting intervals $\hat{\boldsymbol{M}}_p$ and $\hat{\boldsymbol{M}}_q$ at two different nodes p and q, one finds that the reference points cannot be further apart than $\pi^H/2$ (centerpoint setting assumed), since $\hat{\boldsymbol{M}}_p \cup \hat{\boldsymbol{M}}_q$ has length at most π^H by item (2) of Lemma 4.

Since the above considerations are only meaningful for maintaining precision, i.e., setting the reference point of $\boldsymbol{C}_p$, the optimality of $\mathcal{M}$ nevertheless recommends its use for determining left and right edge of $\boldsymbol{C}_p$. However, this requires some care since the reference point computed via the associated precision intervals might lie outside of the accuracy interval. This is demonstrated by the following example: Reconsider our system of four nodes A, B, C, D, now with the following characteristics:

- A's interval clock is deteriorated by ± 1 unit and is actually one unit ahead of real-time after one period P,
- B's, C's interval clocks are deteriorated by ± 1 unit and are actually one unit behind real-time after period P,
- D's clock exhibits symmetric faults.

Assuming fault-free, zero-delay communication and initially synchonized clocks satisfying $\boldsymbol{C}_p(t_0) = \hat{\boldsymbol{C}}_p(t_0)$ for $p = A, B, C$ and $\boldsymbol{\pi}_0 = [-1, 1]$, so that $\boldsymbol{\pi}^H = [-2, 2]$ (since maximum deterioration during P is ± 1), consider the evolution of accuracy intervals during two rounds depicted in Figure 6.

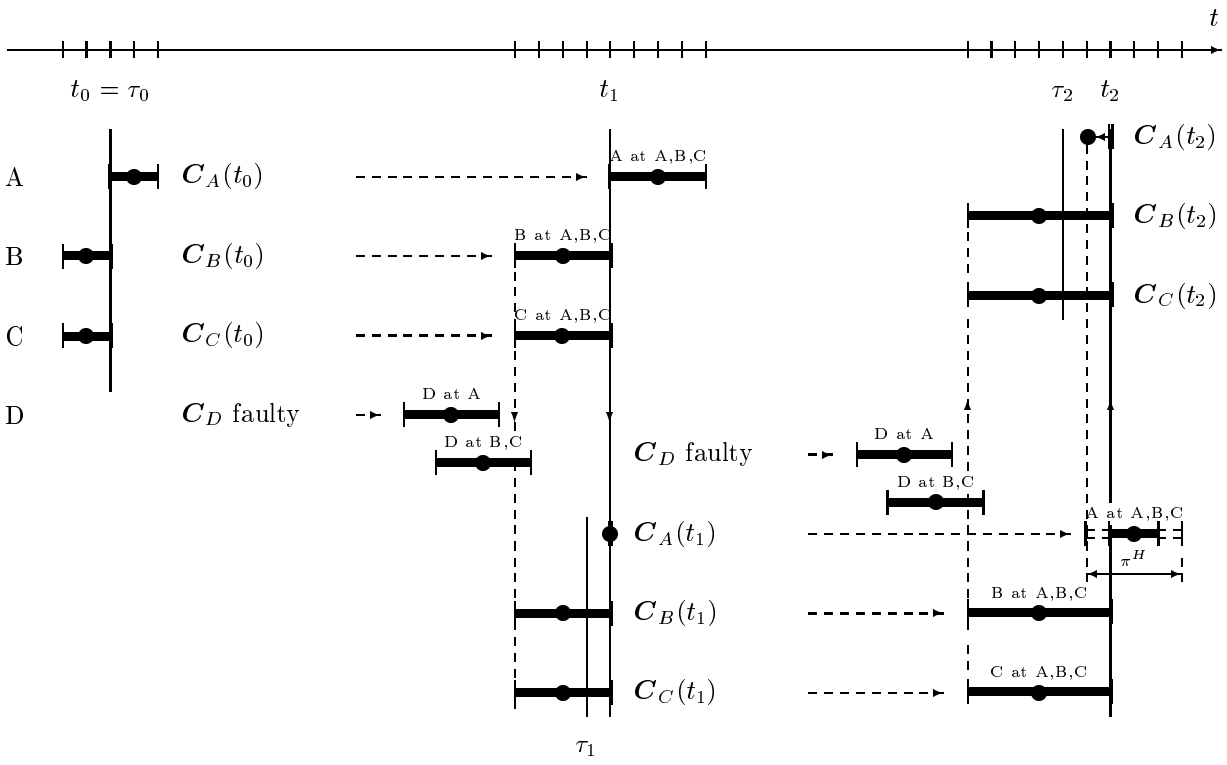

Figure 6: *Example of a reference point setting outside the accuracy interval at t_2. The accuracy interval $\boldsymbol{C}_A(t_2)$ computed by applying $\boldsymbol{\mathcal{M}}_4^3$ to the received intervals does not contain the reference point computed as the centerpoint of $\boldsymbol{\mathcal{M}}_4^3$ applied to the associated $\boldsymbol{\pi}^H$-precision intervals.*

At t_1, applying $\boldsymbol{\mathcal{M}}_4^3$ to the received accuracy intervals resp. the associated $\boldsymbol{\pi}^H$-precision intervals (which satisfy $\boldsymbol{C}_q^p(t_1) = \hat{\boldsymbol{C}}_q^p(t_1)$ for $p, q \in \{A, B, C\}$ here) yields $\boldsymbol{C}_A(t_1) = [t_1 \pm \mathbf{0}]$ and $\boldsymbol{C}_B(t_1) = \boldsymbol{C}_C(t_1) = [t_1 - 2 \pm \mathbf{2}]$ at the respective nodes; recall that the reference point of $\boldsymbol{C}_q$ is computed as the centerpoint of $\boldsymbol{\mathcal{M}}$ applied to the $\hat{\boldsymbol{C}}_q^p$'s. In order to ensure $\boldsymbol{\pi}_0$-correctness of clock A and B, internal global time τ_1 must be set to $t_1 - 1$ to lie in the intersection of the renewed $\boldsymbol{\pi}_0$-precision intervals. Although τ_1 does not lie in $\boldsymbol{C}_A(t_1)$, this situation is still feasible due to the fact that we decoupled precision and accuracy in the definition of $\boldsymbol{\pi}$-precision intervals. However, the problem shows up when setting the reference point of clock A at t_2: Applying $\boldsymbol{\mathcal{M}}$ to the received accuracy intervals at A provides $[t_2 \pm \mathbf{0}]$, but the reference point evaluates to $t_2 - 1$, which lies outside. If we ignored this, that is, if we set the reference point to t_2, precision would be violated. Therefore, $[t_2 \pm \mathbf{0}]$ must be enlarged to the left by 1 to include the reference point. Note also that internal global time τ_2 must be set to $t_2 - 2$ here.

Viewed from a different angle, this problem can be seen as a consequence of the fact that internal global time may drift away from real-time. In Figure 6, it is node D's faulty behavior that slows down the overall progress of internal global time relative to real-time. For that reason, we eventually decoupled accuracy and precision in Definition 2, viewing them as *orthogonal* issues. Note that orthogonality actually opens up the possibility of employing virtually any internal synchronization algorithm for maintaining precision, and to enlarge accuracy as needed, see Remark 3 on Theorem 2.

We should add that a similar approach was taken for the adaptive intersection algorithm [Mil95] of NTP, which is also based upon $\boldsymbol{\mathcal{M}}$. Besides trying to dynamically figure out a reasonable f (thereby sacrificing guaranteed accurateness), it extends the resulting interval appropriately to include the reference points of all apparently non-faulty input intervals. This way, it is guaranteed that the final reference point (computed according to maximum likelihood principles) lies within the final accuracy interval.

In order to formalize our orthogonal accuracy convergence function, we first have to introduce a discrete, asymmetric reference point setting operation $\boldsymbol{\pi}$-center$_{G_S}$. It is a generalization of centerpoint setting, which partitions an interval according to the proportion of $\pi^- : \pi^+$ while accounting for the fact that the CPU used for computing the reference point has integer arithmetic only: According to Definition 9 in Appendix B, we require all quantities manipulated by our clock synchronization algorithm to be integer multiples of the *clock-setting granularity* $G_S > 0$. Therefore, an integer division (rather than an exact one) is employed in $\boldsymbol{\pi}$-center$_{G_S}$, so that the analysis must deal with the truncation error.

Unfortunately, we cannot simply use exact (= non-discrete) reference point setting plus a remainder term $\mathcal{O}(G_S)$ in our analysis. Since we aiming at hardware-assisted clock synchronization with worst case accuracy and precision in the μs-range and below, see [SKM+99], this simplification would spoil the very accurate generic analysis of [SS97b]: Although G_S is smaller than clock granularity G for most adjustable clock implementations, it is nevertheless much larger than the $\mathcal{O}(.)$-terms already present in the results of [SS97b], see Remark 1 on Theorem 3. Therefore, we have to take the trouble of tracking the truncation errors explicitly.

Definition 7 (Discrete Reference Point Setting) *Let an interval $\boldsymbol{I} = [a, b]$ with a, b being integer multiples of $G_S > 0$ and some arbitrary $\boldsymbol{\pi} = [-\pi^-, \pi^+]$ satisfying $\pi = \pi^- + \pi^+ > 0$ be given. With $\lfloor x \rfloor_{G_S}$ denoting truncation of x to the next integer multiple of G_S being $\leq x$, and $\lceil x \rceil_{G_S}$ denoting rounding up x to the next integer multiple of G_S being $\geq x$, we define*

$$\boldsymbol{\pi}\text{-}center_{G_S}(\boldsymbol{I}) = \Big\lfloor \frac{\pi^- b + \pi^+ a}{\pi} \Big\rfloor_{G_S}. \tag{19}$$

A few technical lemmas dealing with the properties of $\boldsymbol{\pi}$-center$_{G_S}$ are provided in Appendix A.

Now we are ready for the formal definition of the orthogonal accuracy convergence function $\boldsymbol{\mathcal{OA}}$. Basically, the result of $\boldsymbol{\mathcal{OA}}$ is the interval provided by $\boldsymbol{\mathcal{M}}$ applied to the accuracy intervals of the input set $\boldsymbol{\mathcal{J}}$, possibly extended appropriately to include the reference point. The latter is computed independently ("orthogonally") as the $\boldsymbol{\pi}^H$-center$_{G_S}$ of the interval obtained by applying $\boldsymbol{\mathcal{M}}$ to the associated $\boldsymbol{\pi}^H$-precision intervals of the input set $\hat{\boldsymbol{\mathcal{J}}}$.

Definition 8 (Convergence Function $\boldsymbol{\mathcal{OA}}$) *Let $\boldsymbol{\mathcal{I}}$ be a set of n compatible accuracy intervals and $\boldsymbol{\mathcal{J}} \subseteq \boldsymbol{\mathcal{I}}$ with $|\boldsymbol{\mathcal{J}}| = n' \leq n$ be the set of non-empty intervals among them. With $\hat{\boldsymbol{\mathcal{J}}}$ denoting the set of associated $\boldsymbol{\pi}^H$-precision intervals for some given $\boldsymbol{\pi}^H$ with π^{H-}, π^{H+} being integer multiples of G_S, and f denoting a given fault-tolerance parameter, the* orthogonal accuracy convergence function $\boldsymbol{\mathcal{OA}}^{\boldsymbol{\pi}^H}_{n-f}(\boldsymbol{\mathcal{J}})$ *(abbreviated by $\boldsymbol{\mathcal{OA}}$) is defined by*

$$ref(\boldsymbol{\mathcal{OA}}(\boldsymbol{\mathcal{J}})) = \boldsymbol{\pi}^H\text{-}center_{G_S}(\boldsymbol{\mathcal{M}}^{n-f}_{n'}(\hat{\boldsymbol{\mathcal{J}}})) \qquad and \tag{20}$$

$$\boldsymbol{\mathcal{OA}}(\boldsymbol{\mathcal{J}}) = \boldsymbol{\mathcal{M}}^{n-f}_{n'}(\boldsymbol{\mathcal{J}}) \cup ref(\boldsymbol{\mathcal{OA}}(\boldsymbol{\mathcal{J}})). \tag{21}$$

In order to analyze the performance of the interval-based clock synchronization algorithm of Definition 9 employing $\mathcal{OA}$, it is sufficient to evaluate the characteristic functions of $\mathcal{OA}$ according to Definition 11. Plugging them into the generic results of [SS97b], as done in Section 5, the final expressions for precision, accuracy, etc. follow immediately. To improve readability, we provide $\mathcal{OA}$'s characteristic functions via two theorems: All precision-related results can be found in Theorem 1, whereas the more complicated derivations for accuracy-related quantities are covered by Theorem 2.

The first Theorem 1 states how $\mathcal{OA}$ affects precision. It determines, for any pair of nodes p and q, (1) how the application of $\mathcal{OA}$ affects precision in the current round, and (2) what precision is obtained at the beginning of the next round. As an input, our theorem takes [1] a bound $\boldsymbol{\pi}^H$ on the precision of all non-faulty input intervals, and [2] the maximum "difference" $\boldsymbol{\pi}_I$ of the intervals received from a single non-faulty sender s at node p resp. q. We particularly emphasize the quite simple proof of Theorem 1, which can be attributed to the power of our generic analysis based upon internal global time.

Theorem 1 (Precision $\mathcal{OA}$) *Let $\mathcal{I}_p = \{\boldsymbol{I}_p^1, \ldots, \boldsymbol{I}_p^n\}$ resp. $\mathcal{I}_q = \{\boldsymbol{I}_q^1, \ldots, \boldsymbol{I}_q^n\}$ be two ordered sets of n compatible accuracy intervals (all representing the same real-time t) obtained at nodes p resp. q at the end of a round, which are in accordance with the fault model of Assumption 1. Moreover, let the subsets of non-empty accuracy intervals among them be $\mathcal{J}_p \subseteq \mathcal{I}_p$, $|\mathcal{J}_p| = n_p \leq n$ and $\mathcal{J}_q \subseteq \mathcal{I}_q$, $|\mathcal{J}_q| = n_q \leq n$ and assume that*

[1] any non-faulty $\boldsymbol{I}_p^i \in \mathcal{I}_p$ as well as any non-faulty $\boldsymbol{I}_q^i \in \mathcal{I}_q$ is $\boldsymbol{\pi}^H$-correct for some given $\boldsymbol{\pi}^H$ with π^{H+}, π^{H-} being integer multiples of G_S,

[2] any pair of non-faulty intervals $\{\boldsymbol{I}_p^i, \boldsymbol{I}_q^i\}$ is $\boldsymbol{\pi}_I$-precise for some given $\boldsymbol{\pi}_I \subseteq \boldsymbol{\pi}^H$, where π_I is an integer multiple of G_S with $\zeta = (\pi^H - \pi_I)\min\{\pi^{H-}, \pi^{H+}\}/\pi^H \geq G_S$.

The convergence function $\mathcal{OA}_{n-f_s-f_a}^{\boldsymbol{\pi}^H}$ applied to $\mathcal{J}_p$ resp. $\mathcal{J}_q$ at node p resp. q provides intervals $\boldsymbol{R}_p = \mathcal{OA}(\mathcal{J}_p) = [T_p' \pm \boldsymbol{\alpha}_p']$ resp. $\boldsymbol{R}_q = \mathcal{OA}(\mathcal{J}_q) = [T_q' \pm \boldsymbol{\alpha}_q']$ with reference points being integer multiples of G_S. Its precision-related characteristic functions, which are monotonic with respect to any interval argument as long as π^{H-}/π^H remain invariant, are as follows:

(1) The precision preservation function $\boldsymbol{\Phi}(.)$*, which ensures that $\boldsymbol{R}_p$ and $\boldsymbol{R}_q$ are $\boldsymbol{\Phi}(\boldsymbol{\pi}^H)$-correct, is*

$$\boldsymbol{\Phi}(\boldsymbol{\pi}^H) = \boldsymbol{\pi}^H. \tag{22}$$

(2) The precision enhancement function $\Pi(.)$*, which ensures that the set $\{\boldsymbol{R}_p, \boldsymbol{R}_q\}$ is $\boldsymbol{\pi}_0$-precise with $\pi_0 = \Pi(\boldsymbol{\pi}^H, \boldsymbol{\pi}_I) < \pi^H$, evaluates to*

$$\Pi(\boldsymbol{\pi}^H, \boldsymbol{\pi}_I) = \max\Big\{\Big\lceil \pi^{H+} + \frac{\pi^{H-}}{\pi^H}\pi_I \Big\rceil_{G_S}, \Big\lceil \pi^{H-} + \frac{\pi^{H+}}{\pi^H}\pi_I \Big\rceil_{G_S}\Big\}. \tag{23}$$

Proof: From Definition 8 of $\mathcal{OA}$, we know that the reference point of $\boldsymbol{R}_p = \mathcal{OA}(\mathcal{J}_p)$ resp. $\boldsymbol{R}_q = \mathcal{OA}(\mathcal{J}_q)$ is determined by applying $\mathcal{M}$ to the $\boldsymbol{\pi}^H$-precision intervals associated with $\boldsymbol{I}_p^i \in \mathcal{J}_p$ resp. $\boldsymbol{I}_q^i \in \mathcal{J}_q$, resulting in

$$\tilde{\boldsymbol{R}}_p = \boldsymbol{\mathcal{M}}_{n_p}^{n-f_s-f_a}(\hat{\boldsymbol{\mathcal{J}}}_p) \qquad \text{and} \qquad \tilde{\boldsymbol{R}}_q = \boldsymbol{\mathcal{M}}_{n_q}^{n-f_s-f_a}(\hat{\boldsymbol{\mathcal{J}}}_q).$$

Since any non-faulty $\boldsymbol{I}_p^i$ is $\boldsymbol{\pi}^H$-correct according to precondition [1], it is guaranteed that $\hat{\boldsymbol{I}}_p^i$ contains internal global time τ, and that $|\hat{\boldsymbol{I}}_p^i| \leq \pi^H$, so that any intersection of such intervals has these properties as well. Lemma 2 applies with $n := n_p$ and $f := f_s + f_a - (n - n_p)$, hence it follows that $\tau \in \tilde{\boldsymbol{R}}_p$ by its item (1) and $|\tilde{\boldsymbol{R}}_p| \leq \pi^H$ by its item (2) in conjunction with

$$n_p - 2f - f'_u \geq n - 2f_s - 3f_a + n - n_p \geq n - 2f_s - 3f_a \geq 1, \tag{24}$$

recall Assumption 1. Since $\boldsymbol{\mathcal{OA}}$ sets the reference point T'_p of $\tilde{\boldsymbol{R}}_p = [T'_p \pm \boldsymbol{\pi}_p]$ to $\boldsymbol{\pi}^H$-center$_{G_S}(\tilde{\boldsymbol{R}}_p)$, applying Lemma 6 provided in Appendix A yields

$$\pi_p^- = \Big\lfloor \frac{\pi^{H-}}{\pi^H} |\tilde{\boldsymbol{R}}_p| \Big\rfloor_{G_S} \leq \frac{\pi^{H-}}{\pi^H} |\tilde{\boldsymbol{R}}_p| \leq \pi^{H-} \quad \text{and} \quad \pi_p^+ = \Big\lceil \frac{\pi^{H+}}{\pi^H} |\tilde{\boldsymbol{R}}_p| \Big\rceil_{G_S} \leq \Big\lceil \frac{\pi^{H+}}{\pi^H} \pi^H \Big\rceil_{G_S} = \pi^{H+};$$

recall that π^{H+} was assumed to be an integer multiple of G_S. Since $\mathrm{ref}(\boldsymbol{R}_p) = \mathrm{ref}(\tilde{\boldsymbol{R}}_p)$, the asserted $\boldsymbol{\pi}^H$-correctness of $\boldsymbol{R}_p$ and hence expression (22) for $\boldsymbol{\Phi}(.)$ follows. The required monotonicity of $\boldsymbol{\Phi}(\boldsymbol{\pi}^H)$ with respect to $\boldsymbol{\pi}^H$ is immediately apparent.

Of course, exactly the same reasoning holds for $\boldsymbol{R}_q$, completing the proof of item (1).

As far as item (2) is concerned, we first recall that any pair of $\boldsymbol{\pi}^H$-correct intervals $\boldsymbol{I}_p^i \in \boldsymbol{\mathcal{J}}_p$ and $\boldsymbol{I}_q^i \in \boldsymbol{\mathcal{J}}_q$ was assumed to be $\boldsymbol{\pi}_I$-precise in precondition [2]. Hence it follows that $|\hat{\boldsymbol{I}}_p^i \cup \hat{\boldsymbol{I}}_q^i| \leq \pi^H + \pi_I$, since $|\hat{\boldsymbol{I}}_p^i|, |\hat{\boldsymbol{I}}_q^i| \leq \pi^H$ and $|\mathrm{ref}(\boldsymbol{I}_p^i) - \mathrm{ref}(\boldsymbol{I}_q^i)| \leq \pi_I$ by item (2) of Lemma 1. This implies $|\tilde{\boldsymbol{R}}_p \cup \tilde{\boldsymbol{R}}_q| \leq \pi^H + \pi_I$ due to $\tilde{\boldsymbol{R}}_p \cup \tilde{\boldsymbol{R}}_q \subseteq \bigcap_{k=1}^{n-2f_s-3f_a} \hat{\boldsymbol{I}}_p^{i_k} \cup \hat{\boldsymbol{I}}_q^{i_k}$, recall item (2) of Lemma 5 and (24). Now we can apply Lemma 7 provided in Appendix A with $\pi_p = \pi_q = \bar{\pi}_p = \bar{\pi}_q := \pi^H$ and $\pi := \pi^H + \pi_I$, which eventually yields $|\mathrm{ref}(\boldsymbol{R}_p) - \mathrm{ref}(\boldsymbol{R}_q)| = |\mathrm{ref}(\tilde{\boldsymbol{R}}_p) - \mathrm{ref}(\tilde{\boldsymbol{R}}_q)| \leq \Pi(\boldsymbol{\pi}^H, \boldsymbol{\pi}_I)$ as given by (23). This ensures $\boldsymbol{\pi}_0$-precision for any $\boldsymbol{\pi}_0$ with $|\boldsymbol{\pi}_0| = \pi_0$ as asserted. The required relation $\pi_0 < \pi^H$ is verified via

$$\pi_0 - G_S \leq \max\Big\{\pi^{H+} + \frac{\pi^{H-}}{\pi^H}\pi_I, \pi^{H-} + \frac{\pi^{H+}}{\pi^H}\pi_I\Big\} = \pi^H - \zeta \leq \pi^H - G_S,$$

which follows from (23) and the condition imposed on π_I in precondition [2] of our theorem.

Finally, the required monotonicity of $\Pi(\boldsymbol{\pi}^H, \boldsymbol{\pi}_I)$ with respect to $\boldsymbol{\pi}^H$ and $\boldsymbol{\pi}_I$ is obvious from (23) by recalling that the potentially problematic fractions $\frac{\pi^{H-}}{\pi^H}$ and $\frac{\pi^{H+}}{\pi^H}$ were assumed to be invariant. This eventually completes the proof of Theorem 1. □

Remarks:

1. The above theorem considers the most general case where $|\tilde{\boldsymbol{R}}_p|$ and $|\tilde{\boldsymbol{R}}_q|$ may even be zero. A smaller $\Pi(.)$ could be derived if a (reasonably large) non-zero lower bound on $|\tilde{\boldsymbol{R}}_p|$ and $|\tilde{\boldsymbol{R}}_q|$ could be guaranteed.

2. Observe that $\Pi(\boldsymbol{\pi}^H, \boldsymbol{\pi}_I)$ given in (23) is minimized when $\boldsymbol{\pi}^H$ is symmetric, i.e., when $\pi^{H+} = \pi^{H-}$. In that case, we obtain

$$\pi_0 = \Big\lceil \frac{\pi^H + \pi_I}{2} \Big\rceil_{G_S} \leq \frac{\pi^H + \pi_I}{2} + \frac{G_S}{2}.$$

   This gives the maximum *precision enhancement* of our convergence function, cf. [Sch87]. The convergence factor is 1/2, which is the same as provided by the well-known fault-tolerant midpoint (FTM) convergence function, see [LWL88].

3. The fact that $\boldsymbol{R}_p$ is $\boldsymbol{\pi}^H$-correct easily provides the *"accuracy"* α of our convergence function in the terminology of [Sch87], which gives the maximum amount the computed clock value can differ from any non-faulty input clock value. More specifically, since any non-faulty input interval is $\boldsymbol{\pi}^H$-correct, it follows (see [SS97b, Lem. 7]) that $|\alpha| \leq \pi^H$. Hence, $\boldsymbol{\mathcal{OA}}$ provides the same "accuracy" as FTM and most other convergence functions, however, with the notable exception of the optimal algorithms [FC95b] and [Sch97a].

The following Theorem 2 shows how $\boldsymbol{\mathcal{OA}}$ affects accuracy intervals, i.e., the on-line bound on a node's maximum deviation from real-time. As an input, it takes the same precision-related quantities [1], [2] as Theorem 1, the intersection of certain precision intervals [3], and the accuracies of all non-faulty input intervals [4]. Consult the discussion prior to Definition 11 in Appendix C for details.

Theorem 2 (Accuracy $\boldsymbol{\mathcal{OA}}$) *Let $\boldsymbol{\mathcal{I}}_p = \{\boldsymbol{I}_p^1, \ldots, \boldsymbol{I}_p^n\}$ resp. $\boldsymbol{\mathcal{I}}_q = \{\boldsymbol{I}_q^1, \ldots, \boldsymbol{I}_q^n\}$ be two ordered sets of n compatible accuracy intervals (all representing the same real-time t) obtained at nodes p resp. q at the end of a round, which are in accordance with the fault model of Assumption 1. Moreover, let the subsets of non-empty accuracy intervals among them be $\boldsymbol{\mathcal{J}}_p \subseteq \boldsymbol{\mathcal{I}}_p$, $|\boldsymbol{\mathcal{J}}_p| = n_p \leq n$ and $\boldsymbol{\mathcal{J}}_q \subseteq \boldsymbol{\mathcal{I}}_q$, $|\boldsymbol{\mathcal{J}}_q| = n_q \leq n$ and assume that*

[1] any non-faulty $\boldsymbol{I}_p^i \in \boldsymbol{\mathcal{I}}_p$ as well as any non-faulty $\boldsymbol{I}_q^i \in \boldsymbol{\mathcal{I}}_q$ is $\boldsymbol{\pi}^H$-correct for some given $\boldsymbol{\pi}^H$ with π^{H+}, π^{H-} being integer multiples of G_S,

[2] any pair of non-faulty intervals $\{\boldsymbol{I}_p^i, \boldsymbol{I}_q^i\}$ is $\boldsymbol{\pi}_I$-precise for some given $\boldsymbol{\pi}_I \subseteq \boldsymbol{\pi}^H$, where π_I is an integer multiple of G_S with $\zeta = (\pi^H - \pi_I)\min\{\pi^{H-}, \pi^{H+}\}/\pi^H \geq G_S$.

[3] for any s with both $\boldsymbol{I}_p^s$ and $\boldsymbol{I}_q^s$ being non-faulty, the intersection of the associated precision intervals $\hat{\boldsymbol{I}}_p^s \cap \hat{\boldsymbol{I}}_q^s \cap \hat{\boldsymbol{I}}_p^{\min_p} \cap \hat{\boldsymbol{I}}_q^{\min_q}$ resp. $\hat{\boldsymbol{I}}_p^s \cap \hat{\boldsymbol{I}}_q^s \cap \hat{\boldsymbol{I}}_p^{\max_p} \cap \hat{\boldsymbol{I}}_q^{\max_q}$, where $\min_x$ resp. $\max_x$ represents that non-faulty node that leads to the leftmost right($\hat{\boldsymbol{I}}_x^{\min_x}$) resp. the rightmost left($\hat{\boldsymbol{I}}_x^{\max_x}$) for $x \in \{p, q\}$, has length at least $\iota_s^+ \geq 0$ resp. $\iota_s^- \geq 0$ (integer multiples of G_S),

[4] the accuracies of any non-faulty $\boldsymbol{I}_p^i = [T_p^i \pm \boldsymbol{\alpha}_p^i]$ are integer multiples of G_S satisfying $\boldsymbol{\alpha}_p^i \subseteq \boldsymbol{\beta}_p^i \in \boldsymbol{\mathcal{B}}_p$ for a given set of accuracy bounds $\boldsymbol{\mathcal{B}}_p = \{\boldsymbol{\beta}_p^1, \ldots, \boldsymbol{\beta}_p^n\}$ (and analogous for $\boldsymbol{I}_q^i$ with set of accuracy bounds $\boldsymbol{\mathcal{B}}_q$).

The convergence function $\boldsymbol{\mathcal{OA}}_{n-f_s-f_a}^{\boldsymbol{\pi}^H}$ applied to $\boldsymbol{\mathcal{J}}_p$ resp. $\boldsymbol{\mathcal{J}}_q$ at node p resp. q provides accurate intervals $\boldsymbol{R}_p = \boldsymbol{\mathcal{OA}}(\boldsymbol{\mathcal{J}}_p) = [T_p' \pm \boldsymbol{\alpha}_p']$ resp. $\boldsymbol{R}_q = \boldsymbol{\mathcal{OA}}(\boldsymbol{\mathcal{J}}_q) = [T_q' \pm \boldsymbol{\alpha}_q']$ with reference point and accuracies being integer multiples of G_S. Its accuracy-related characteristic functions, which are monotonic with respect to any interval argument as long as π^{H-}/π^H remain invariant, are as follows:

(1) The conditional accuracy preservation functions $\aleph^-(.)$, $\aleph^+(.)$, *which guarantee $\boldsymbol{\alpha}_p' \subseteq [-\aleph^-(\boldsymbol{\mathcal{B}}_p, \boldsymbol{\pi}^H, \forall s : \iota_s^-), \aleph^+(\boldsymbol{\mathcal{B}}_p, \boldsymbol{\pi}^H, \forall s : \iota_s^+)]$ read*

$$\aleph^-(\boldsymbol{\mathcal{B}}_p, \boldsymbol{\pi}^H, \forall s : \iota_s^-) = \beta_p^{x,-} + \left\lfloor \frac{\pi^{H+}}{\pi^H}(\pi^H - \iota_x^-) \right\rfloor_{G_S}, \tag{25}$$

$$\aleph^+(\boldsymbol{\mathcal{B}}_p, \boldsymbol{\pi}^H, \forall s : \iota_s^+) = \beta_p^{x,+} + \left\lceil \frac{\pi^{H-}}{\pi^H}(\pi^H - \iota_x^+) \right\rceil_{G_S}, \tag{26}$$

where x denotes the node with the $n-2f_s-3f_a$-largest accuracy bounds among $\boldsymbol{\mathcal{B}}_p$, i.e., $\beta_p^{x,-} = \max_{i:n-2f_s-3f_a}\{\beta_p^{i,-}\}$ resp. $\beta_p^{x,+} = \max_{i:n-2f_s-3f_a}\{\beta_p^{i,+}\}$ with $\max_{i:m}\mathcal{B}$ denoting the m-th largest element of the set $\mathcal{B} = \{\beta^i : 1 \leq i \leq n\}$.

(2) The conditional intersection enhancement functions $\Im_-(.)$ *resp.* $\Im_+(.)$*, which ensure that the set $\{\boldsymbol{R}_p, \boldsymbol{R}_q\}$ is $\boldsymbol{\pi}_0^{\iota_{pq}^-}$-precise resp. $\boldsymbol{\pi}_0^{\iota_{pq}^+}$-precise with $\pi_0^{\iota_{pq}^-} = \Im^-(\boldsymbol{\pi}^H, \boldsymbol{\pi}_I)$ resp. $\pi_0^{\iota_{pq}^+} = \Im^+(\boldsymbol{\pi}^H, \boldsymbol{\pi}_I)$ for worst case settings with respect to negative resp. positive accuracy, evaluate to*

$$\Im^-(\boldsymbol{\pi}^H, \boldsymbol{\pi}_I) = \Big\lceil \frac{\pi^{H-}}{\pi^H}\pi_I \Big\rceil_{G_S} \qquad \textit{resp.} \qquad \Im^+(\boldsymbol{\pi}^H, \boldsymbol{\pi}_I) = \Big\lceil \frac{\pi^{H+}}{\pi^H}\pi_I \Big\rceil_{G_S}. \tag{27}$$

Proof: Moreover, from Definition 8, it is evident that the accuracies α'^+_p, α'^-_p in $\boldsymbol{R}_p$ and α'^+_q, α'^-_q in $\boldsymbol{R}_q$ —as well as their reference points— are integer multiples of G_S. Finally, item (1) of Lemma 2 applied to (21) makes sure that $\boldsymbol{R}_p$ and $\boldsymbol{R}_q$ are accurate.

Therefore, it only remains to bound $\boldsymbol{\alpha'}_p$ (without loss of generality, since the analogous result for $\boldsymbol{R}_q$ is obtained by exchanging p and q in our theorem). For that purpose, we need an arrangement of input intervals that maximizes, say, the positive accuracy α'^+_p subject to the given lower bounds $\forall s : \iota_s^+$ on the intersection of certain non-faulty input precision intervals. Note that this worst case scenario must also cover situations where the reference point lies outside of the accuracy interval, recall our considerations at the beginning of this section.

Abbreviating the $n-2f_s-3f_a$-largest of p's accuracy bounds by $\beta = \max_{i:n-2f_s-3f_a}\{\beta_p^{i+}\}$, the worst case scenario for α'^+_p is depicted in Figure 7. Note that we will provide the detailed argument for α'^+_p only; α'^-_p is derived analogously.

First, let $\vec{\boldsymbol{I}} = [T - \pi^-, T + \alpha^+]$ be the *mixed interval* of an arbitrary accuracy interval $\boldsymbol{I} = [T \pm \boldsymbol{\alpha}]$ with its associated $\boldsymbol{\pi}$-precision interval $\hat{\boldsymbol{I}} = [T \pm \boldsymbol{\pi}]$. Mixed intervals are in fact ideally suited for attacking our problem: Since left and right edge of the result of $\boldsymbol{\mathcal{M}}$ (and hence $\boldsymbol{\mathcal{OA}}$) are computed independently of each other, and $\text{right}(\vec{\boldsymbol{I}}_p^s) = \text{right}(\boldsymbol{I}_p^s)$ resp. $\text{left}(\vec{\boldsymbol{I}}_p^s) = \text{left}(\hat{\boldsymbol{I}}_p^s)$, it follows that $\text{right}(\vec{\boldsymbol{R}}_p) = \text{right}(\boldsymbol{R}_p)$ resp. $\text{left}(\vec{\boldsymbol{R}}_p) = \text{left}(\hat{\boldsymbol{R}}_p)$ as well. Analyzing the result of $\boldsymbol{\mathcal{M}}$ in terms of mixed intervals, however, is easy since the hybrid fault model in Assumption 1 guarantees that the sets of mixed input intervals $\vec{\boldsymbol{\mathcal{I}}}_p^s$, $\vec{\boldsymbol{\mathcal{I}}}_q^s$ are in accordance with the fault model of Definition 5, which underlies the results on Marzullo's function $\boldsymbol{\mathcal{M}}$ derived in Section 3.

To see that Figure 7 provides a worst case scenario for $\alpha'^+_p = \beta'$, we first argue that $|\vec{\boldsymbol{R}}_p|$ cannot be larger than $\beta + \pi^{H-}$, since item (2) of Lemma 2 in conjunction with (24) reveals that $\vec{\boldsymbol{R}}_p$ lies within at least $n - 2f_s - 3f_a \geq 1$ non-faulty input intervals $\vec{\boldsymbol{I}}_p^{b_k} \in \boldsymbol{\mathcal{J}}_p$. Note that it is assumed in Figure 7 that $\boldsymbol{I}_p^{u_i} = \boldsymbol{I}_p^{b_i}$ has the i-th largest accuracy bound β_p^{i+} and $\vec{\boldsymbol{I}}_p^{u_{i+1}} \subseteq \vec{\boldsymbol{I}}_p^{u_i}$. This implies that $\text{left}(\vec{\boldsymbol{R}}_p)$ cannot be further left. Moreover, in an attempt to maximize β', it sould not be set further right either due to the monotonicity property (58) of $\boldsymbol{\pi}$-center$_{G_S}$.

Next, from item (2) of Lemma 5 it follows that there are at least $n - 2f_s - 3f_a$ pairs of non-faulty intervals $\{\boldsymbol{I}_p^{u_k}, \boldsymbol{I}_q^{u_k}\}$ (present in $\boldsymbol{\mathcal{J}}_p$ resp. $\boldsymbol{\mathcal{J}}_q$) such that $\vec{\boldsymbol{R}}_p \cup \vec{\boldsymbol{R}}_q \subseteq \bigcap_{k=1}^{n-2f_s-3f_a} \vec{\boldsymbol{I}}_p^{u_k} \cup \vec{\boldsymbol{I}}_q^{u_k}$. Since the reference points (and hence the left edges of the mixed intervals) of any non-faulty pair $\boldsymbol{I}_p^s$, $\boldsymbol{I}_q^s$ can be at most π_I apart according to precondition [2], $\text{left}(\vec{\boldsymbol{R}}_q)$ cannot be further left than shown in Figure 7.

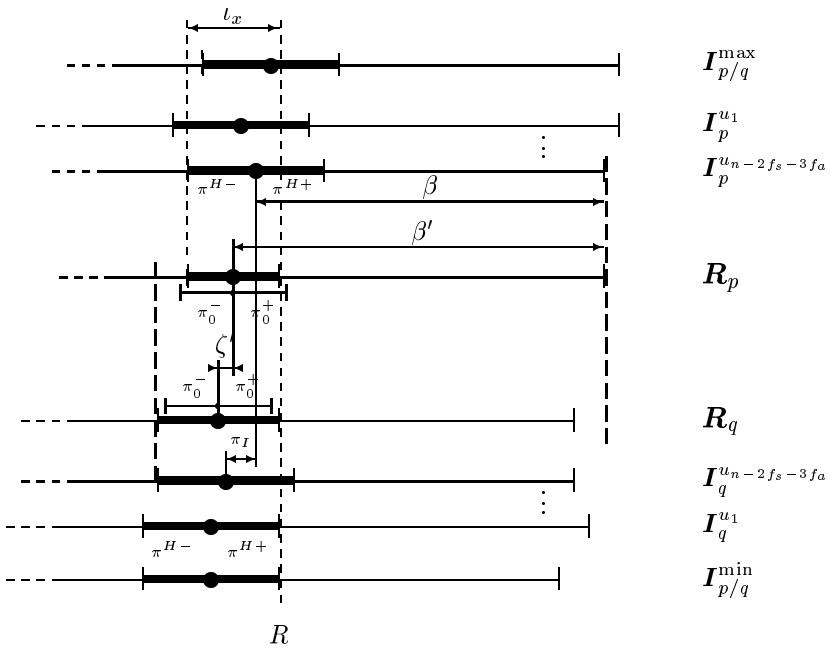

Figure 7: *Worst case scenario for $\alpha_p'^{+} = \beta'$ with resulting precision ζ'. The intersection of the leftmost non-faulty precision interval $\hat{\boldsymbol{I}}_{p/q}^{\min}$ with any non-faulty $\hat{\boldsymbol{I}}_p^x$ has length ι_x. The $n-2f_s-3f_a$ unions of input intervals $\vec{\boldsymbol{I}}_p^{u_i} \cup \vec{\boldsymbol{I}}_q^{u_i}$ contain the resulting $\vec{\boldsymbol{R}}_p \cup \vec{\boldsymbol{R}}_q$.*

Finally, as far as the worst postition of right($\tilde{\boldsymbol{R}}_p$) is concerned, we know from precondition [3] of our theorem that no non-faulty precision interval $\hat{\boldsymbol{I}}_p^s$, $\hat{\boldsymbol{I}}_q^s$ can have a right edge left of R in Figure 7. Hence, from the distributed minimal intersection property of $\boldsymbol{\mathcal{M}}$ in item (1) of Lemma 5, it follows that both $\tilde{\boldsymbol{R}}_p$ and $\tilde{\boldsymbol{R}}_q$ must have this property as well. It follows that we have to consider right($\tilde{\boldsymbol{R}}_p$) $= R$ and right($\tilde{\boldsymbol{R}}_q$) $= R$ for worst case settings with respect to β' only, since the monotonicity property (58) of $\boldsymbol{\pi}$-center$_{G_S}$ implies again that setting right($\tilde{\boldsymbol{R}}_p$) further right would provide a smaller β' only. Similarily, setting right($\tilde{\boldsymbol{R}}_q$) further right could only decrease ζ', which would enlarge ι_x' in the next round and hence provide a smaller β' then, see below.

Evaluating β' for the above situation is simple: By using (59) and $-\lceil x \rceil = \lfloor -x \rfloor$ (see [Knu73, Sec. 1.2.4, Ex. 4]), we find

$$\beta' = \beta + \pi^{H-} - \Big\lfloor \frac{\pi^{H-}}{\pi^H} \iota_x \Big\rfloor_{G_S} = \beta + \Big\lceil \frac{\pi^H}{\pi^H} \pi^{H-} - \frac{\pi^{H-}}{\pi^H} \iota_x \Big\rceil_{G_S}, \tag{28}$$

which easily yields expression (26) for $\aleph^+(.)$. The required monotonicity is immediately apparent, given that π^{H+}/π^H and π^{H-}/π^H were assumed to be invariant.

Next, we have to prove the expression for $\Im_+(.) = \text{ref}(\boldsymbol{R}_p) - \text{ref}(\boldsymbol{R}_q)$ given in item (2) of our theorem, which is bounded by ζ' in Figure 7. Noting that $\text{left}(\boldsymbol{R}_p) - \text{left}(\boldsymbol{R}_q) = \pi_I$ here, we obtain

$$\zeta' \;=\; \pi_I + \Big\lfloor \frac{\pi^{H-}}{\pi^H} \iota_x \Big\rfloor_{G_S} - \Big\lfloor \frac{\pi^{H-}}{\pi^H} (\pi_I + \iota_x) \Big\rfloor_{G_S}$$

$$
\begin{aligned}
&= \left\lfloor \frac{\pi^{H-}}{\pi^H}\iota_x \right\rfloor_{G_S} + \left\lceil \frac{\pi^{H+}}{\pi^H}\pi_I - \frac{\pi^{H-}}{\pi^H}\iota_x \right\rceil_{G_S} \\
&\leq \left\lceil \frac{\pi^{H+}}{\pi^H}\pi_I \right\rceil_{G_S}, \qquad (29)
\end{aligned}
$$

where we used the "triangle inequality" $\lceil x+y \rceil \leq \lceil x \rceil + \lceil y \rceil$ (see [Knu73, Sec. 1.2.4, Ex. 7]). This eventually confirms the expression for $\Im^+(.)$ in (27). The required monotonicity of $\Im^+(.)$ is again immediately apparent.

The analogous expressions for $\aleph^-(.)$ resp. $\Im_-(.)$ can be obtained by considering Figure 7 mirrored at the (dashed) vertical line R and exchanging π^{H+} and π^{H-} etc. Comparing (59) and (60) reveals that $\aleph^-(.)$ given by (25) follows from replacing $\lfloor . \rfloor$ by $\lceil . \rceil$ in (28). Similarily, we find

$$
\begin{aligned}
\zeta' &= \pi_I + \left\lceil \frac{\pi^{H+}}{\pi^H}\iota_x \right\rceil_{G_S} - \left\lceil \frac{\pi^{H+}}{\pi^H}(\pi_I + \iota_x) \right\rceil_{G_S} \\
&= \left\lceil \frac{\pi^{H-}}{\pi^H}\pi_I + \frac{\pi^{H+}}{\pi^H}(\pi_I + \iota_x) \right\rceil_{G_S} - \left\lceil \frac{\pi^{H+}}{\pi^H}(\pi_I + \iota_x) \right\rceil_{G_S} \\
&\leq \left\lceil \frac{\pi^{H-}}{\pi^H}\pi_I \right\rceil_{G_S} + \left\lceil \frac{\pi^{H+}}{\pi^H}(\pi_I + \iota_x) \right\rceil_{G_S} - \left\lceil \frac{\pi^{H+}}{\pi^H}(\pi_I + \iota_x) \right\rceil_{G_S} \\
&= \left\lceil \frac{\pi^{H-}}{\pi^H}\pi_I \right\rceil_{G_S},
\end{aligned}
$$

which confirms the expression for $\Im_-(.)$ in (27).

To complete the proof of our theorem, it only remains to show that the worst case scenario considered here is also *globally* valid. It is of course locally valid, in the sense that there is no scenario that provides a worse β' for the given ι_x. However, the resulting ζ' is quite small, and since ζ' determines ι_x —and hence β'— in the next round[6], the question arises whether a locally suboptimal setting could provide a worse overall accuracy. In view of (26) and (25), it is sufficient to check whether $\beta' + \zeta' = \text{right}(\boldsymbol{R}_p) - \text{ref}(\boldsymbol{R}_q)$ is maximal in the scenario of Figure 7, which is of course true by construction. This eventually completes the proof of Theorem 2. □

Remarks:

1. There is a straightforward modification of $\boldsymbol{\mathcal{OA}}$ that improves accuracy by exploiting further information from reference point setting. More specifically, any interval $\boldsymbol{I} \in \boldsymbol{\mathcal{J}}$ satisfying $\hat{\boldsymbol{I}} \cap \boldsymbol{\mathcal{M}}_{n'}^{n-f_s-f_a}(\hat{\boldsymbol{\mathcal{J}}}) = \emptyset$ is faulty since its associated $\boldsymbol{\pi}^H$-precision interval does not contain internal global time τ and may thus be discarded prior to computing the accuracy interval in (21). Item (3) of Lemma 3 implies that the accuracy of this modified version is not worse than $\boldsymbol{\mathcal{OA}}$'s, and should in fact be better in most executions, in particular, if accuracies are large.

2. Apart from clock state synchronization elaborated on in this paper, Theorems 1 and 2 can also be used immediately for clock rate synchronization. The latter is an interesting alternative to high-performance quartz clocks when targeting clock synchronization with very high precision, where decreasing any clock's drift rate to, say, $\rho \leq 10^{-7}$ is mandatory. Interestingly enough, this problem is also tractable

[6]Note that ι_x does actually not depend upon x since ζ' is uniformly valid, see Theorem 4.

by interval-based techniques: A generic analysis for clock rate synchronization was conducted in [Sch97c], which relies on the same paradigm as [SS97b]. It shows that any convergence function suitable for interval-based clock synchronization —like $\boldsymbol{\mathcal{OA}}$— can be reused in the rate setting as well.

3. It is interesting to note that statements and proof of our major Theorems 1 and 2 apply literally when the Marzullo function employed for computing the reference point (20) in $\boldsymbol{\mathcal{OA}}$'s Definition 8 is replaced by certain other interval-based functions [SS99]. For example, it is possible to recast the well-known fault-tolerant midpoint convergence function (FTM) of [LWL88] into an equivalent interval-based version FTM-I (similar to the one employed in [Lam87]) that operates on equally-sized intervals. Recall that the input precision intervals in (20) are obtained by mounting the same interval $\boldsymbol{\pi}^H$ at the accuracy intervals' reference points, resulting in identical length. Suitable pigeonhole principle proofs can then be used to show that the pivotal upper bounds of Lemma 2 and 4 for $\boldsymbol{\mathcal{M}}$ apply to FTM-I as well.

5 Orthogonal Accuracy Algorithm

In this section, we will plug in the characteristic functions of the orthogonal accuracy convergence function $\boldsymbol{\mathcal{OA}}$ into the generic expressions for precision, accuracy, etc. of Theorem 5 in Appendix C. This provides a complete characterization of the worst case performance of the orthogonal accuracy algorithm OA. In order to briefly introduce the various parameters of our system model arising in the resulting expressions, we also restate the generic algorithm's definition [SS97b, Def. 7] in Definition 9; consult [SS97b, Assum. 1–4] for further information.

The first of our major theorems describes the worst case performance of OA with respect to precision. It assumes instantaneous clock correction in Step (T) of OA, although most results carry over literally to continuous amortization, see [SS97b, Thm. 2] for details.

Theorem 3 (Precision Algorithm OA) *For the system model complying to [SS97b, Assum. 1–4] and the fault model in Assumption 1, the orthogonal accuracy algorithm with instantaneous clock correction, transmission delay compensation*

$$\begin{aligned}\Delta \;\geq\; & 2\varepsilon_{\max} + \varepsilon^{+}_{\max} + (B+3)u_{\max} + 2G + G_S + \delta_{\max}(1+\rho^{-}_{\max}) \\ & +(2P + \Lambda + \Omega + 2\Gamma_{\max} - 2\Gamma_{\min} - 2\delta_{\min})\rho_{\max} \\ & +\mathcal{O}(P\rho^2_{\max} + G\rho_{\max} + \varepsilon_{\max}\rho_{\max}),\end{aligned} \tag{30}$$

and the (symmetric!)

$$\begin{aligned}\boldsymbol{\pi}^H \;=\; & \boldsymbol{\pi}_0 + 2\boldsymbol{u}_{\max} + \overline{\boldsymbol{G}} + \boldsymbol{\varepsilon}_{\max} + (P + \Gamma_{\max} - \Gamma_{\min} - \delta_{\min})\boldsymbol{\rho}_{\max} \\ & +\mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\boldsymbol{\rho}_{\max}\end{aligned} \tag{31}$$

used in $\boldsymbol{\mathcal{OA}}^{\boldsymbol{\pi}^H}_{n-f_s-f_a}$ *requires* $\mathcal{O}(n \log n)$ *computation time to synchronize the (non-faulty) clocks of* n *nodes as follows:*

- *Initial worst-case precision (i.e., the precision at the beginning of each round of the slowest non-faulty clock)*

$$\pi_{0,\max} = \pi_0 + u_{\max} + G + (\Gamma_{\max} - \Gamma_{\min})\rho_{\max} + \mathcal{O}(P\rho^2_{\max} + G\rho_{\max} + \varepsilon_{\max}\rho_{\max}) \tag{32}$$

where $\boldsymbol{\pi}_0 = [-\pi_0^-, \pi_0^+]$ *is given by*

$$\begin{aligned}
\pi_0^- &= \frac{1}{2}\Big(\varepsilon_{\max} + Bu_{\max} + G + G_S + (\Lambda + \Omega + \Delta + \Gamma_{\max} - \delta_{\min})\rho_{\max}\Big) \\
&\quad + 2u_{\max}^+ + \varepsilon_{\max}^+ + (P + \Gamma_{\max} - \Gamma_{\min} - \delta_{\min})\rho_{\max}^+ \\
&\quad + \mathcal{O}(P\rho_{\max}^2 + G\rho_{\max} + \varepsilon_{\max}\rho_{\max}) \qquad (33) \\
\pi_0^+ &= \frac{1}{2}\Big(\varepsilon_{\max} + Bu_{\max} + G + G_S + (\Lambda + \Omega + \Delta + \Gamma_{\max} - \delta_{\min})\rho_{\max}\Big) \\
&\quad + 2u_{\max}^- + \varepsilon_{\max}^- + G + (P + \Gamma_{\max} - \Gamma_{\min} - \delta_{\min})\rho_{\max}^- \\
&\quad + \mathcal{O}(P\rho_{\max}^2 + G\rho_{\max} + \varepsilon_{\max}\rho_{\max}), \qquad (34)
\end{aligned}$$

so that

$$\begin{aligned}
\pi_0 &= 2\varepsilon_{\max} + (B+2)u_{\max} + 2G + G_S \qquad (35) \\
&\quad + (P + \Lambda + \Omega + \Delta + 2\Gamma_{\max} - \Gamma_{\min} - 2\delta_{\min})\rho_{\max} \\
&\quad + \mathcal{O}(P\rho_{\max}^2 + G\rho_{\max} + \varepsilon_{\max}\rho_{\max}). \qquad (36)
\end{aligned}$$

- *Overall worst-case precision* $\pi_{\max}$ *satisfying*

$$\begin{aligned}
\pi_{\max} &= 2\varepsilon_{\max} + (B+3)u_{\max} + 3G + G_S \\
&\quad + (2P + \Lambda + \Omega + \Delta + 2\Gamma_{\max} - \Gamma_{\min} - 2\delta_{\min})\rho_{\max} \\
&\quad + \max\Big\{\varepsilon_{\max}^+ + 2u_{\max}^+ + (2\Gamma_{\max} - 2\Gamma_{\min} - \delta_{\min})\rho_{\max}^+, \\
&\qquad \varepsilon_{\max}^- + 2u_{\max}^- + G + (2\Gamma_{\max} - 2\Gamma_{\min} - \delta_{\min})\rho_{\max}^-, 0\Big\} \\
&\quad + \mathcal{O}(P\rho_{\max}^2 + G\rho_{\max} + \varepsilon_{\max}\rho_{\max}), \qquad (37)
\end{aligned}$$

- *Any two non-faulty nodes* p, q *resynchronize within real-time* $t_p^R - t_q^R$ *satisfying*

$$\Gamma_p - \Gamma_q - \pi_P \leq t_p^R - t_q^R \leq \Gamma_p - \Gamma_q + \pi_P$$

for

$$\pi_P = \pi_0 + u_{\max} + P\rho_{\max} + \mathcal{O}(P\rho_{\max}^2 + G\rho_{\max} + \varepsilon_{\max}\rho_{\max}), \qquad (38)$$

- *Adjustments of at most* Υ *are applied to the local clock of any non-faulty node, which are bounded according to* $-\pi^- \leq \Upsilon \leq \pi^+$ *with*

$$\begin{aligned}
\pi^- &= 2\varepsilon_{\max} + (B+3)u_{\max} + 2G + G_S \\
&\quad + (2P + \Lambda + \Omega + \Delta + 2\Gamma_{\max} - \Gamma_{\min} - 2\delta_{\min})\rho_{\max} + \varepsilon_{\max}^+ + u_{\max}^+ \\
&\quad + (\Gamma_{\max} - \Gamma_{\min} - \delta_{\min})\rho_{\max}^+ + \mathcal{O}(P\rho_{\max}^2 + G\rho_{\max} + \varepsilon_{\max}\rho_{\max}), \qquad (39) \\
\pi^+ &= 2\varepsilon_{\max} + (B+3)u_{\max} + 3G + G_S \\
&\quad + (2P + \Lambda + \Omega + \Delta + 2\Gamma_{\max} - \Gamma_{\min} - 2\delta_{\min})\rho_{\max} + \varepsilon_{\max}^- + u_{\max}^- \\
&\quad + (\Gamma_{\max} - \Gamma_{\min} - \delta_{\min})\rho_{\max}^- + \mathcal{O}(P\rho_{\max}^2 + G\rho_{\max} + \varepsilon_{\max}\rho_{\max}). \qquad (40)
\end{aligned}$$

Proof: The computational complexity of the orthogonal accuracy algorithm is primarily determined by the complexity of computing the convergence function $\boldsymbol{\mathcal{OA}}$, which is $\mathcal{O}(n \log n)$ due to the two evaluations of $\boldsymbol{\mathcal{M}}$ employed in $\boldsymbol{\mathcal{OA}}$, recall Definition 8 and 6.

By item (1) of Theorem 5, $\boldsymbol{\pi}_0$ is the solution of the equation $|\boldsymbol{\pi}_0| = \Pi(\boldsymbol{\pi}^H, \boldsymbol{\pi}_I)$ involving $\mathcal{OA}$'s precision enhancement function $\Pi(.)$ given by (23). Precision enhancement is optimal if $\boldsymbol{\pi}^H$ given by (82) is a symmetric interval, recall the remarks following Theorem 1. Note that the condition imposed on π_I in precondition [2] of Theorem 1 is also amply fulfilled in this case, just compare (83) and (31). Hence, writing $\boldsymbol{\pi}^H = \boldsymbol{\pi}_0 + \boldsymbol{\pi}_1$ with

$$\boldsymbol{\pi}_1 = 2\boldsymbol{u}_{\max} + \overline{\boldsymbol{G}} + \boldsymbol{\varepsilon}_{\max} + (P + \Gamma_{\max} - \Gamma_{\min} - \delta_{\min})\boldsymbol{\rho}_{\max} + \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\boldsymbol{\rho}_{\max}, \quad (41)$$

we exploit the freedom of choosing an arbitrary reference point of $\boldsymbol{\pi}_0$ to enforce this symmetry: Setting

$$\boldsymbol{\pi}_0 = \Big[-\Big(\lceil \pi_I/2 \rceil_{G_S} + \pi_1^+\Big), \lceil \pi_I/2 \rceil_{G_S} + \pi_1^-\Big] \quad (42)$$

such that $\pi_0 = 2\lceil \pi_I/2 \rceil_{G_S} + \pi_1$ provides a symmetric interval

$$\boldsymbol{\pi}^H = \boldsymbol{\pi}_0 + \boldsymbol{\pi}_1 = \Big[-\Big(\lceil \pi_I/2 \rceil_{G_S} + \pi_1\Big), \lceil \pi_I/2 \rceil_{G_S} + \pi_1\Big] \quad (43)$$

with $\pi^H = 2\lceil \pi_I/2 \rceil_{G_S} + 2\pi_1$. Evaluating $\Pi(\boldsymbol{\pi}^H, \boldsymbol{\pi}_I)$ given by (23) yields

$$\begin{aligned} \max\Big\{\Big\lceil \pi^{H+} + \frac{\pi^{H-}}{\pi^H}\pi_I\Big\rceil_{G_S}, \Big\lceil \pi^{H-} + \frac{\pi^{H+}}{\pi^H}\pi_I\Big\rceil_{G_S}\Big\} &= \Big\lceil \frac{\pi^H + \pi_I}{2}\Big\rceil_{G_S} \\ &= \Big\lceil \lceil \pi_I/2 \rceil_{G_S} + \pi_1 + \pi_I/2\Big\rceil_{G_S} \\ &= 2\lceil \pi_I/2 \rceil_{G_S} + \pi_1 = \pi_0 \end{aligned} \quad (44)$$

as required; recall that π_1 and all other precision values are integer multiples of G_S since its constituting parameters have this property, see Definition 9. Plugging in π_1^+, π_1^- resulting from (41) and $\lceil \pi_I/2 \rceil_{G_S} \leq \pi_I/2 + G_S/2$ with $\pi_I = |\boldsymbol{\pi}_I|$ from (83) into (42) confirms the values of π_0^-, π_0^+ given in (33) and (34). Addition or, alternatively, substitution in (44) provides the value of π_0 stated in (36). Last but not least, (32) providing $\pi_{0,\max}$ is only a restatement of (78).

Next, the value π_P given in (38) is obtained by plugging in (conservative) maximum bounds for $\boldsymbol{u}_{p/q}$ and $\boldsymbol{\rho}_{p/q}$ in (86). Plugging in the swapped $\overline{\boldsymbol{\Phi}}(\boldsymbol{\pi}^H) = \overline{\boldsymbol{\pi}}^H$ according to item (1) of Theorem 1 and the definition (31) of $\boldsymbol{\pi}^H$ into (85) provides

$$\begin{aligned} \boldsymbol{\pi} &= \overline{\boldsymbol{\pi}}_0 + 2\overline{\boldsymbol{u}}_{max} + \boldsymbol{G} + \overline{\boldsymbol{\varepsilon}}_{\max} + \Big(P + \Gamma_{\max} - \Gamma_{\min} - \delta_{\min} + \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\Big)\overline{\boldsymbol{\rho}}_{\max} \\ &\quad + \boldsymbol{\pi}_0 + \boldsymbol{u}_{\max} + \Big(P + \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\Big)\boldsymbol{\rho}_{\max}, \end{aligned}$$

from where the values of π^-, π^+ given in (39) and (40) follow easily. Now it is possible to evaluate (84) in Theorem 5, which confirms the value of $\pi_{\max}$ stated in (37).

The proof of Theorem 3 is almost completed; we only have to justify the value of Δ given in (30). Plugging in the expressions for π_0 and $\pi = \mathcal{O}(\varepsilon_{\max} + G + P\rho_{\max})$ into the definition of Δ in (68), we easily obtain

$$\begin{aligned} \Delta &\geq \frac{1}{1 + \rho_{\max}^+}\Big(2\varepsilon_{\max} + (B+3)u_{\max} + 2G + G_S + \delta_{\max} + \varepsilon_{\max}^+ \\ &\qquad + (2P + \Lambda + \Omega + \Delta + 2\Gamma_{\max} - 2\Gamma_{\min} - 2\delta_{\min})\rho_{\max} \\ &\qquad + \mathcal{O}(P\rho_{\max}^2 + G\rho_{\max} + \varepsilon_{\max}\rho_{\max})\Big) \\ &= \frac{Z + \Delta\rho_{\max}}{1 + \rho_{\max}^+} \end{aligned}$$

Solving this for Δ yields

$$\begin{aligned}\Delta &\geq \frac{Z}{1-\rho^-_{\max}} = Z\Big(1+\rho^-_{\max}+\mathcal{O}((\rho^-_{\max})^2)\Big)\\ &= 2\varepsilon_{\max}+\varepsilon^+_{\max}+(B+3)u_{\max}+2G+G_S+\delta_{\max}\\ &\quad +(2P+\Lambda+\Omega+2\Gamma_{\max}-2\Gamma_{\min}-2\delta_{\min})\rho_{\max}\\ &\quad +\delta_{\max}\rho^-_{\max}+\mathcal{O}(P\rho^2_{\max}+G\rho_{\max}+\varepsilon_{\max}\rho_{\max}).\end{aligned}$$

This eventually completes the proof of Theorem 3. □

Remarks:

1. For usual settings, the $\mathcal{O}(.)$-terms in our theorem are very small and can be neglected in practice. Their only purpose is to specify the order of magnitude of the neglected terms.

2. The precision results above are valid for <u>any</u> setting of the parameters defined in the system model. For that purpose, our analysis provides the worst case behaviour of the algorithm under the worst setting of parameters, like $\boldsymbol{\rho}_p = \boldsymbol{\rho}_{\max}$ for any node p. One could derive improved worst case results for more relaxed parameterizations, although our algorithm does not benefit much from such situations unless refined worst case bounds are compiled into it — after all, $\boldsymbol{\mathcal{OA}}$ depends on $\boldsymbol{\pi}^H$!

3. With respect to worst case precision, our orthogonal accuracy algorithm is equivalent to the well-known fault tolerant midpoint algorithm (FTM) of [LWL88], cf. Remark 3 following Theorem 2. FTM has the same computational complexity $\mathcal{O}(n \log n)$ and the same worst-case precision $\pi_{\max} \approx 5\varepsilon + 4P\rho$ (in a comparable setting); our terminology relates to the one of [LWL88] by $\pi_{\max} = \gamma$, $\varepsilon_{\max} = 2\varepsilon$ and $\rho_{\max} = 2\rho$. Both algorithms require initially synchronized clocks as well.

   Our algorithm is hence slightly suboptimal with respect to worst case precision: $\pi_{\max}$ exceeds the provably necessary and tight lower bound $4\varepsilon + 4P\rho$ (see [FC95a], [FC95b]) by ε. The maximum correction $\Upsilon \approx 5\varepsilon + 4P\rho$ applied to the clock of a non-faulty node, however, exceeds the tight lower bound of $2P\rho$ considerably.

The next theorem provides algorithm OA's worst case behavior with respect to accuracy:

Theorem 4 (Accuracy Algorithm OA) *For the system model complying to [SS97b, Assum. 1–4] and the fault model in Assumption 1, the accuracies $\alpha_q^{k+1,-}$, $\alpha_q^{k+1,+}$ of a non-faulty node q's accuracy interval $\boldsymbol{A}_q^{k+1}(t_q^{k+1}) = [T_q^{k+1} \pm \boldsymbol{\alpha}_q^{k+1}]$ at the beginning of round $k+1$, $k \geq 0$, as computed by the orthogonal accuracy algorithm OA with transmission delay compensation Δ given by (30) and $\boldsymbol{\pi}^H$ given by (31), are integer multiples of G_S satisfying the following properties:*

- *The interval of accuracy satisfies $\boldsymbol{\alpha}_q^{k+1} \subseteq \boldsymbol{\beta}_q^{k+1}$ with*

$$\begin{aligned}\beta_q^{k+1,-} &= \max_{p:n-2f_s-3f_a}\Big\{\Big\{\beta_p^{k,-}+u_p^-+u_q^-+G+G_A+\varepsilon^-_{pq}+(P-\Delta-\Gamma_p)\rho_p^-\\ &\qquad +(\Gamma_q+\Delta-\delta_{pq})\rho_q^-+(\Lambda+\Omega)\max\{\rho_q^--\rho_p^-,0\}\Big\}_{p\neq q}\end{aligned}$$

$$
\begin{aligned}
&\cup\Big\{\beta_q^{k,-}+u_q^-+P\rho_q^-\Big\}\Big\}\\
&+\lfloor\pi^H/4\rfloor_{G_S}+\mathcal{O}(P\rho_{\max}^2+G\rho_{\max}+\varepsilon_{\max}\rho_{\max}) &(45)\\
\leq\ &\max_{p:n-2f_s-3f_a}\{\beta_p^{k,-}\}+u_{\max}^-+G+G_A+\varepsilon_{\max}^-+(P-\Delta-\Gamma_{\min})\rho_{\max}^-\\
&+u_q^-+(\Gamma_q+\Delta-\delta_{\min})\rho_q^-\\
&+\lfloor\pi^H/4\rfloor_{G_S}+\mathcal{O}(P\rho_{\max}^2+G\rho_{\max}+\varepsilon_{\max}\rho_{\max}), &(46)\\
\beta_q^{0,-}\ =\ &\alpha_q^{0,-}+\lceil\pi^H/4\rceil_{G_S}, &(47)\\
\beta_q^{k+1,+}\ =\ &\max_{p:n-2f_s-3f_a}\Big\{\Big\{\beta_p^{k,+}+u_p^++u_q^++G_A+\varepsilon_{pq}^++(P-\Delta-\Gamma_p)\rho_p^+\\
&\qquad+(\Gamma_q+\Delta-\delta_{pq})\rho_q^++(\Lambda+\Omega)\max\{\rho_q^+-\rho_p^+,0\}\Big\}_{p\neq q}\\
&\cup\Big\{\beta_q^{k,+}+u_q^++P\rho_q^+\Big\}\Big\}\\
&+\lceil\pi^H/4\rceil_{G_S}+\mathcal{O}(P\rho_{\max}^2+G\rho_{\max}+\varepsilon_{\max}\rho_{\max}) &(48)\\
\leq\ &\max_{p:n-2f_s-3f_a}\{\beta_p^{k,+}\}+u_{\max}^++G_A+\varepsilon_{\max}^++(P-\Delta-\Gamma_{\min})\rho_{\max}^+\\
&+u_q^++(\Gamma_q+\Delta-\delta_{\min})\rho_q^+\\
&+\lceil\pi^H/4\rceil_{G_S}+\mathcal{O}(P\rho_{\max}^2+G\rho_{\max}+\varepsilon_{\max}\rho_{\max}), &(49)\\
\beta_q^{0,+}\ =\ &\alpha_q^{0,+}+\lfloor\pi^H/4\rfloor_{G_S}, &(50)
\end{aligned}
$$

where $\max_{p:m}\mathcal{B}$ *denotes the* m*-th largest element of the set* $\mathcal{B}=\{\beta_p:1\leq p\leq n\}$, $\boldsymbol{\alpha}_q^0\subseteq\boldsymbol{\pi}_0$ *is the initial interval of accuracies, and*

$$
\begin{aligned}
\pi^H\ =\ &3\varepsilon_{\max}+(B+4)u_{\max}+3G+G_S\\
&+(2P+\Lambda+\Omega+\Delta+3\Gamma_{\max}-2\Gamma_{\min}-3\delta_{\min})\rho_{\max}\\
&+\mathcal{O}(P\rho_{\max}^2+G\rho_{\max}+\varepsilon_{\max}\rho_{\max}); &(51)
\end{aligned}
$$

note that $\lfloor\pi^H/4\rfloor_{G_S}\leq\pi^H/4$ *and* $\lceil\pi^H/4\rceil_{G_S}\leq\pi^H/4+G_S/2$.

- *The maximum deviation of node* q*'s reference point* T_q^{k+1} *from real-time* t_q^{k+1} *("traditional accuracy") at the beginning of round* $k+1$, $k\geq 0$, *reads*

$$
\begin{aligned}
T_q^{k+1}-t_q^{R,k}\ &\leq\ \pi_0^-+(k+1)\Big(2u_{\max}^-+G+\varepsilon_{\max}^-+(P+\Gamma_{\max}-\Gamma_{\min}-\delta_{\min})\rho_{\max}^-\\
&\qquad+\mathcal{O}(P\rho_{\max}^2+G\rho_{\max}+\varepsilon_{\max}\rho_{\max})\Big), &(52)\\
T_q^{k+1}-t_q^{R,k}\ &\geq\ -\pi_0^+-(k+1)\Big(2u_{\max}^++\varepsilon_{\max}^++(P+\Gamma_{\max}-\Gamma_{\min}-\delta_{\min})\rho_{\max}^+\\
&\qquad+\mathcal{O}(P\rho_{\max}^2+G\rho_{\max}+\varepsilon_{\max}\rho_{\max})\Big), &(53)
\end{aligned}
$$

where π_0^- *resp.* π_0^+ *are given by (33) resp. (34).*

- *The inverse rate* $r_q=\lim_{k\to\infty}\frac{t_q^{R,k}-t_q^0}{T_q^{k+1}-T_q^0}$, *where* $T_q^0=C_q(t_q^0)$ *is node* q*'s local time at the beginning of round* $k=0$, *of the synchronized clock at node* q *lies within*

$$
\left[1\pm\Big(\boldsymbol{\rho}_{\max}+\frac{2\boldsymbol{u}_{\max}+\overline{\boldsymbol{G}}+\boldsymbol{\varepsilon}_{\max}+(\Gamma_{\max}-\Gamma_{\min}-\delta_{\min})\boldsymbol{\rho}_{\max}}{P}+\mathcal{O}(P\rho_{\max}+G+\varepsilon_{\max})\boldsymbol{\rho}_{\max}\Big)\right]. \tag{54}
$$

Proof: We showed in the proof of Theorem 3 that $\boldsymbol{\pi}^H$ given by (31), which immediately leads to (51), is a symmetric interval, hence $\pi^{H+}/\pi^H = \pi^{H-}/\pi^H = 1/2$. Since both π^{H-} and π^{H+} are integer multiples of G_S, it follows that the error in estimating $\lceil \pi^H/4 \rceil_{G_S}$ by $\pi^H/4$ is at most $G_S/2$ as asserted.

According to item (0) of Theorem 5, the bounds $\beta_q^{k+1,-}$ resp. $\beta_q^{k+1,+}$ are just $\boldsymbol{\mathcal{OA}}$'s accuracy preservation functions $\aleph^-(.)$ resp. $\aleph^+(.)$ derived in item (1) of Theorem 2: For round $k = 0$, evaluating (77) according to the initial synchronization assumption in item (0) of Definition 9 forces us to assume a zero-length common intersection $\iota_s^{0,-} = \iota_s^{0,+} = 0$ for any node s. Plugging this into (25) resp. (26) governed by (70) and using (71) provides

$$\begin{aligned}
\aleph^-(\boldsymbol{\mathcal{B}}_q^1, \boldsymbol{\pi}^H, \forall s : 0) &= \max_{i:n-2f_s-3f_a} \{\beta_i^{0,-} + \omega_q^{i,-}\} + \pi^{H+} \\
&= \max_{i:n-2f_s-3f_a} \{\alpha_i^{0,-} + \omega_q^{i,-}\} + \lceil \pi^H/4 \rceil_{G_S} + \lfloor \pi^H/4 \rfloor_{G_S} \\
\aleph^+(\boldsymbol{\mathcal{B}}_q^1, \boldsymbol{\pi}^H, \forall s : 0) &= \max_{i:n-2f_s-3f_a} \{\beta_i^{0,+} + \omega_q^{i,+}\} + \pi^{H-} \\
&= \max_{i:n-2f_s-3f_a} \{\alpha_i^{0,+} + \omega_q^{i,+}\} + \lceil \pi^H/4 \rceil_{G_S} + \lfloor \pi^H/4 \rfloor_{G_S} \qquad (55)
\end{aligned}$$

where we used $\pi^H/2 - \lceil \pi^H/4 \rceil_{G_S} = \lfloor \pi^H/4 \rfloor_{G_S}$. For round $k \geq 1$, we have $\iota_s^k = \iota_s^{k,-} = \iota_s^{k,+} = \pi_0 - \lceil \pi_I/2 \rceil_{G_S}$ uniformly for any node s due to (75) resp. (76) in conjunction with (27). Therefore,

$$\frac{1}{2}(\pi^H - \iota_s^k) = \frac{\pi^H - \pi_0 + \lceil \pi_I/2 \rceil_{G_S}}{2} = \frac{\pi_1 + \lceil \pi_I/2 \rceil_{G_S}}{2} = \pi^H/4,$$

where we used (43). Plugging this into (25) resp. (26) and using (70) and (71) again provides

$$\begin{aligned}
\aleph^-(\boldsymbol{\mathcal{B}}_q^{k+1}, \boldsymbol{\pi}^H, \forall s : \iota_s^{k,-}) &= \max_{i:n-2f_s-3f_a} \{\beta_i^{k,-} + \omega_q^{i,-}\} + \lfloor \pi^H/4 \rfloor_{G_S} \\
\aleph^+(\boldsymbol{\mathcal{B}}_q^{k+1}, \boldsymbol{\pi}^H, \forall s : \iota_s^{k,+}) &= \max_{i:n-2f_s-3f_a} \{\beta_i^{k,+} + \omega_q^{i,+}\} + \lceil \pi^H/4 \rceil_{G_S}. \qquad (56)
\end{aligned}$$

Replacing i by p and plugging in the definition of $\boldsymbol{\omega}_q^p$ given in (71) and (72), expressions (45) and (48) follow. Uniformly bounding the terms that depend on p by their maximum values and applying Lemma 8 readily confirms (46) and (49); note that a certain technical condition [SS97b, (16)], namely, $\delta_{\min}\boldsymbol{\rho}_{\max} \subseteq \boldsymbol{\varepsilon}_{\max}$, ensures that the derived bound is valid for $p = q$ as well. To justify the initial values (47) and (50), we just note that (55) and (56) differ only by $\lfloor \pi^H/4 \rfloor_{G_S}$, which can conveniently be put into $\beta_q^{0,+}$ and $\beta_q^{0,-}$.

Turning our attention to traditional accuracy, we first plug in $\boldsymbol{\mathcal{OA}}$'s precision preservation function $\boldsymbol{\Phi}(\boldsymbol{\pi}^H) = \boldsymbol{\pi}^H$ into (88) to obtain

$$T_q^{k+1} - t_q^{R,k} \in \overline{\boldsymbol{\pi}}_0 + (k+1)(\overline{\boldsymbol{\pi}}^H - \overline{\boldsymbol{\pi}}_0); \qquad (57)$$

inserting the swapped expression for $\boldsymbol{\pi}^H - \boldsymbol{\pi}_0$ obtained from (31) easily yields (52) and (53). Finally, plugging in $\boldsymbol{\pi}^H - \boldsymbol{\pi}_0$ into (89) also justifies expression (54) for the inverse rate of the synchronized clock, completing the proof of Theorem 4. □

Remarks:

1. OA provides the same, slightly suboptimal worst case traditional accuracy and drift as the FTM algorithm of [LWL88]: The synchronized clocks drift away from real-time by at most the maximum drift $\rho_{\max}$ of the (worst) physical clock plus some smaller terms. Note that it has been proved in [ST87] that the worst case drift cannot be better than the drift of the physical clocks. Algorithms that are optimal in this respect have been provided in [ST87], [FC95b] and [Sch97a].

2. It is important to note that accuracy intervals can grow much faster than traditional accuracy, which reveals the major weakness of the orthogonal accuracy algorithm: Basically, traditional accuracy increases resp. decreases at most by the maximum enlargement of negative resp. positive accurcay $\boldsymbol{\pi}^H - \boldsymbol{\pi}_0 = \max_{p,q}\{\boldsymbol{\omega}_q^p\}$ given by (71) during each round. However, (56) reveals that the growth of both positive and negative accuracy exceeds this value by $\approx \pi^H/4$. This implies that both $\alpha_q^{k+1,-}$ and $\alpha_q^{k+1,+}$ could get larger than the sum of positive (53) and negative (52) traditional accuracy (although this cannot happen simultaneously). Intuitively, this is due to the fact that the reference point and one edge of the current accuracy interval could move into opposite directions due to faulty intervals that affect the accuracy and precision algorithm differently. Recall that it can even happen that the reference point is placed outside the originally computed accuracy interval.

3. There are several directions one can follow in order to devise convergence functions that behave better with respect to the accuracy bounds:

   - One should consider algorithms that maintain precision and accuracy not orthogonally but rather in an integrated way.
   - One could try to reduce the "power" of the precision algorithm with respect to causing internal global time to drift from real-time. In fact, OA enforces precision by quite radical corrections (suboptimal maximum clock corrections Υ, see Remark 3 following Theorem 3) that are incompatible with the progress of real-time. The *optimal precision convergence function* $\boldsymbol{\mathcal{OP}}$ analyzed in [Sch97a] shows how far one can get with this approach.
   - There might be ways of limiting the adverse effects of faulty clocks.

6 Conclusions

In this paper, we introduced and rigorously analyzed a novel convergence function-based orthogonal accuracy clock synchronization algorithm OA. Belonging to the class of interval-based algorithms, it both guarantees bounded internal synchronization precision and provides on-line bounds on the instantaneous accuracy with respect to external time as well. OA employs the orthogonal accuracy convergence function $\boldsymbol{\mathcal{OA}}$ in the generic algorithm of [SS97b], which encodes two (almost) independent algorithms for maintaining precision and accuracy based on the Marzullo function $\boldsymbol{\mathcal{M}}$.

Our comprehensive analysis, which relies upon the powerful interval-based framework established in [SS97b] and incorporates a thorough investigation of the worst case performance of $\boldsymbol{\mathcal{M}}$ as well, provides accurate expressions for all worst case performance measures, like precision, maximum clock correction, accuracy, etc. With respect to worst case precision, it turned out that OA performs equivalent to the well-known *fault-tolerant*

midpoint (FTM) algorithm of [LWL88]: Maximum precision $5\varepsilon + 4P\rho$, maximum clock correction $5\varepsilon+4P\rho$, and global drift ρ plus smaller terms are hence all slightly sub-optimal. The major weakness of OA, however, are its relatively large worst case accuracy bounds, which could even be larger than twice its instantaneous accuracy.

A particular advantage of our results over traditional ones is that they are suitable for very high-accuracy clock synchronization as well. This is primarily a consequence of a very detailed system model, which incorporates several non-standard issues like non-zero clock granularity and broadcast latencies. An expressive hybrid fault model incorporating both arbitrary and more restricted node/communication faults is also utilized; among its particular strengths is proper modeling of receive omissions. The most important fact revealed by our detailed formulas is that clock granularity (G) and, in particular, rate adjustment uncertainity (usually $u = G/m$, for some small positive integer m) of discrete rate-adjustment techniques have a considerable impact (as much as $12u + 4G$) upon achievable worst case precision and accuracy. This makes clear that any attempt to approach 1 μs worst case precision —as targeted by our SynUTC-project, see [SKM+99]— must utilize clocks with $G, u \ll 1\mu s$.

Future work in this area will primarily be devoted to an improvement of the definitely sub-optimal accuracy interval of any orthogonal accuracy algorithm. We are currently looking at a promising candidate algorithm that maintains precision and accuracy in an integrated way, which will hopefully provide considerably improved accuracy bounds.

Acknowledgements

I am grateful to Klaus Schossmaier for his invaluable comments on the many versions of this paper. A long list of enhancements was also provided by the anonymous referees, which greatly helped me to improve the exposition of the analysis. Their considerable efforts spent on scrutinizing the long manuscript are acknowledged.

References

[AK96] M. H. Azadmanesh and Roger M. Kieckhafer. New hybrid fault models for asynchronous approximate agreement. *IEEE Transactions on Computers*, 45(4):439–449, 1996.

[BI96] R. R. Brooks and S. S. Iyengar. Robust distributed computing and sensing algorithms. *IEEE Computer*, pages 53–60, June 1996.

[Dan97] Peter H. Dana. Global Positioning System (GPS) time dissemination for real-time applications. *Real-Time Systems*, 12(1):9–40, January 1997.

[DB93] Rogério Drummond and Özalp Babaoğlu. Low-cost clock synchronization. *Distributed Computing*, 6:193–203, 1993.

[DHS86] Danny Dolev, Joseph Y. Halpern, and H. Raymond Strong. On the possibility and impossibility of achieving clock synchronization. *Journal of Computer and System Sciences*, 32:230–250, 1986.

[FC95a] Christof Fetzer and Flaviu Cristian. Lower bounds for function based clock synchronization. In *Proc. 14th ACM Symposium on Principles of Distributed Computing*, Ottawa, CA, August 1995.

[FC95b] Christof Fetzer and Flaviu Cristian. An optimal internal clock synchronization algorithm. In *Proc. 10th Annual IEEE Conference on Computer Assurance*, Gaithersburg, MD, June 1995.

[HS97] Dieter Höchtl and Ulrich Schmid. Long-term evaluation of GPS timing receiver failures. In *Proc. of the 29th IEEE PTTI Systems and Application Meeting*, pages 165–180, Long Beach, California, December 1997.

[Knu73] Donald E. Knuth. *Fundamental Algorithms*, volume 1 of *The Art of Computer Programming*. Addison-Wesley, Reading, Massachusetts, 2nd edition, 1973.

[Lam87] Leslie Lamport. Synchronizing time servers. Technical Report 18, Digital System Research Center, 1987.

[Lis93] Barbara Liskov. Practical uses of synchronized clocks in distributed systems. *Distributed Computing*, 6:211–219, 1993.

[LWL88] Jennifer Lundelius-Welch and Nancy A. Lynch. A new fault-tolerant algorithm for clock synchronization. *Information and Computation*, 77(1):1–36, 1988.

[Mar84] Keith A. Marzullo. *Maintaining the Time in a Distributed System: An Example of a Loosely-Coupled Distributed Service*. PhD dissertation, Stanford University, Department of Electrical Engineering, February 1984.

[Mar90] Keith A. Marzullo. Tolerating failures of continuous-valued sensors. *ACM Transactions on Computer Systems*, 8(4):284–304, 1990.

[Mil95] David L. Mills. Improved algorithms for synchronizing computer network clocks. *IEEE Transactions on Networks*, pages 245–254, June 1995.

[MS85] Stephen R. Mahaney and Fred B. Schneider. Inexact agreement: Accuracy, precision, and graceful degradation. In *Proc. 4th ACM Symposium on Principles of Distributed Computing*, pages 237–249, Minaki, Canada, August 1985.

[OSF92] OSF. *Introduction to OSF DCE*. Prentice Hall, Englewood Cliffs, NJ, 1992.

[Sch86] Fred B. Schneider. A paradigm for reliable clock synchronization. In *Proc. Advanced Seminar of Local Area Networks*, pages 85–104, Bandol, France, April 1986.

[Sch87] Fred B. Schneider. Understanding protocols for byzantine clock synchronization. Technical Report 87-859, Cornell University, Department of Computer Science, August 1987.

[Sch95] Ulrich Schmid. Synchronized universal time coordinated for distributed real-time systems. *Control Engineering Practice*, 3(6):877–884, 1995. (Reprint from Proc. IFAC Workshop on Real-Time Programming WRTP'94, Lake Reichenau/Germany, 1994, p. 101–107.).

[Sch97a] Ulrich Schmid. Interval-based clock synchronization with optimal precision. Technical Report 183/1-78, Department of Automation, Technische Universität Wien, July 1997.

[Sch97b] Ulrich Schmid, editor. *Special Issue on The Challenge of Global Time in Large-Scale Distributed Real-Time Systems*, J. Real-Time Systems 12(1–3), 1997.

[Sch97c] Klaus Schossmaier. An interval-based framework for clock rate synchronization algorithms. In *Proc. 16th ACM Symposium on Principles of Distributed Computing*, pages 169–178, St. Barbara, USA, August 1997.

[SKM+99] Ulrich Schmid, Johann Klasek, Thomas Mandl, Herbert Nachtnebel, Gerhard R. Cadek, and Nikolaus Kerö. A Network Time Interface M-Module for distributing GPS-time over LANs. *J. Real-Time Systems*, 17(2), 1999. (to appear).

[SS97a] Ulrich Schmid and Klaus Schossmaier. Interval-based clock state synchronization revisited. Technical Report 183/1-80, Technische Universität Wien, Department of Automation, July 1997.

[SS97b] Ulrich Schmid and Klaus Schossmaier. Interval-based clock synchronization. *J. Real-Time Systems*, 12(2):173–228, March 1997.

[SS99] Ulrich Schmid and Klaus Schossmaier. Fault-tolerant interval algorithms. Technical Report 183/1-96, Department of Automation, TU Vienna, September 1999. (forthcoming).

[ST87] T. K. Srikanth and Sam Toueg. Optimal clock synchronization. *Journal of the ACM*, 34(3):626–645, July 1987.

[VR92] Paulo Veríssimo and Luís Rodrigues. A posteriori agreement for fault-tolerant clock synchronization on broadcast networks. In *Proc. 22nd International Symosium on Fault-Tolerant Computing*, Boston, Massachusetts, July 1992.

[VRC97] Paulo Veríssimo, Luís Rodrigues, and Antonio Casimiro. Cesiumspray: a precise and accurate global clock service for large-scale systems. *J. Real-Time Systems*, 12(3):243–294, 1997.

[YM93] Z. Yang and T. A. Marsland. Annotated bibliography on global states and times in distributed systems. *ACM SIGOPS Operating Systems Review*, pages 55–72, June 1993.

A Technical Lemmas

Lemma 6 (Properties Discrete Reference Point Setting) *Let* $\boldsymbol{I} = [a, b]$ *with* $0 \leq a \leq b$ *being integer multiples of* $G_S > 0$ *and an arbitrary interval* $\boldsymbol{\pi} = [-\pi^-, \pi^+]$ *with* $\pi = \pi^- + \pi^+ > 0$ *be given. Then,*

$$\boldsymbol{\pi}\text{-}center_{G_S}([a, b]) \leq \boldsymbol{\pi}\text{-}center_{G_S}([a + x, b + y]) \tag{58}$$

for any $x, y \geq 0$ being integer multiples of G_S, and the accuracies in the interval $[r \pm \boldsymbol{\alpha}]$ obtained from $\boldsymbol{I}$ by setting the reference point to $r = \boldsymbol{\pi}\text{-center}_{G_S}(\boldsymbol{I})$ satisfy

$$\alpha^- = \Big\lfloor \frac{\pi^-}{\pi} |\boldsymbol{I}| \Big\rfloor_{G_S}, \tag{59}$$

$$\alpha^+ = \Big\lceil \frac{\pi^+}{\pi} |\boldsymbol{I}| \Big\rceil_{G_S}. \tag{60}$$

Proof: Monotonicity (58) follows immediately from the definition (19) of $\boldsymbol{\pi}\text{-center}_{G_S}$ and monotonicity of $\lfloor x \rfloor$. The expressions for positive and negative accuracy yield

$$\alpha^- = \Big\lfloor \frac{\pi^- b + \pi^+ a}{\pi} \Big\rfloor_{G_S} - a = \Big\lfloor \frac{\pi^- b + \pi^+ a - \pi a}{\pi} \Big\rfloor_{G_S} = \Big\lfloor \frac{\pi^-}{\pi} |\boldsymbol{I}| \Big\rfloor_{G_S}$$

and

$$\alpha^+ = b - \Big\lfloor \frac{\pi^- b + \pi^+ a}{\pi} \Big\rfloor_{G_S} = b + \Big\lceil -\frac{\pi^- b + \pi^+ a}{\pi} \Big\rceil_{G_S} = \Big\lceil \frac{\pi^+}{\pi} |\boldsymbol{I}| \Big\rceil_{G_S},$$

where we used the well-known fact $-\lceil x \rceil = \lfloor -x \rfloor$ (see e.g. [Knu73, Sec. 1.2.4, Ex. 4]). □

Lemma 7 (Precision Enhancement) *Let $\boldsymbol{I}_p$, $\boldsymbol{I}_q$ be two consistent intervals with length $0 \leq |\boldsymbol{I}_p| \leq \bar{\pi}_p$, $0 \leq |\boldsymbol{I}_q| \leq \bar{\pi}_q$ being integer multiples of G_S, and ref$(\boldsymbol{I}_p) = \boldsymbol{\pi}_p\text{-center}_{G_S}(\boldsymbol{I}_p)$, ref$(\boldsymbol{I}_q) = \boldsymbol{\pi}_q\text{-center}_{G_S}(\boldsymbol{I}_q)$ for some $\boldsymbol{\pi}_p = [-\pi_p^-, \pi_p^+]$ and $\boldsymbol{\pi}_q = [-\pi_q^-, \pi_q^+]$. If $|\boldsymbol{I}_p \cup \boldsymbol{I}_q| \leq \pi$ with $\max\{\bar{\pi}_p, \bar{\pi}_q\} \leq \pi \leq \bar{\pi}_p + \bar{\pi}_q$, then*

$$|\textit{ref}(\boldsymbol{I}_p) - \textit{ref}(\boldsymbol{I}_q)| \leq \begin{cases} \max\Big\{ \Big\lceil \frac{\pi_q^-}{\pi_q}\bar{\pi}_q + \frac{\pi_p^+}{\pi_p}(\pi - \bar{\pi}_q) \Big\rceil_{G_S}, \Big\lceil \frac{\pi_p^-}{\pi_p}\bar{\pi}_p + \frac{\pi_q^+}{\pi_q}(\pi - \bar{\pi}_p) \Big\rceil_{G_S} \Big\} & \textit{if } \frac{\pi_q^-}{\pi_q} \geq \frac{\pi_p^+}{\pi_p}, \\ \max\Big\{ \Big\lceil \frac{\pi_p^+}{\pi_p}\bar{\pi}_p + \frac{\pi_q^-}{\pi_q}(\pi - \bar{\pi}_p) \Big\rceil_{G_S}, \Big\lceil \frac{\pi_q^+}{\pi_q}\bar{\pi}_q + \frac{\pi_p^-}{\pi_p}(\pi - \bar{\pi}_q) \Big\rceil_{G_S} \Big\} & \textit{otherwise.} \end{cases}$$

Proof: Apart from discreteness of $\boldsymbol{\pi}_p\text{-center}_{G_S}$, this is a straightforward linear programming problem. One has to look out for an arrangement of the consistent intervals $\boldsymbol{I}_p = [a, b]$ and $\boldsymbol{I}_q = [c, d]$ that maximizes the distance of their reference points. More specifically, recalling the definition of $\boldsymbol{\pi}\text{-center}_{G_S}$ from Definition 7, we are interested in

$$u = \min\Big\{ \Big\lfloor \frac{\pi_p^+ a + \pi_p^- b}{\pi_p} \Big\rfloor_{G_S} - \Big\lfloor \frac{\pi_q^+ c + \pi_q^- d}{\pi_q} \Big\rfloor_{G_S} \Big\} \tag{61}$$

subject to $b - a \leq \bar{\pi}_p$, $a - b \leq 0$, $d - c \leq \bar{\pi}_q$, $c - d \leq 0$, $b \leq \pi$, $d \leq \pi$, $c \leq b$, $a \leq d$, where $a, b, c, d \geq 0$ are integer multiples of G_S. Note that (61) covers arrangements where ref$(\boldsymbol{I}_p) \leq$ ref$(\boldsymbol{I}_q)$ only; miminizing provides the maximum distance for negative values only. However, since the problem is symmetric in p and q, the maximum value for reverse arragements follows immediately by exchanging p and q.

Fortunately, it is not difficult to identify the worst case scenario. Ignoring discreteness for the moment, i.e., assuming $G_S = 0$, if

$$\frac{\pi_q^-}{\pi_q} \geq \frac{\pi_p^+}{\pi_p}, \tag{62}$$

the maximum distance of the reference points occurs when $\boldsymbol{I}_q$ is made as large as possible, i.e., $|\boldsymbol{I}_q| = \bar{\pi}_q$, leaving at most $|\boldsymbol{I}_p| = \pi - \bar{\pi}_q$ for $\boldsymbol{I}_p$; recall that we assumed $\pi \leq \bar{\pi}_p + \bar{\pi}_q$.

Returning to the discrete case, this suggests that the minimum value of (61) satisfies

$$-u = \Big\lceil \frac{\pi_p^+}{\pi_p} \lfloor \pi - \bar{\pi}_q \rfloor_{G_S} \Big\rceil_{G_S} + \lfloor \frac{\pi_q^-}{\pi_q} \bar{\pi}_q \rfloor_{G_S} \leq \Big\lceil \frac{\pi_p^+}{\pi_p} (\pi - \bar{\pi}_q) + \frac{\pi_q^-}{\pi_q} \bar{\pi}_q \Big\rceil_{G_S}, \tag{63}$$

recall Lemma 6; the upper bound follows easily from moving the second term into the first one and then omitting $\lfloor . \rfloor_{G_S}$.

However, we have to confirm that discreteness of $\boldsymbol{\pi}$-center$_{G_S}$ does not impose a different maximum value. More specifically, we have to show that —starting from (63)— shrinking $|\boldsymbol{I}_q| = \bar{\pi}_q$ by kG_S for some integer $k \geq 0$ and simultaneously enlarging $|\boldsymbol{I}_p| = \lfloor \pi - \bar{\pi}_q \rfloor_{G_S} \leq \pi - \bar{\pi}_q$ by the same amount does not lead to a larger value of $-u$. For any $k \geq 0$, we find

$$\begin{aligned} -u(k) &= \Big\lceil \frac{\pi_p^+}{\pi_p} (\lfloor \pi - \bar{\pi}_q \rfloor_{G_S} + kG_S) \Big\rceil_{G_S} + \Big\lfloor \frac{\pi_q^-}{\pi_q} (\bar{\pi}_q - kG_S) \Big\rfloor_{G_S} \\ &\leq \Big\lceil \frac{\pi_p^+}{\pi_p} (\lfloor \pi - \bar{\pi}_q \rfloor_{G_S} + kG_S) + \frac{\pi_q^-}{\pi_q} (\bar{\pi}_q - kG_S) \Big\rceil_{G_S} \\ &\leq \Big\lceil \frac{\pi_p^+}{\pi_p} \lfloor \pi - \bar{\pi}_q \rfloor_{G_S} + \frac{\pi_q^-}{\pi_q} \bar{\pi}_q \Big\rceil_{G_S}, \end{aligned} \tag{64}$$

by virtue of condition (62), thus confirming our conjecture (63).

If, on the other hand, condition (62) is not true, the minimum (considering $G_S = 0$ first) occurs when $\boldsymbol{I}_p$ is made as large as possible, providing the value

$$-u' = \frac{\pi_p^+}{\pi_p} \bar{\pi}_p + \frac{\pi_q^-}{\pi_q} (\pi - \bar{\pi}_p). \tag{65}$$

As before, this suggests the minimum value

$$-u' = \lfloor \frac{\pi_p^+}{\pi_p} \bar{\pi}_p \rfloor_{G_S} + \Big\lceil \frac{\pi_q^-}{\pi_q} \lfloor \pi - \bar{\pi}_p \rfloor_{G_S} \Big\rceil_{G_S} \leq \Big\lceil \frac{\pi_p^+}{\pi_p} \bar{\pi}_p + \frac{\pi_q^-}{\pi_q} (\pi - \bar{\pi}_p) \Big\rceil_{G_S} \tag{66}$$

for the discrete case. For any $k \geq 0$, we find

$$\begin{aligned} -u'(k) &= \Big\lfloor \frac{\pi_p^+}{\pi_p} (\bar{\pi}_p - kG_S) \Big\rfloor_{G_S} + \Big\lceil \frac{\pi_q^-}{\pi_q} (\lfloor \pi - \bar{\pi}_p \rfloor_{G_S} + kG_S) \Big\rceil_{G_S} \\ &\leq \Big\lceil \frac{\pi_p^+}{\pi_p} (\bar{\pi}_p - kG_S) + \frac{\pi_q^-}{\pi_q} (\lfloor \pi - \bar{\pi}_p \rfloor_{G_S} + kG_S) \Big\rceil_{G_S} \\ &\leq \Big\lceil \frac{\pi_p^+}{\pi_p} \bar{\pi}_p + \frac{\pi_q^-}{\pi_q} \lfloor \pi - \bar{\pi}_p \rfloor_{G_S} \Big\rceil_{G_S} \end{aligned} \tag{67}$$

,

which also confirms (66) and completes the case of arrangements $\text{ref}(\boldsymbol{I}_p) \leq \text{ref}(\boldsymbol{I}_q)$.

To establish the result for reverse arrangements, we only have to exchange p and q. Rewriting condition (62) appropriately reads

$$\frac{\pi_p^-}{\pi_p} = 1 - \frac{\pi_p^+}{\pi_p} \geq 1 - \frac{\pi_q^-}{\pi_q} = \frac{\pi_q^+}{\pi_q},$$

which is fulfilled iff (62) is satisfied. Exchanging p and q in (63) and (66) and taking the maximum of the appropriate values eventually completes the proof of the lemma. □

Lemma 8 (Uniform Bounds m-Maxima) *Let $\mathcal{S} = \{z_i\}_{1 \le i \le n}$ with $z_i = x_i + y_i$ and y be given, such that $x_i \le x_j$ for $i < j$ and $y_i \le y$ for $1 \le i, j \le n$. If $\overline{\mathcal{S}} = \{w_i\}_{1 \le i \le n}$ with $w_i = x_i + y$ for $1 \le i \le n$, then the respective m-th largest elements satisfy $\max_m \mathcal{S} \le \max_m \overline{\mathcal{S}}$.*

Proof: If $\max_m \mathcal{S} = z_v$ for some $1 \le v \le n$, we claim that there exists some $p \le n-m+1$ such that $z_v \le z_p$ since otherwise $z_1 < z_v, \ldots, z_{n-m+1} < z_v$. However, z_v is the m-largest element in $\mathcal{S}$, which means that there are only $n - m$ z_i's that could possible satisfy $z_i < z_v$, providing the required contradiction. Since of course $w_i \le w_j$ for $i < j$, we have $\max_m \mathcal{S} = z_v \le z_p \le w_p \le w_{n-m+1} = \max_m \overline{\mathcal{S}}$. □

B Generic Algorithm

This appendix contains a restatement of the generic interval-based clock synchronization algorithm of [SS97b], along with a brief mentioning of the many parameters found in our system model. It is provided for the ease of reference only; please consult the original paper for details.

Definition 9 (Generic Algorithm [SS97b, Def. 7]) *With the parameters required for the instance of the algorithm running at node q,*

- *node q's intrinsic inverse rate deviation bound $\boldsymbol{\rho}_q$ and uniform bound $\boldsymbol{\rho}_{\max} \supseteq \bigcup_p \boldsymbol{\rho}_p$ with $\rho_{\max} = |\boldsymbol{\rho}_{\max}| = \rho^-_{\max} + \rho^+_{\max}$ ([SS97b, Assum. 2]),*
- *clock granularity G, clock setting granularity G_S, node q's maximum rate adjustment uncertainty $\boldsymbol{u}_q = [-u_q^-, u_q^+]$, and uniform maximum rate adjustment uncertainty $\boldsymbol{u}_{\max} \supseteq \bigcup_p \boldsymbol{u}_p$ with $u_{\max} = u^-_{\max} + u^+_{\max}$ ([SS97b, Assum. 2]),*
- *transmission delay characteristics δ_{sq}, $\boldsymbol{\varepsilon}_{sq}$ for all nodes $s \ne q$, uniform bounds $0 \le \delta_{\min} \le \min_{p,q}\{\delta_{pq}\}$, $\delta_{\max} \ge \max_{p,q}\{\delta_{pq}\}$, $\boldsymbol{\varepsilon}_{\max} \supseteq \bigcup_{p,q} \boldsymbol{\varepsilon}_{pq}$ with $\varepsilon_{\max} = |\boldsymbol{\varepsilon}_{\max}| = \varepsilon^-_{\max} + \varepsilon^+_{\max}$, "indicator" of broadcast network $B \in \{1, 2\}$, and accuracy transmission loss G_A ([SS97b, Assum. 4]),*
- *computation delay compensation Γ_q (integer multiple of G) guaranteeing node q's maximum computation time γ_q ([SS97b, Assum. 1]), chosen according to*

$$\Gamma_q \ge \frac{\gamma_q + u_q^-}{1 - \rho_q^-},$$

  and uniform bounds $\Gamma_{\max} \ge \max_p\{\Gamma_p\}$ and $0 \le \Gamma_{\min} \le \min_p\{\Gamma_p\}$; usually $\Gamma_p = \Gamma_{\max} = \Gamma_{\min}$ is the same for all nodes p,
- *broadcast delay compensation $\Lambda + \Omega$ (integer multiple of G), chosen to satisfy*

$$\Lambda + \Omega \ge \frac{\lambda_{\max} + \omega_{\max} + u^-_{\max}}{1 - \rho^-_{\max}};$$

  in conjunction with Δ below, it ensures that resynchronization starts only after all CSMs broadcast by non-faulty nodes during an FME have arrived ([SS97b, Assum. 4]),

- *transmission delay compensation* Δ *(integer multiple of* G*) chosen according to*

$$\Delta \;\geq\; \frac{\pi_0 + u_{\max} + \delta_{\max} + (P - \Gamma_{\min} + \pi^-)\rho_{\max} + \varepsilon^+_{\max}}{1 + \rho^+_{\max}}, \tag{68}$$

  where π_0 *and* π^- *depend on the particular convergence function employed (provided by the appropriate analysis),*

- *round period* $P \geq \Lambda + \Omega + \Delta + \Gamma_{\max}$ *(integer multiple of* G*),*

where all parameters are integer multiples of G_S *unless otherwise specified, our generic algorithm is defined as follows:*

(0) **Initial Synchronization***: At each node* q*, the local interval clock* $\boldsymbol{C}_q$ *has to be initialized to the accuracy interval* $\boldsymbol{A}^0_q = [T^0_q - \alpha^{0-}_q, T^0_q + \alpha^{0+}_q]$ *at some synchronous real-time* t^0_q *by some external means. This initialization must ensure*

- $t^0_q \in \boldsymbol{A}^0_q$,
- $T^0_q \in [\Lambda + \Omega + \Delta + \Gamma_q \pm \boldsymbol{\pi}]$,
- $\boldsymbol{\alpha}^0_q \subseteq \boldsymbol{\pi}_0$,

where $\boldsymbol{\pi}$ *and* $\boldsymbol{\pi}_0$ *depend on the particular convergence function employed (provided by the appropriate analysis).*

(1) **Periodic Synchronization***: Close to the end of each round* $k \geq 0$*, every node* q *in the system performs the following operations (the dependency of* T^I*,* t^I_q*, etc. upon round* k *is suppressed for brevity):*

(S) *CSM Send: Periodically at times* $C_q(t^I_q) = T^I = (k+1)P$*, node* q *initiates a broadcast. The message* M_{qp} *sent to node* p *at some real-time* t^A_{qp} *during that broadcast operation contains the accuracy interval* $\boldsymbol{A}_{qp} = [T^A_{qp} \pm \boldsymbol{\alpha}^A_{qp}] = \boldsymbol{C}_q(t^A_{qp})$*. For the zero-delay "loop-back transmission" to the own node* q*,* $t^A_{qq} = t^I_q$ *so that* $T^A_{qq} = T^I = (k+1)P$*.*

(R) *CSM Reception: If a clock synchronization message* M_{pq} *from node* p *arrives at node* q *at real-time* t^p_q*, when* $C_q(t^p_q) = T^p_q$*, the interval*

$$\boldsymbol{I}^p_q = \begin{cases} \boldsymbol{A}_{pq} + [T^R_q - T^p_q + \delta_{pq} \pm \boldsymbol{2G}_A + \boldsymbol{\varepsilon}_{pq}] + (T^R_q - T^p_q)\boldsymbol{\rho}_q + \boldsymbol{u}_q + \overline{\boldsymbol{G}} \\ \boldsymbol{A}_{qq} + T^R_q - T^A_{qq} + (T^R_q - T^A_{qq})\boldsymbol{\rho}_q \qquad \text{for } p = q \end{cases}$$

is computed and stored in a set $\boldsymbol{\mathcal{I}}_q$*. For the definition of the resynchronization time* T^R_q*, see Step (T).*

(C) *Computation: At real-time* $t^{\Lambda+\Omega+\Delta}_q$ *defined by* $C_q(t^{\Lambda+\Omega+\Delta}_q) = T^{\Lambda+\Omega+\Delta} = (k+1)P + \Lambda + \Omega + \Delta$*, the convergence function* $\boldsymbol{\mathcal{CV}}$ *is applied to the compatible intervals stored in* $\boldsymbol{\mathcal{I}}_q$*, yielding the interval* $\boldsymbol{R}_q$*. In addition,* $\boldsymbol{\mathcal{I}}_q$ *is re-initialized to the empty set for the next round.*

(T) *Termination and Resynchronization: At real-time* t^R_q *defined by* $C_q(t^R_q) = T^R_q = (k+1)P+\Lambda+\Omega+\Delta+\Gamma_q$*, node* q*'s interval clock* $\boldsymbol{C}_q$ *is set to* $\boldsymbol{R}_q$ *(instantaneously or by continuous amortization).*

C Generic Analysis

This appendix provides an extension of the definition of the generic convergence function [SS97b, Def. 7] and the major theorem [SS97b, Thm. 1], which considerably improve (and hence replace) the original versions. They have been developed in the technical report [SS97a] and are restated here, along with their justification and proofs, for further referencing.

We mentioned already that all results obtained in [SS97b] are generic in the sense that they are expressed in terms of a few characteristic parameters of the convergence function defined in [SS97b, Def. 11], namely,

- accuracy preservation function $\aleph(.)$, giving bounds on the provided interval of accuracies,
- precision preservation function $\Phi(.)$, giving the precision of the provided accuracy interval w.r.t. τ^k,
- precision enhancement function $\Pi(.)$, giving the precision of the provided accuracy interval w.r.t. τ^{k+1}.

In our original Definition [SS97b, Def. 11], we required $\mathcal{CV}$ to be both translation invariant and weakly monotonic according to the following definition:

Definition 10 (Translation Invariance & Weak Monotonicity [SS97b, Def. 10]) *Given two sets $\mathcal{I} = \{\boldsymbol{I}_1, \ldots, \boldsymbol{I}_n\}$ and $\mathcal{J} = \{\boldsymbol{J}_1, \ldots, \boldsymbol{J}_n\}$ of $n \geq 1$ accuracy intervals, an interval-valued function $\boldsymbol{f}()$ of $n \geq 1$ interval arguments is called*

(1) weakly monotonic *iff $\boldsymbol{I}_j \subseteq \boldsymbol{J}_j$ with ref$(\boldsymbol{I}_j) =$ ref$(\boldsymbol{J}_j)$ for all $1 \leq j \leq n$ implies $\boldsymbol{f}(\mathcal{I}) \subseteq \boldsymbol{f}(\mathcal{J})$,*

(2) translation invariant *iff $\boldsymbol{f}(\boldsymbol{I}_1 + \Delta, \ldots, \boldsymbol{I}_n + \Delta) = \boldsymbol{f}(\boldsymbol{I}_1, \ldots, \boldsymbol{I}_n) + \Delta$ for any real Δ.*

However, we recognized that there is no need to require weak monotonicity for $\mathcal{CV}$ itself: Monotonicity is used in [SS97b] to justify that bounds on input intervals carry over to bounds on the result. Still, there is no need to establish this property for $\mathcal{CV}$, since it suffices to show that all characteristic functions are monotonic w.r.t. their interval arguments. The modified definition of the generic convergence function in Definition 11 now accounts for this fact.

More importantly, in the course of analyzing particular convergence functions, we eventually recognized that applying the generic framework of [SS97b] literally provides overly conservative[7] worst case accuracy bounds. Intuitively, this is due to the fact that we simply added up the worst case enlargement of, say, α_p^+ —provided by the convergence function's accuracy preservation function $\aleph^+(.)$— in every round. The worst case enlargement of α_p^+, however, cannot occur successively in consecutive rounds. This is due to the fact that the occurence of the worst case setting usually yields an initial precision (i.e.,

[7] We do not mean the obvious lacking of a good worst case bound on $\alpha_p = \alpha_p^- + \alpha_p^+$ here. In fact, since the worst case scenario for negative and positive accuracy cannot occur simultaneously, simply adding the bounds on α_p^+ and α_p^- provides an overly conservative bound on α_p only.

after resnchronization) that is better than the worst case one. This means that the initial precision in the next round is smaller than the worst case one assumed for the previous round, which in turn prohibits the occurence of the worst case enlargement of α_p^+ then.

The key idea used for improving our analysis is to take into account the common intersection of those $\boldsymbol{\pi}$-precision intervals associated with the convergence function's input intervals that eventually determine the worst case setting for, say, α_p^+. By feeding its length ι as an additional parameter into the accuracy preservation function $\aleph^+(.)$, the worst case enlargement of α_p^+ can be conditioned on the common intersection actually present. Note that a lower bound was found to be sufficient for this purpose, since excessive adjustments happen for small intersections only.[8] An additional characteristic function $\Im^+(.)$ resp. $\Im^-(.)$ is introduced for keeping track of how the convergence function affects ι in worst case scenarios w.r.t. positive resp. negative accuracy. This (conditional) intersection enhancement function effectively determines ι for the next round, i.e., propagates the required information over multiple rounds.

Making this idea working in practice, however, is tricky for several reasons: First of all, it is the particular convergence function that determines how many/which input intervals are involved in the worst case accuracy setting. Leaving $\boldsymbol{\mathcal{CV}}$ unspecified in our generic framework thus makes it necessary to supply all the information required for constructing the particular ι from our general expressions. Moreover, different ι^- resp. ι^+ are usually required for computing the worst case bound for α_p^- resp. α_p^+. In fact, it may even be the case that some ι depending upon p is required for determining the worst case accuracy at node p, see [Sch97a] for an example. It is worth noting, however, that something like $\iota^p < \iota^q$ implies that the computed worst case accuracy bounds for nodes p and q cannot both be tight.

To cope with those problems, we actually utilize individual lower bounds ι_s^+ resp. ι_s^-, $1 \leq s \leq n$, on the length of the intersection $\hat{\boldsymbol{I}}_p^s \cap \hat{\boldsymbol{I}}_q^s \cap \hat{\boldsymbol{I}}_p^{\min_p} \cap \hat{\boldsymbol{I}}_q^{\min_q}$ resp. $\hat{\boldsymbol{I}}_p^s \cap \hat{\boldsymbol{I}}_q^s \cap \hat{\boldsymbol{I}}_p^{\max_p} \cap \hat{\boldsymbol{I}}_q^{\max_q}$, where $\min_x$ resp. $\max_x$ represents that non-faulty node that leads to the leftmost right($\hat{\boldsymbol{I}}_x^{\min_x}$) resp. the rightmost left($\hat{\boldsymbol{I}}_x^{\max_x}$) for $x \in \{p, q\}$. Clearly, ι_s^+ resp. ι_s^- is meant for the worst case accuracy setting for α_p^+ resp. α_p^-, so $\forall s : \iota_s^+$ resp. $\forall s : \iota_s^-$ are supplied as parameters to $\aleph^+(.)$ resp. $\aleph^-(.)$ and $\Im^+(.)$ resp. $\Im^-(.)$ in Definition 11.

In order to be able to inductively compute an expression for, say, ι_p^+, the conditional intersection enhancement function $\Im^+(.)$ provides an upper bound $\pi_0^{\iota_{pq}^+}$ on the mutual precision of the $\boldsymbol{\mathcal{CV}}$'s results $\{\boldsymbol{R}_p, \boldsymbol{R}_q\}$ computed at node p, q under the worst case accuracy setting for α_p^+. Since $\boldsymbol{R}_p$ and $\boldsymbol{R}_q$ are both $\boldsymbol{\pi}_0$-correct, this implies that the mutual intersection $\hat{\boldsymbol{R}}_p \cap \hat{\boldsymbol{R}}_q$ for any q (including $q = \min_x$) has length at least $\overline{\tau}_p^+ = \pi_0 - \max_q \pi_0^{\iota_{pq}^+}$. Now, since the drift and delay compensation operations applied for computing the input intervals fed into the convergence function are accuracy preserving, it follows that any initial intersection is necessarily preserved during a round. Thus, any intersection of source intervals like $\hat{\boldsymbol{I}}_p^s \cap \hat{\boldsymbol{I}}_q^{\min_q}$ must also have length at least $\overline{\tau}_s^+$, which implies that we can simply use $\iota_s^+ = \overline{\tau}_s^+$.

The above considerations lead to the following modified definition of the generic convergence function of [SS97b]:

[8] At least for any particular convergence function considered up to now.

Definition 11 (Generic Convergence Function $\mathcal{CV}$, cf. [SS97b, Def. 7]) *Let $\mathcal{I}_p = \{\boldsymbol{I}_p^1, \ldots, \boldsymbol{I}_p^n\}$ resp. $\mathcal{I}_q = \{\boldsymbol{I}_q^1, \ldots, \boldsymbol{I}_q^n\}$, $q \neq p$, be two ordered sets of n compatible intervals (all representing the same real-time t) obtained at nodes p resp. q at the end of a round, which are in accordance with a given fault model $\mathcal{F}$. Assuming that*

[1] any non-faulty $\boldsymbol{I}_p^i$ is $\boldsymbol{\pi}_p^i$-correct for $\boldsymbol{\pi}_p^i \in \boldsymbol{\mathcal{P}}_p = \{\boldsymbol{\pi}_p^1, \ldots, \boldsymbol{\pi}_p^n\}$ denoting a given set of precision bounds (and analogously for $\boldsymbol{I}_q^i$ with set of precision bounds $\boldsymbol{\mathcal{P}}_q = \{\boldsymbol{\pi}_q^1, \ldots, \boldsymbol{\pi}_q^n\}$),

[1'] $\boldsymbol{\mathcal{P}} = \{\boldsymbol{\pi}^1, \ldots, \boldsymbol{\pi}^n\}$ with $\boldsymbol{\pi}_p^i \cup \boldsymbol{\pi}_q^i \subseteq \boldsymbol{\pi}^i \subseteq \boldsymbol{\pi}^H$, for some suitable $\boldsymbol{\pi}^H$, denotes a set of uniform precision bounds ensuring $\boldsymbol{\pi}^i$-correctness of both $\boldsymbol{I}_p^i$ and $\boldsymbol{I}_q^i$ (if non-faulty),

[2] any pair of non-faulty intervals $\{\boldsymbol{I}_p^i, \boldsymbol{I}_q^i\}$ is $\boldsymbol{\pi}_I$-precise for some $\boldsymbol{\pi}_I \subseteq \boldsymbol{\pi}^H$,

[3] for any s with both $\boldsymbol{I}_p^s$ and $\boldsymbol{I}_q^s$ being non-faulty, the intersection of the associated precision intervals $\hat{\boldsymbol{I}}_p^s \cap \hat{\boldsymbol{I}}_q^s \cap \hat{\boldsymbol{I}}_p^{\min_p} \cap \hat{\boldsymbol{I}}_q^{\min_q}$ resp. $\hat{\boldsymbol{I}}_p^s \cap \hat{\boldsymbol{I}}_q^s \cap \hat{\boldsymbol{I}}_p^{\max_p} \cap \hat{\boldsymbol{I}}_q^{\max_q}$, where $\min_x$ resp. $\max_x$ represents that non-faulty node that leads to the leftmost right($\hat{\boldsymbol{I}}_x^{\min_x}$) resp. the rightmost left($\hat{\boldsymbol{I}}_x^{\max_x}$) for $x \in \{p, q\}$, has length at least $\iota_s^+ \geq 0$ resp. $\iota_s^- \geq 0$ (integer multiples of G_S),

[4] the accuracies of any non-faulty $\boldsymbol{I}_p^i = [T_p^i \pm \boldsymbol{\alpha}_p^i]$ are integer multiples of G_S satisfying $\boldsymbol{\alpha}_p^i \subseteq \boldsymbol{\beta}_p^i \in \boldsymbol{\mathcal{B}}_p$ for a given set of accuracy bounds $\boldsymbol{\mathcal{B}}_p = \{\boldsymbol{\beta}_p^1, \ldots, \boldsymbol{\beta}_p^n\}$ (and analogously for $\boldsymbol{I}_q^i$ with set of accuracy bounds $\boldsymbol{\mathcal{B}}_q$),

let

$$\begin{aligned} \boldsymbol{R}_p &= \boldsymbol{\mathcal{CV}}(\boldsymbol{\mathcal{I}}_p) = [T_p' \pm \boldsymbol{\alpha}_p'] \\ \boldsymbol{R}_q &= \boldsymbol{\mathcal{CV}}(\boldsymbol{\mathcal{I}}_q) = [T_q' \pm \boldsymbol{\alpha}_q'] \\ &. \end{aligned} \tag{69}$$

The generic convergence function $\boldsymbol{\mathcal{CV}}$ must be translation invariant and should provide accurate intervals with reference point and accuracies being integer multiples of G_S. Its behavior is characterized by the following functions, which must be monotonic w.r.t. any interval argument:

(1) Precision preservation function $\boldsymbol{\Phi}(.)$, *so that $\boldsymbol{R}_p$ is $\boldsymbol{\Phi}(\boldsymbol{\mathcal{P}}_p, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I)$-correct and $\boldsymbol{R}_q$ is $\boldsymbol{\Phi}(\boldsymbol{\mathcal{P}}_q, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I)$-correct, with $|\boldsymbol{\Phi}(\boldsymbol{\mathcal{P}}, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I)| = \mathcal{O}(\pi^H)$ for $\pi^H = |\boldsymbol{\pi}^H|$.*

(2) Precision enhancement function $\Pi(.)$, *so that the set $\{\boldsymbol{R}_p, \boldsymbol{R}_q\}$ is $\boldsymbol{\pi}_0$-precise for any $\boldsymbol{\pi}_0$ satisfying $|\boldsymbol{\pi}_0| = \pi_0 = \Pi(\boldsymbol{\mathcal{P}}, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I)$, with $\Pi(\boldsymbol{\mathcal{P}}, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I) < \pi^H = |\boldsymbol{\pi}^H|$.*

(3) Conditional intersection enhancement functions $\Im^-(.)$ *resp.* $\Im^+(.)$, *so that the set $\{\boldsymbol{R}_p, \boldsymbol{R}_q\}$ is $\boldsymbol{\pi}_0^{\iota_{pq}^-}$-precise resp. $\boldsymbol{\pi}_0^{\iota_{pq}^+}$-precise with*

$$\begin{aligned} \pi_0^{\iota_{pq}^-} &= \Im^-(\boldsymbol{\mathcal{B}}_p, \boldsymbol{\mathcal{B}}_q, \boldsymbol{\mathcal{P}}_p, \boldsymbol{\mathcal{P}}_q, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I, \forall s : \iota_s^-) \\ \pi_0^{\iota_{pq}^+} &= \Im^+(\boldsymbol{\mathcal{B}}_p, \boldsymbol{\mathcal{B}}_q, \boldsymbol{\mathcal{P}}_p, \boldsymbol{\mathcal{P}}_q, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I, \forall s : \iota_s^+) \end{aligned}$$

for worst case accuracy settings w.r.t. α_p^- resp. α_p^+.

(4) Conditional accuracy preservation functions $\aleph^-(.)$, $\aleph^+(.)$, *so that*

$$\boldsymbol{\alpha}'_p \;\subseteq\; \left[-\aleph^-(\boldsymbol{\mathcal{B}}_p, \boldsymbol{\mathcal{P}}_p, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I, \forall s : \iota_s^-), \aleph^+(\boldsymbol{\mathcal{B}}_p, \boldsymbol{\mathcal{P}}_p, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I, \forall s : \iota_s^+)\right]$$

Note that both the precision enhancement function $\Pi(.)$ and the intersection enhancement functions $\Im^-(.)$, $\Im^+(.)$ provide a worst case precision, although for different classes of input scenarios: Whereas $\Pi(.)$ provides the precision for any input scenario, $\Im^-(.)$ resp. $\Im^+(.)$ is valid for scenarios leading to worst case α_p^- resp. α_p^+ according to item (4) only.

Another issue that was left open by the analysis of [SS97b] is the determination of *traditional accuracy*, which gives the amount local time may drift from real-time during a given time interval Δt. Although worst case bounds on accuracy intervals obviously provide an upper bound on traditional accuracy as well, this usually leads to overly conservative estimates. Therefore, an explicit expression for traditional accuracy is added to our major theorem. Moreover, taking the limit $\Delta t \to \infty$, traditional accuracy leads to the rate (and hence drift) of the synchronized clocks, which is more convenient for comparison. Note that appropriate worst case bounds for are available for most existing internal synchronization algorithms, see e.g. [LWL88], [MS85], [Sch87], [ST87], [VR92], [DB93], [FC95b], [VRC97].

Our notion of internal global time makes it easy to deal with traditional accuracy. We only have to bound the maximum "jump" internal global time can experience when switching from one round to the next. This is sufficient since internal global time progresses as real-time does during a round, so that no additional deviation from real-time occurs in between synchronization instants.

The appropriately modified and extended major result of the generic analysis in [SS97b] for the case of instantaneous adjustment[9] of the local clocks in step (T) of the algorithm in Definition 9 reads as follows:

Theorem 5 (Instantaneous Correction, cf. [SS97b, Thm. 1]) *Running in a system complying to [SS97b, Assum. 1–4], the clock synchronization algorithm of Definition 9 using the generic convergence function* $\boldsymbol{\mathcal{CV}}$ *—characterized by accuracy preservation* $\aleph^\pm(.)$*, precision preservation* $\boldsymbol{\Phi}(.)$*, precision enhancement* $\Pi(.)$*, and intersection enhancement* $\Im^\pm(.)$ *subject to a given fault model* $\mathcal{F}$*— guarantees accuracy and precision for all rounds* $k \geq 0$ *as follows:*

(0) The accuracy interval $\boldsymbol{A}_q^{k+1} = \boldsymbol{A}_q^{k+1}(t_q^{R,k}) = [T_q^{k+1} \pm \boldsymbol{\alpha}_q^{k+1}]$ *provided by the local interval clock of a non-faulty node* q *at the beginning of round* $k+1$*,* $k \geq 0$*, satisfies* $\boldsymbol{\alpha}_q^{k+1} \subseteq \boldsymbol{\beta}_q^{k+1}$ *with*

$$\begin{aligned} \boldsymbol{\beta}_q^{k+1} &= \left[-\aleph^-(\boldsymbol{\mathcal{B}}_q^{k+1}, \boldsymbol{\mathcal{P}}_q, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I, \forall s : \iota_s^{k,-}), \aleph^+(\boldsymbol{\mathcal{B}}_q^{k+1}, \boldsymbol{\mathcal{P}}_q, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I, \forall s : \iota_s^{k,+})\right] \quad (70) \\ \boldsymbol{\beta}_q^0 &= \boldsymbol{\alpha}_q^0, \end{aligned}$$

where $\boldsymbol{\mathcal{B}}_q^{k+1} = \{\boldsymbol{\beta}_q^{1,k+1}, \ldots, \boldsymbol{\beta}_q^{n,k+1}\}$ *is defined by* $\boldsymbol{\beta}_q^{p,k+1} = \boldsymbol{\beta}_p^k + \boldsymbol{\omega}_q^p$ *with*

$$\begin{aligned} \boldsymbol{\omega}_q^p &= \boldsymbol{u}_p + \boldsymbol{u}_q + \overline{\boldsymbol{G}} + \boldsymbol{2G}_A + \boldsymbol{\varepsilon}_{pq} \\ &\quad + (P - \Delta - \Gamma_p)\boldsymbol{\rho}_p + (\Gamma_q + \Delta - \delta_{pq})\boldsymbol{\rho}_q \end{aligned}$$

[9] Consult [SS97b] for how to carry over this result to clock adjustment via continuous amortization.

$$
\begin{aligned}
&+ (\Lambda+\Omega)[-\max\{\rho_q^- - \rho_p^-, 0\}, \max\{\rho_q^+ - \rho_p^+, 0\}] \\
&+ \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\boldsymbol{\rho}_{\max} \qquad \textit{for } p \neq q, \qquad (71)\\
\boldsymbol{\omega}_q^q &= \boldsymbol{u}_q + P\boldsymbol{\rho}_q + \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\boldsymbol{\rho}_q, \qquad (72)
\end{aligned}
$$

the set $\boldsymbol{\mathcal{P}}_q = \{\boldsymbol{\pi}_q^1, \ldots, \boldsymbol{\pi}_q^n\}$ *of node* q*'s precision bounds* $\boldsymbol{\pi}_q^p \subseteq \boldsymbol{\pi}^p \subseteq \boldsymbol{\pi}^H$ *—see item (1)— is defined by*

$$
\begin{aligned}
\boldsymbol{\pi}_q^p &= \boldsymbol{\pi}_0 + \boldsymbol{u}_p + \boldsymbol{u}_q + \overline{\boldsymbol{G}} + \boldsymbol{\varepsilon}_{pq} + (P - \Delta - \Gamma_p)\boldsymbol{\rho}_p + (\Gamma_q + \Delta - \delta_{pq})\boldsymbol{\rho}_q \\
&\quad + (\Lambda+\Omega)[-\max\{\rho_q^- - \rho_p^-, 0\}, \max\{\rho_q^+ - \rho_p^+, 0\}] \\
&\quad + \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\boldsymbol{\rho}_{\max} \qquad \textit{for } p \neq q, \qquad (73)\\
\boldsymbol{\pi}_q^q &= \boldsymbol{\pi}_0 + \boldsymbol{u}_q + P\boldsymbol{\rho}_q + \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\boldsymbol{\rho}_q, \qquad (74)
\end{aligned}
$$

and

$$
\begin{aligned}
\iota_q^{k+1,-} &= \pi_0 - \max_j \Im^-(\boldsymbol{\mathcal{B}}_q^{k+1}, \boldsymbol{\mathcal{B}}_j^{k+1}, \boldsymbol{\mathcal{P}}_q, \boldsymbol{\mathcal{P}}_j, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I, \forall s : \iota_s^{k,-}) \geq 0, \qquad (75)\\
\iota_q^{k+1,+} &= \pi_0 - \max_j \Im^+(\boldsymbol{\mathcal{B}}_q^{k+1}, \boldsymbol{\mathcal{B}}_j^{k+1}, \boldsymbol{\mathcal{P}}_q, \boldsymbol{\mathcal{P}}_j, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I, \forall s : \iota_s^{k,+}) \geq 0, \qquad (76)\\
\iota_q^{0,-} = \iota_q^{0,+} &= \pi_0 - \max_j\{\alpha_j^0\} \geq 0. \qquad (77)
\end{aligned}
$$

(1) The interval clocks of non-faulty nodes are synchronized to the (observable) initial worst case precision (i.e., the precision at the beginning of each round of the slowest non-faulty clock)

$$
\begin{aligned}
\pi_{0,\max} &= \pi_0 + u_{\max} + G + (\Gamma_{\max} - \Gamma_{\min})\rho_{\max} \\
&\quad + \mathcal{O}(P\rho_{\max}^2 + G\rho_{\max} + \varepsilon_{\max}\rho_{\max}) \qquad (78)
\end{aligned}
$$

with $\pi_0 = |\boldsymbol{\pi}_0|$*, where* $\boldsymbol{\pi}_0$ *is a solution of the equation*

$$
|\boldsymbol{\pi}_0| = \Pi(\boldsymbol{\mathcal{P}}, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I) \qquad (79)
$$

for the set $\boldsymbol{\mathcal{P}} = \{\boldsymbol{\pi}^1, \ldots, \boldsymbol{\pi}^n\}$ *of uniform precision bounds* $\boldsymbol{\pi}^p \subseteq \boldsymbol{\pi}^H$ *defined by*

$$
\begin{aligned}
\boldsymbol{\pi}^p &= \boldsymbol{\pi}_0 + \boldsymbol{u}_p + \boldsymbol{u}_{\max} + \overline{\boldsymbol{G}} + \boldsymbol{\varepsilon}_{\max} \\
&\quad + (P - \Delta - \Gamma_p)\boldsymbol{\rho}_p + (\Gamma_{\max} + \Delta - \delta_{\min})\boldsymbol{\rho}_{\max} \\
&\quad + (\Lambda+\Omega)[-(\rho_{\max}^- - \rho_p^-), \rho_{\max}^+ - \rho_p^+] \\
&\quad + \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\boldsymbol{\rho}_{\max}, \qquad (80)\\
\boldsymbol{\pi}^{\overline{q}} &= \boldsymbol{\pi}_0 + \boldsymbol{u}_q + P\boldsymbol{\rho}_q + \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\boldsymbol{\rho}_q \qquad (81)\\
\boldsymbol{\pi}^H &= \boldsymbol{\pi}_0 + 2\boldsymbol{u}_{\max} + \overline{\boldsymbol{G}} + \boldsymbol{\varepsilon}_{\max} + (P + \Gamma_{\max} - \Gamma_{\min} - \delta_{\min})\boldsymbol{\rho}_{\max} \\
&\quad + \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\boldsymbol{\rho}_{\max}, \qquad (82)\\
\boldsymbol{\pi}_I &= \boldsymbol{\varepsilon}_{\max} + B\boldsymbol{u}_{\max} + \overline{\boldsymbol{G}} + (\Lambda + \Omega + \Delta + \Gamma_{\max} - \delta_{\min})\boldsymbol{\rho}_{\max} \\
&\quad + \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\boldsymbol{\rho}_{\max}, \qquad (83)
\end{aligned}
$$

where $\boldsymbol{\pi}^{\overline{q}} \subseteq \boldsymbol{\pi}^q$ *denotes node* q*'s own (i.e., non-remote) precision bound.*

(2) The (observable) worst-case precision $\pi_{\max}$ *satisfies*

$$
\begin{aligned}
\pi_{\max} &= \max\Big\{\pi^- + u_{\max}^+ + (\Gamma_{\max} - \Gamma_{\min})\rho_{\max}^+, \\
&\qquad \pi^+ + u_{\max}^- + (\Gamma_{\max} - \Gamma_{\min})\rho_{\max}^-,\ \pi_0 + u_{\max} + P\rho_{\max}\Big\} \\
&\quad + G + \mathcal{O}(P\rho_{\max}^2 + G\rho_{\max} + \varepsilon_{\max}\rho_{\max}) \qquad (84)
\end{aligned}
$$

with

$$\begin{aligned}\boldsymbol{\pi} &= \overline{\boldsymbol{\Phi}}(\boldsymbol{\mathcal{P}}, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I) + \boldsymbol{\pi}_0 + \boldsymbol{u}_{\max} + P\boldsymbol{\rho}_{\max} \\ &\quad + \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\boldsymbol{\rho}_{\max}, \end{aligned} \tag{85}$$

where $\overline{\boldsymbol{\Phi}}(.)$ *denotes the result of* $\boldsymbol{\Phi}(.)$ *with swapped positive and negative accuracy.*

(3) Resynchronization of any two non-faulty nodes p, q *occurs within real-time* $t_p^R - t_q^R$ *satisfying*

$$\begin{aligned} t_p^R - t_q^R &\subseteq \Gamma_p - \Gamma_q + [-\pi_0, \pi_0] + \boldsymbol{u}_p + \overline{\boldsymbol{u}}_q + P(\boldsymbol{\rho}_p + \overline{\boldsymbol{\rho}}_q) \\ &\quad + \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})[-\rho_{\max}, \rho_{\max}], \end{aligned} \tag{86}$$

where clock adjustments Υ_q *of at most* $\Upsilon_q \in \boldsymbol{\pi}_q \subseteq \boldsymbol{\pi}$ *defined by*

$$\boldsymbol{\pi}_q = \overline{\boldsymbol{\Phi}}(\boldsymbol{\mathcal{P}}_q, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I) + \boldsymbol{\pi}_0 + \boldsymbol{u}_q + P\boldsymbol{\rho}_q + \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})\boldsymbol{\rho}_{\max} \tag{87}$$

are applied to the clock of a non-faulty node q.

(4) Let $\boldsymbol{\Phi} = \bigcup_q \boldsymbol{\Phi}(\boldsymbol{\mathcal{P}}_q, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I) \subseteq \boldsymbol{\Phi}(\boldsymbol{\mathcal{P}}, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I)$. *For any round* $k \geq 0$, $\tau^{k+1}(t) - \tau^k(t) \in \overline{\boldsymbol{\Phi}} - \overline{\boldsymbol{\pi}}_0$, *and the traditional accuracy at the beginning of round* $k+1$ *satisfies*

$$T_q^{k+1} - t_q^{R,k} \in \overline{\boldsymbol{\pi}}_0 + (k+1)(\overline{\boldsymbol{\Phi}} - \overline{\boldsymbol{\pi}}_0). \tag{88}$$

The inverse rate $r_{q,syn}^{-1}$ *of the synchronized clock at any node* q *evaluates to*

$$r_{q,syn}^{-1} = \lim_{k\to\infty} \frac{t_q^{R,k} - t_q^0}{T_q^{k+1} - T_q^0} \in \Big[1 \pm \frac{\boldsymbol{\Phi} - \boldsymbol{\pi}_0}{P}\Big], \tag{89}$$

where $T_q^0 = C_q(t_q^0)$ *is node* q*'s local time at the beginning of round* $k = 0$.

Proof: With the exception of item (0) and item (4), the proof of Theorem 1 in [SS97b] applies without modification. We just replaced π in certain remainder terms like (80) by the coarse bound $\pi = \mathcal{O}(P\rho_{\max} + G + \varepsilon_{\max})$ established in the proof of [SS97b, Lem. 12], and added (81) according to [SS97b, Eq. (58)], which gives the precision bound $\boldsymbol{\pi}^{\overline{q}}$ that applies for local (non-remote) intervals $\boldsymbol{I}_q^q$.

As far as item (0) is concerned, we found it convenient to introduce the additional abbreviation $\boldsymbol{\omega}_q^p$ in (71)/(72). Moreover, according to the modified Definition 11 of the generic convergence function, we splitteded up $\boldsymbol{\aleph}(.)$ in $\aleph^-(.)$ resp. $\aleph^+(.)$ and extended the list of parameters to include $\forall s : \iota_s^{k,-}$ resp. $\forall s : \iota_s^{k,+}$. The justification of the appropriate definitions (75)–(77) follows the line of reasoning outlined prior to Definition 11: An upper bound $\pi_0^{\iota_{qj}^+}$ on the precision of the set $\mathrm{shift}_{t'}(\{\boldsymbol{A}_q^k, \boldsymbol{A}_j^k\})$ (shifting the intervals to some common point in time t' just makes them compatible) at the beginning of a round k implies a lower bound $\iota_{qj}^+ = \pi_0 - \pi_0^{\iota_{qj}^+}$ on the length of the mutual intersection of the associated $\boldsymbol{\pi}_0$-precision intervals $\hat{\boldsymbol{A}}_q^k \cap \hat{\boldsymbol{A}}_j^k$. Taking the minimum over all j gives a lower bound $\iota_q^+ = \pi_0 - \max_j \pi_0^{\iota_{qj}^+}$ on that mutual intersection for arbitrary nodes j, including $j = \min_x$. The required $\pi_0^{\iota_{qj}^+}$, however, is given by the conditional intersection enhancement funtion

$\Im^+(.)$ for rounds $k > 0$, and by the initial synchronization assumption in Definition 9 for round $k = 0$.

It only remains to confirm that the precision intervals associated with any (non-faulty) $\boldsymbol{I}_p^s$ and $\boldsymbol{I}_q^s$ —fed into the convergence function at node p and q— have a mutual intersection with either $\boldsymbol{I}_p^{\min_p}$ and $\boldsymbol{I}_q^{\min_q}$ of length at least ι_s^+ as well. However, the drift and delay compensation operations used to obtain round k's $\boldsymbol{I}_x^s$ from $\boldsymbol{A}_s^k$ have been explicitly designed to preserve accurateness, so they must preserve any initial mutual intersection. Hence, choosing $\iota_q^{k+1,+}$ and $\iota_q^{k+1,-}$ according to (75)–(77) is legitimate.

Turning our attention to item (4), we know that $\boldsymbol{A}_q^{k+1}$ is $\boldsymbol{\Phi}(\boldsymbol{\mathcal{P}}_q, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I)$-correct w.r.t. τ^k by virtue of item (2) of Definition 11 and hence $\boldsymbol{\Phi}$-correct. Note that $\boldsymbol{\Phi} \subseteq \boldsymbol{\Phi}(\boldsymbol{\mathcal{P}}, \boldsymbol{\pi}^H, \boldsymbol{\pi}_I)$ is a simple consequence of monotonicity of $\boldsymbol{\Phi}(.)$. Moreover, from item (3) of Definition 11, it follows that all $\boldsymbol{A}_p^{k+1}$ are $\boldsymbol{\pi}_0$-precise. Hence it is possible to choose $\tau^{k+1} \in \bigcap_p \hat{\boldsymbol{A}}_p^{k+1}$, as justified by Definition 2, and we claim that we may in fact choose τ^{k+1} so that

$$-(\Phi^+ - \pi_0^+) \leq \tau^{k+1} - \tau^k \leq \Phi^- - \pi_0^-.$$

If this was not feasible, there would exist an interval $\hat{\boldsymbol{A}}_r^{k+1}$ of length π_0 with $\tau^{k+1} = \text{left}(\hat{\boldsymbol{A}}_r^{k+1})$ or else $\tau^{k+1} = \text{right}(\hat{\boldsymbol{A}}_r^{k+1})$ that satisfies $\tau^k \notin \hat{\boldsymbol{A}}_r^{k+1} + (\boldsymbol{\Phi} - \boldsymbol{\pi}_0) = \boldsymbol{\Phi}$. This, however, would contradict $\boldsymbol{\Phi}$-correctness of $\boldsymbol{A}_r^{k+1}$.

The bound (88) is a simple consequence of the fact that internal global time progresses as real-time does during a round, so that the maximum deviation between internal global time and real-time remains the same during any round. Therefore, we just have to add up the worst case internal global time "jumps" at each round. The initial deviation (in round 0) is zero since choosing $\tau^0(t) = t$ is legitimate due to the initial synchronization assumption in Definition 9. Hence, to complete the proof of (88), it only remains to add the maximum deviation between τ^{k+1} and the reference point T_q^{k+1} of $\boldsymbol{A}_q^{k+1}$, which is trivial since the latter is $\boldsymbol{\pi}_0$-correct.

To derive expression (89) for the rate of the synchronized clocks, we multiply (88) by -1 to arrive at

$$t_q^{R,k} - t_q^0 \in T_q^{k+1} - T_q^0 - (t_q^0 - T_q^0) + \boldsymbol{\pi}_0 + (k+1)(\boldsymbol{\Phi} - \boldsymbol{\pi}_0). \tag{90}$$

From the initial synchronization assumptions in item (0) of Definition 9, we gather $t_q^0 - T_q^0 \in \boldsymbol{\alpha}_q^0 \subseteq \boldsymbol{\pi}_0$. Moreover, from step (T) of the algorithm in conjunction with the fact that the maximum clock adjustment was shown to satisfy $\Upsilon_q \in \boldsymbol{\pi}$ in item (3) of this theorem, we obtain $T_q^{k+1} - T_q^0 = (k+1)P + \mathcal{O}(\pi)$. Plugging this into (90), we find

$$\frac{t_q^{R,k} - t_q^0}{T_q^{k+1} - T_q^0} \in 1 + \frac{[-\pi_0, \pi_0] + (k+1)(\boldsymbol{\Phi} - \boldsymbol{\pi}_0)}{(k+1)P + \mathcal{O}(\pi)}.$$

Taking the limit for $k \to \infty$ eventually provides (89) and completes the proof of our theorem. □

Overview of the Analysis

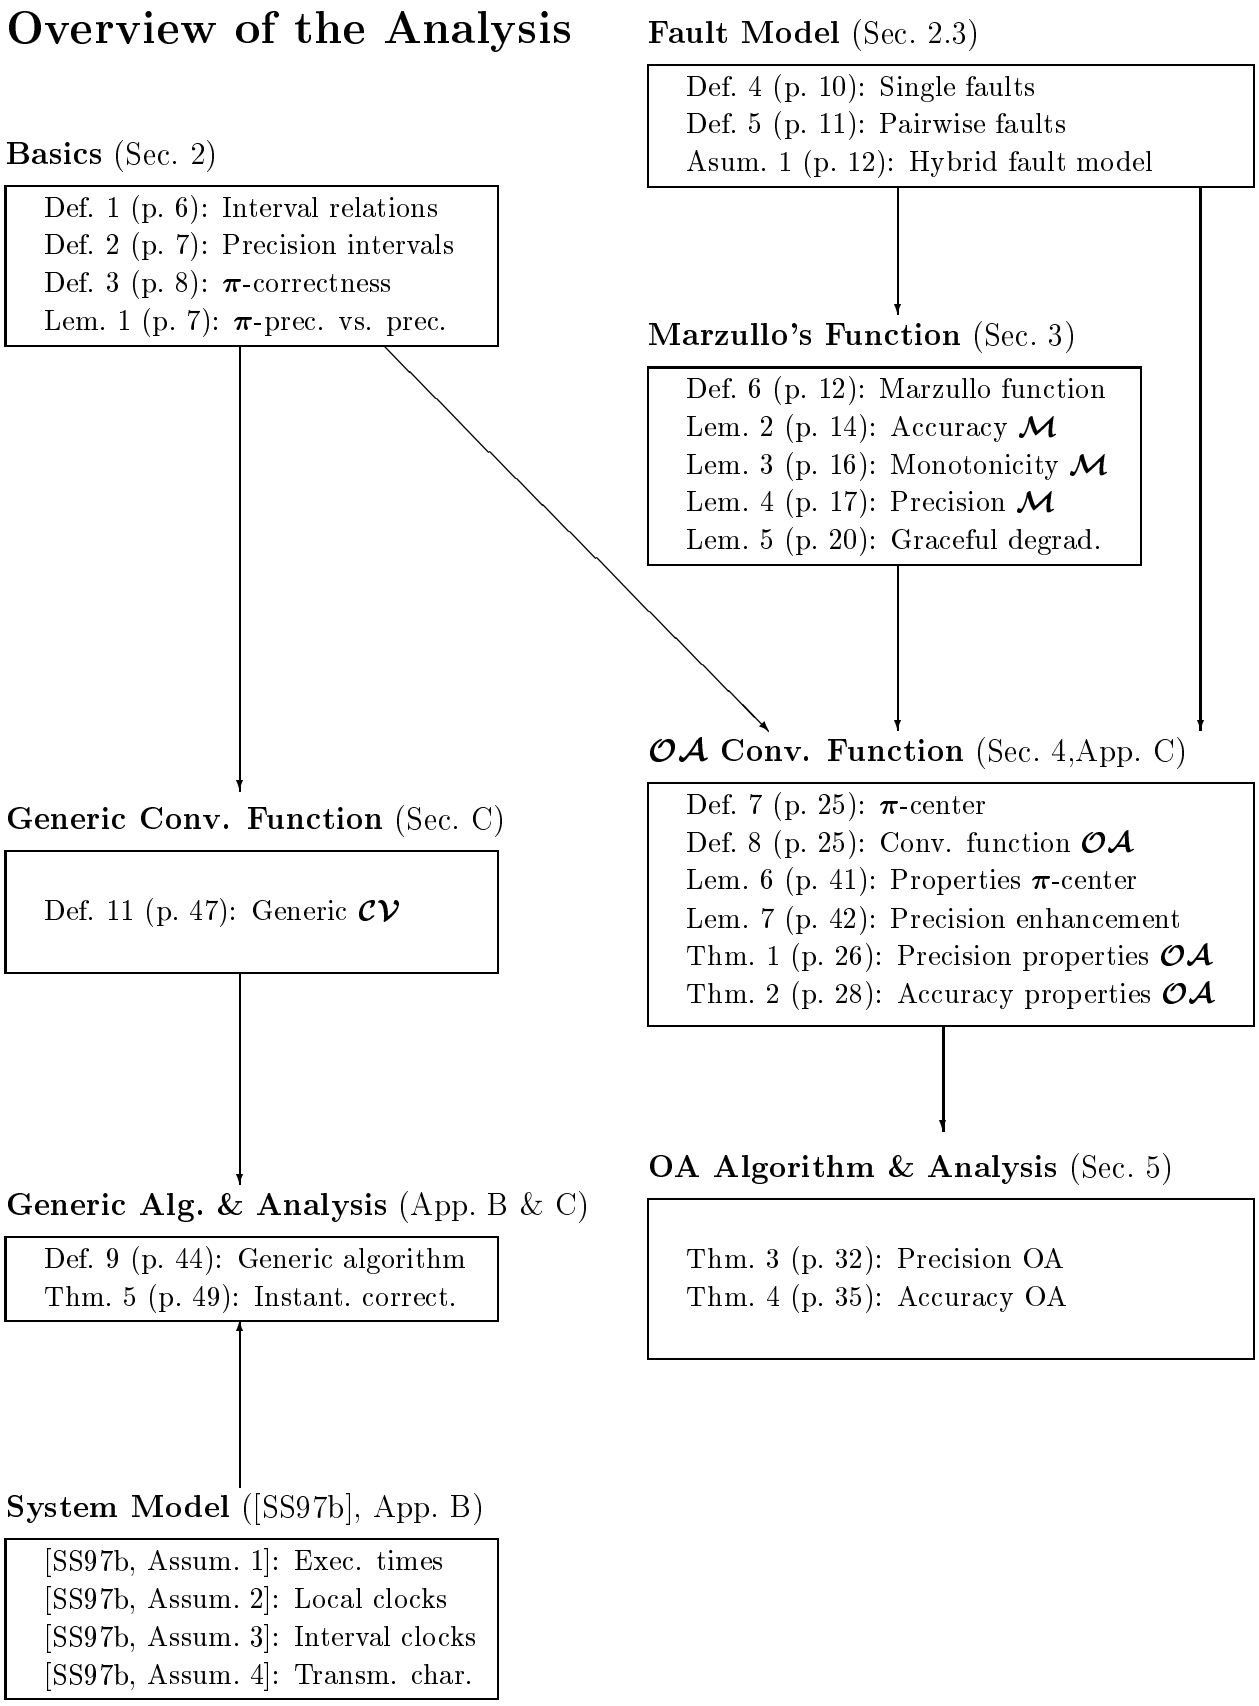

Glossary

Name	Meaning	Def.	Pg.
$\aleph^+(.)$, $\aleph^-(.)$	accuracy preservation function	Def. 11	47
$\boldsymbol{\alpha} = [-\alpha^-, \alpha^+]$	interval of accuracies of $\boldsymbol{I} = [r \pm \boldsymbol{\alpha}]$	Def. 2.2	6
$\alpha = \alpha^- + \alpha^+$	length of $\boldsymbol{\alpha}$	Sec. 2.2	6
$\boldsymbol{\alpha}_p^0 = [-\alpha_p^{0-}, \alpha_p^{0+}]$	initial accurracies at node p	Def. 9	44
$C(t), C_p(t)$	ordinary clock (node p)	Def. 2	3
$\boldsymbol{C}_p(t) = [C_p(t) \pm \boldsymbol{\alpha}_p(t)]$	local interval clock of node p	Def. 2	3
$\delta_{\max}, \delta_{\min}$	uniform bounds det. part trans. delay	Def. 9	44
Δ	transmission delay compensation	Def. 9	44
$\boldsymbol{\varepsilon}_{\max}$	max. transmission delay uncertainity	Def. 9	44
f_a	number of asymmetric resp. arbitrary faults	Asm. 1	12
f_c	number of crash faults	Thm. 4	17
f_s	number of symmetric resp. simple faults	Asm. 1	12
f_u	number of unbounded accuracy faults	Thm. 2	14
G	clock granularity	Def. 9	44
G_S	clock setting granularity	Def. 9	44
$\Gamma_{\max}, \Gamma_{\min}$	uniform logical time exec. time bounds	Def. 9	44
$\Im_-(.)$, $\Im_+(.)$	intersection enhancement function	Def. 11	47
$\boldsymbol{I}$, $\boldsymbol{A}$	(accuracy) intervals	Sec. 2.2	6
$\overline{\boldsymbol{I}}$	swapped interval $[r \mp \boldsymbol{\alpha}]$	Equ. 2	6
$\hat{\boldsymbol{I}}$	$\boldsymbol{\pi}$-precision interval associated with $\boldsymbol{I}$	Def. 2	7
$\boldsymbol{\mathcal{I}}_p$	ordered set received intervals at p	Def. 9	44
$\Lambda_{\max}$	logical time maximum broadcast latency	Def. 9	44
$\boldsymbol{\mathcal{M}}$	Marzullo's function	Def. 6	12
$\max_{i:m}\{s_i\}$	m-th largest element among $\{s_i\}$	Thm. 2	28
n	number of nodes	Sec. 2.1	3
$\Omega_{\max}$	log. time max. broadcast oper. delay	Def. 9	44
$\boldsymbol{\mathcal{OA}}$	orthogonal accuracy convergence function	Def. 8	25
P	round period	Def. 9	44
$\boldsymbol{\Phi}(.)$	precision preservation function	Def. 11	47
$\boldsymbol{\pi} = [-\pi^-, \pi^+]$	(generic) interval of precision	Def. 2	7
$\pi = \pi^- + \pi^+$	length of $\boldsymbol{\pi}$	Def. 2	7
$\boldsymbol{\pi}_0 = [-\pi_0^-, \pi_0^+]$	ideal initial precision	Thm. 5	49
$\boldsymbol{\pi}^H = [-\pi^{H-}, \pi^{H+}]$	uniform precision exchanged intervals	Thm. 5	49
$\boldsymbol{\pi}_I$	uniform precision of perceptions	Thm. 5	49
$\boldsymbol{\pi}$-accurate	correct w.r.t. t and $\boldsymbol{\pi}$-precise	Def. 2	7
$\boldsymbol{\pi}$-center	asymmetric reference point setting	Def. 7	25
$\boldsymbol{\pi}$-correct	correct w.r.t. both t and τ	Def. 3	8
$\boldsymbol{\pi}$-precise	precise interval set	Def. 2	7
$\Pi(.)$	precision enhancement function	Def. 11	47
$\mathrm{ref}(\boldsymbol{I})$	reference point of interval $\boldsymbol{I}$	Sec. 2.2	6
τ, τ^k	internal global time (of round k)	Def. 3	8