Onboard ensemble timescale realization for synchronization of future G2G-like constellations

Christian Trainotti
Institute for Communications and Navigation
German Aerospace Center
Weßling, Germany
christian.trainotti@dlr.de

Lotfi Massarweh
Mathematical Geodesy and Positioning
Delft University of Technology
Delft, Netherlands
l.massarweh@tudelft.nl

Giulia Schievano
Institute for Communications and Navigation
German Aerospace Center
Neustrelitz, Germany
giulia.schievano@dlr.de

Gabriele Giorgi
Institute for Communications and Navigation
German Aerospace Center
Weßling, Germany
gabriele.giorgi@dlr.de

***Abstract*—This work analyzes the feasibility of an onboard satellite synchronization scheme in constellations with intersatellite links. Via the link, the satellite can perform data relay and two-way time transfer, enabling a local, onboard computation of the system timescale and satellite clock offsets. This is achieved by means of a synchronization algorithm establishing an ensemble of onboard clocks. The impact on the synchronization performance caused by reduced connectivity, limited data dissemination, data gaps, and link schedule is assessed with respect to centralized solutions, using laboratory measurements of rubidium clocks. Data broadcasting and filter backtracking is proposed as a solution to reduce this impact on satellite synchronization levels.**

***Index Terms*—GNSS, clock ensemble, synchronization, ISL**

I. Introduction

Satellite synchronization is an important operational aspect of constellations for navigation and communications, enabling to timestamp and align the broadcast signals and ensuring the quality of the service provided. Current approach rely on a ground network to monitor the satellite clocks and to predict the satellite time offsets with respect to a system timescale. This method sets stringent requirements on the extension of the ground infrastructure and on the onboard clock stability. For example, satellites of current global navigation satellite systems (GNSSs) carry long-term-stable frequency references to ensure a bounded drift of the onbard timescale during the validity period of predicted corrections. Future satellite constellations will benefit from inter-satellite links (ISLs), performing ranging, time transfer, and data relay between connected satellites. These capabilities enable a radical change in the synchronization approach, shifting from a ground-based satellite time offset estimation via a post-processing chain, to an onboard real-time estimation. This relaxes the aforementioned requirements, improving service quality, system robustness, and autonomy.

979-8-3315-8115-2/26/$31.00 ©2026 IEEE

A two-way time transfer (TWTT) scheme between connected satellites over ISLs produces real-time clock offsets measurements. These observables can be distributed across the constellation with low latency. Each satellite receives all time offset measurements and processes them on board with a dedicated synchronization algorithm. This scheme results in a distributed implementation of a clock ensemble composed of all satellite clocks. Assuming perfect and real-time data distribution to all satellites and identical onboard implementations, each satellite can independently compute its time offset to the same, locally computed, system timescale. By correcting this offset, all satellites are kept synchronized to each other. The performance of this synchronization scheme was simulated for the proposed "Kepler" architecture [1], [2], showing satellite synchronization at the picosecond level [3], [4].

Complete data distribution (meaning that every satellite receives data from every other satellite) is achieved with low latency only in case of perfect satellite connectivity, meaning that two-way ISLs are continuously established between every satellite. In this case, a continuous data channel is established across the constellation. If one or more ISLs cannot be established, this assumption is not met, and the constellation becomes fragmented in several groups of interconnected satellites. While the satellite clock offset measurements and their distribution is achieved within the groups, full data distribution across the constellation is not possible without switching between different link topologies. This scenario arises for different reasons: the link might be interrupted due to technical limitations, faults, etc.; satellites in small constellations might become unreachable due to Earth masking or limited link power; satellites might be equipped with a single ISL terminal. For example, satellites of the Galileo 2nd Generation (G2G) will be equipped with a single active K-band terminal to establish ISLs in a time-division duplexing scheme [5], thus each satellite can connect to only one satellite at any given time.

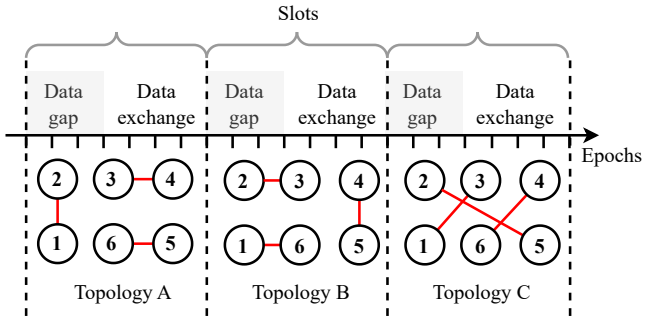

Fig. 1. Topology switching for an example of six satellites: a link topology is maintained for a number of epochs (a *slot*) between switches. At the beginning of each slot, terminal pointing causes a data gap, during which no link is available. Each circle represents a satellite, the red arrows represent the ISLs.

To improve data distribution across the constellation in such scenario, the link topology is switched after a given number of epochs. We denote this period of constant topology as a *slot*. After a switch, the current links are interrupted, and the terminals are realigned to establish a new link to a different satellite. The new links are maintained for a slot until a new topology switch takes place. The terminal pointing, acquisition, and tracking (PAT) after each switch might require some time, during which no data transfer occurs. This effectively produces data gaps further reducing the data dissemination in the constellation. A simplified representation of the temporal evolution is shown in Fig. 1 for six satellites. A dedicated scheduling algorithm (see Section III) computes the sequence of link topologies considering several constraints.

The goal of this work is to assess how the onboard synchronization scheme is affected by the limited connectivity caused by having a single ISL terminal per satellite, and, in particular, by the

- switching between topologies;
- limited data availability given by reduced connectivity;
- data gaps at the beginning of each slot.

After this assessment, a strategy is proposed to mitigate the impact of limited connectivity and topology switches on the proposed constellation synchronization scheme.

II. Data Broadcasting

In a fully-connected constellation, the measurements obtained over each link can be distributed across the constellation by a chain of data relays between satellites. The observations are received with a delay given by the cumulative flight times and hardware delays in the chain. Depending on the number of links needed and the relative stability of the onboard clocks, this delay might be negligible (each clock drift is below the targeted synchronization level during this time), so that each synchronization algorithm on board the satellites assumes the measurements are received in real-time. This assumption cannot be met in case of limited connectivity. In this case, the measurements are shared in near real-time only within the group of satellites currently connected via ISLs. Observations from other satellites are only received after a topology switch. Several slots are required to obtain full data dissemination. This is qualitatively depicted in Fig. 2 for six satellites over three slots. The first row of matrices represent the data distribution on each satellite, with each row containing the number of the satellites whose data has been received. For example, the first row in each matrix contains the data available on satellite 1: during the first slot only data from satellites 1 and 2 are available; in the second slot, the new measurement to 6 is available, and so on. Once the matrix is full, full data dissemination is obtained (irrespective of data age).

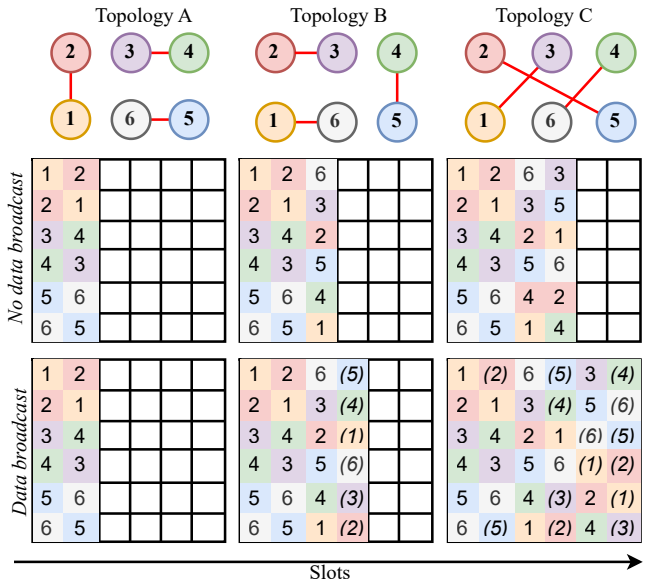

Fig. 2. Data dissemination across six satellites over three slot, only using real-time data (upper row of matrices) and with data broadcasting (lower row). The matrix rows represent the data obtained on the respective satellite. Numbers in parentheses indicate satellite data received indirectly by data broadcast. Full data dissemination is achieved when the matrix is full.

Data broadcasting consists in storing real-time measurements on board each satellite, and then relay them to the newly connected satellite after a topology switch. Although the measurements received by broadcasting refer to a previous slot, this approach improves data distribution across the constellation. Depending on the memory available onboard, data can be stored for an increasing number of slots in the past, thus effectively relaying the same information to different satellites after each slot. The second row of matrices in Fig. 2 shows the data dissemination using data broadcasting over two slots. After a topology switch, each satellite receives data from the new satellite, both in real-time and broadcast. The latter is depicted with the numbers in parentheses. In this example with six satellites, full data dissemination is achieved already after two topologies switches.

III. ISL Scheduler

The scheduler plans and organizes ISL connectivity in time and space considering the number of onboard terminals, their hardware limitations, and physical constraints, such as

Earth masking, terminal pointing angle, link distance, link persistence, etc (for more details see [6]). The primary goal of the scheduler algorithm is to achieve full data dissemination in the lowest possible number of topology switches. The data broadcasting problem is well-know in mathematics: it can be traced back to the so-called *Gossip-problem*, for which numerous algorithms and protocols are already available in the literature [7], [8], [9]. The *Gossip protocol* ensures that the information is disseminated in the smallest amount of switches, denoted as $T = \lceil \log_2 N \rceil + \mod(N, 2)$, where $\lceil \cdot \rceil$ rounds the operand to the closest smaller integer, and $\mod(\cdot)$ denotes the modulo operator. We call this period a *cycle*.

After data dissemination is first reached using the *Gossip protocol*, two different strategies have been applied to plan the the following topologies:

- strategy 1 ("*Frozen schedule*") re-starts the *Gossip protocol* as in the first slot, considering the new visibility conditions (if changed); with this strategy the data broadcasting is reached at the end of every cycle of T slots, by repeating the same sequence of topologies, hence the name "*frozen schedule*";
- strategy 2 ("*Minimum age schedule*") guarantees that full data broadcasting is reached at each slot, by determining the next topology so that the maximum information age is smaller or equal than $T - 1$.

Strategy 1 is simple and agile, and it doesn't require large memory and computation capability, therefore it is a good candidate for onboard computation; on the other hand, strategy 2 is more complex and requires more resources, but it is optimized to bound the age of information relayed between satellites.

IV. Synchronization algorithms

In this paper we compare system synchronization levels obtained by three different algorithms, namely implementation of *local Kalman filters*, a *global Kalman filter*, and a *global batch*. The local filters are the implementations on board all satellites. They compute an iterative real-time solution using only measurements available at each epoch. The measurements are either obtained locally by TWTT with the connected satellite, or they are remote measurements received by data broadcasting. This solution represents the synchronization scheme available in a real case scenario. The global solutions use all clock offset measurements obtained via ISLs, as if they were immediately distributed to a centralized location. They are used as benchmark to assess the performance of the local filters. The global solution is computed either in real-time with a Kalman filter processing the measurements epoch-by-epoch, or in form of a batch solution, using observations accumulated over a period of time in the past. The global batch is thus a post-processed solution.

A. Clock and ensemble models

All three algorithms are based on a clock ensemble model, consisting of one clock per satellite. The total number of clocks in the ensemble is denoted as N. Typically, a clock is modeled with two or three states [10]. The first state represents the phase deviation φ of the clock with respect to an ideal timescale, while the second one represents part of the frequency deviation. In this work, however, to decrease the computational burden of the batch solution, we opt to model each clock using the first state only. This is justified since the frequency deviation of rubidium clocks is negligible within the 1-minute slot length considered in this work.

The time offsets of the single clocks are collected into a single state vector for the entire ensemble $\mathbf{x}^\top = [\varphi_1, \cdots, \varphi_N]$. The general time-discrete dynamics equation of the ensemble is [11], [12]:

$$\mathbf{x}[k+1] = \mathbf{A}(\tau_0)\mathbf{x}[k] + \mathbf{w}[k] \ , \tag{1}$$

where $\mathbf{A}(\tau_0)$ is the propagation matrix, τ_0 is the discretization time between two consecutive samples, and $\mathbf{w}$ is the process noise, assumed normally distributed with covariance $\mathbf{Q}(\tau_0)$. The matrices $\mathbf{A}$ and $\mathbf{Q}(\tau_0)$ are block-diagonal, composed of the propagation and covariance matrices of the single clocks. Since a single state per clock is used in this case, the matrices of the individual clocks are scalar:

$$A_i = 1, \quad Q_i(\tau_0) = \sigma_1^2 \tau_0 \ , \tag{2}$$

where the parameter $\sigma_1 = 2.6 \times 10^{-12} \sqrt{\text{s}}$ is the white frequency noise standard deviation of the clock. Even though the matrix $\mathbf{A}$ is an identity matrix, it is maintained for the sake of generality in case of a future extension to a second state.

The link observables are the geometry-free clock offsets of the connected satellites obtained via TWTT, already corrected for relativity. The observable between satellites i and j is the time difference

$$z_{ij}[k] = \varphi_i[k] - \varphi_j[k] + v_{ij}[k] = \mathbf{H}_{ij}\mathbf{x}[k] + v_{ij}[k] \ , \tag{3}$$

where v_{ij} is zero-mean white measurement noise with standard deviation $\sigma_z = 0.1\,\text{ns}$. We assume that the TWTT scheme results in equal measurements on both sides, up to the measurement noise, $\langle z_{ij}[k] \rangle = \langle z_{ji}[k] \rangle$, where $\langle \cdot \rangle$ denotes the expectation operator.

All ISLs measurements can be combined together in a single vector $\mathbf{z}_\text{G}$, representing the observables for the global algorithms:

$$\mathbf{z}_\text{G}[k] = \mathbf{H}_\text{G}[k]\mathbf{x}[k] + \mathbf{v}_\text{G}[k] \ , \tag{4}$$

where $\mathbf{H}_\text{G}$ is the time-dependent observation matrix mapping the state vector to the global measurement vector, and $\mathbf{v}_\text{G}$ is a vector of measurement noise, zero-mean with covariance $\mathbf{R} = \sigma_z^2 \mathbf{I}$. The observation matrix changes in time according to the topology valid at a given epoch.

On the i-th satellite, only a subset of these observations are available, denoted with $\mathbf{z}_i$:

$$\mathbf{z}_i[k] = \mathbf{H}_i[k]\mathbf{x}[k] + \mathbf{v}_i[k] \ , \tag{5}$$

where the local observation matrix $\mathbf{H}_i$ must be computed accounting only for the measurements obtained locally and received by data broadcasting.

B. Global and local filters

The global and local filters estimate the ensemble state epoch-by-epoch using a Kalman filter. They have the same implementation, and both estimate the clock offset of all satellites, but use different observations. The global filter uses the global measurements $\mathbf{z}_\mathrm{G}$, while the i-th local filter uses the observations $\mathbf{z}_i$.

The filters compute the state estimate $\widehat{\mathbf{x}}$ with the recursion:

$$\widehat{\mathbf{x}}^-[k]=\mathbf{A}\widehat{\mathbf{x}}[k-1] \tag{6}$$

$$\mathbf{P}^-[k]=\mathbf{A}\mathbf{P}[k-1]\mathbf{A}^\top+\mathbf{Q} \tag{7}$$

$$\mathbf{K}[k]=\mathbf{P}^-[k]\mathbf{H}^\top[k]\left(\mathbf{H}[k]\mathbf{P}^-[k]\mathbf{H}^\top[k]+\mathbf{R}\right)^{-1} \tag{8}$$

$$\widehat{\mathbf{x}}[k]=\widehat{\mathbf{x}}^-[k]+\mathbf{K}[k]\left(\mathbf{z}[k]-\mathbf{H}\widehat{\mathbf{x}}^-[k]\right) \tag{9}$$

$$\mathbf{P}[k]=(\mathbf{I}-\mathbf{K}[k]\mathbf{H})\mathbf{P}^-[k] \tag{10}$$

where $\widehat{\mathbf{x}}^-$ is the state prediction, $\widehat{\mathbf{x}}$ the state update, $\mathbf{P}^-$ the covariance matrix associated to the predicted estimation error, $\mathbf{P}$ the covariance matrix associated to the updated estimation error, and $\mathbf{K}$ the Kalman gain. The filter is initialized as $\widehat{\mathbf{x}}[0]=\widehat{\mathbf{x}}_0$, $\widehat{\mathbf{P}}[0]=\mathbf{P}_0$. For a compact notation, the Kalman filter equations are summarized with a compact notation for global and local implementations:

$$\widehat{\mathbf{x}}_\mathrm{G},\mathbf{P}_\mathrm{G}=\mathrm{KF}\left(\widehat{\mathbf{x}}_0,\mathbf{P}_0,\mathbf{z}_\mathrm{G},\mathbf{H}_\mathrm{G}\right) \tag{11}$$

$$\widehat{\mathbf{x}}_i,\mathbf{P}_i=\mathrm{KF}\left(\widehat{\mathbf{x}}_0,\mathbf{P}_0,\mathbf{z}_i,\mathbf{H}_i\right)\ \ \forall i=1\dots N \tag{12}$$

A well-known problem in the application of Kalman filters for clock ensembling is the unbounded growth of the state covariance due to the inherent unobservability of the system [11], [13]. To avoid this, several modifications to reduce the filter covariance have been proposed in literature, however no covariance reduction has been applied in this work.

C. Batch solution

The batch solution computes clock offset estimates over a number of epochs in the past by collecting and processing the measurement over said period. Only a global implementation is presented here. In this work, we assume that the batch length is equal to the current epoch, meaning that all measurements from the first to the current epoch are used. At epoch k, the batch combines equations (1) and (4) as

$$\begin{bmatrix} 0 \\ \vdots \\ 0 \\ \hline \widehat{\mathbf{x}}_0 \\ \hline \mathbf{z}_\mathrm{G}[0] \\ \mathbf{z}_\mathrm{G}[1] \\ \vdots \\ \mathbf{z}_\mathrm{G}[k] \end{bmatrix} = \left[\begin{array}{ccccc} -\mathbf{A} & \mathbf{I} & 0 & \cdots & 0 \\ 0 & -\mathbf{A} & \mathbf{I} & \cdots & 0 \\ & & & \ddots & \\ 0 & \cdots & & -\mathbf{A} & \mathbf{I} \\ \hline \mathbf{I} & 0 & \cdots & & 0 \\ \hline \mathbf{H}_\mathrm{G}[0] & 0 & & \cdots & 0 \\ 0 & \mathbf{H}_\mathrm{G}[1] & & \cdots & 0 \\ & & & \ddots & \\ 0 & \cdots & & & \mathbf{H}_\mathrm{G}[k] \end{array}\right] \begin{bmatrix} \widehat{\mathbf{x}}_\mathrm{B}[0] \\ \widehat{\mathbf{x}}_\mathrm{B}[1] \\ \vdots \\ \widehat{\mathbf{x}}_\mathrm{B}[k] \end{bmatrix} \tag{13}$$

where $\widehat{\mathbf{x}}_\mathrm{B}$ is the vector of estimates computed with the batch solution. The expression (13) can be written as

$$\mathbf{Y}=\mathbf{\Phi}\widehat{\mathbf{X}}_\mathrm{B}\ . \tag{14}$$

The vector $\mathbf{Y}$ contains $Mk+n(k-1)+n$ elements, where M is the number of measurements in $\mathbf{z}_\mathrm{G}$ and n the number of states in $\mathbf{x}_\mathrm{G}$. The vector $\widehat{\mathbf{X}}_\mathrm{B}$ contains nk elements. The covariance matrix of $\mathbf{Y}$ is a block-diagonal matrix:

$$\mathbf{\Sigma}_{\mathbf{YY}}=\mathrm{blkdiag}\left(\mathbf{I}_{k-1}\otimes\mathbf{Q},\mathbf{P}_0,\mathbf{I}_k\otimes\mathbf{R}\right)\ , \tag{15}$$

where $\otimes$ is the Kronecker product. Equation (14) can be solved as a weighted least-squares:

$$\widehat{\mathbf{X}}_\mathrm{B}=\left(\mathbf{\Phi}^\top\mathbf{\Sigma}_{\mathbf{YY}}^{-1}\mathbf{\Phi}\right)^{-1}\mathbf{\Phi}^\top\mathbf{\Sigma}_{\mathbf{YY}}^{-1}\mathbf{Y}\ . \tag{16}$$

D. Measurement backtracking

During a time slot, the local filters on board the satellites can only process the real-time measurement obtained via the link to the connected satellite. Since only information from these two clocks is available, only the corresponding estimates are updated, while the remaining states are predicted without measurement update. Thanks to data broadcasting introduced in Section II, after a topology switch each satellite receives older measurements from the new satellite. This information can be included in the local filter to improve its estimates. First, the filter restores the estimate for the first epoch in the received data. Then, the filter recomputes the current estimate by using the whole received information. Suppose a topology switch occurs at epoch k_s. Before the switch, the i-th filter estimates the clock offsets using the real-time local measurement:

$$\begin{aligned}\widehat{\mathbf{x}}_i[k],\mathbf{P}_i[k]=\ &\mathrm{KF}(\widehat{\mathbf{x}}_i[k-1],\mathbf{P}_i[k-1],\\ &\mathbf{z}_i[k],\mathbf{H}_i[k])\quad\forall\ k<k_\mathrm{s}.\end{aligned} \tag{17}$$

At the switch, the satellite receives measurements over L epochs in the past, $\mathbf{z}_i[k_\mathrm{s}-L,\dots,k_\mathrm{s}]$, so that the procedure can be repeated starting from the local estimates at the epoch $k_\mathrm{s}-L$:

$$\begin{aligned}\widehat{\mathbf{x}}_i[k_\mathrm{s}],\mathbf{P}_i[k_\mathrm{s}]=&\mathrm{KF}\left(\widehat{\mathbf{x}}_i[k_\mathrm{s}-L],\mathbf{P}_i[k_\mathrm{s}-L],\right.\\ &\left.\mathbf{z}_i[k_\mathrm{s}-L,\dots,k_\mathrm{s}],\mathbf{H}_i[k_\mathrm{s}-L,\dots,k_\mathrm{s}]\right)\end{aligned} \tag{18}$$

After the switch, the filter estimates the clock offsets in real-time using the measurement received via the link to the new satellite. The length of the backtracking L is a multiple of the slot length, and depends on the available data storage and relay budgets. Section V-B shows the benefit of measurement backtracking on the local filters for different lengths.

E. Performance parameters

To compare the performance of the different solutions, we introduce performance parameters to associate a quantitative metric to system (de)synchronization. The corrected clocks are defined as the real onboard clocks corrected by their estimated offsets:

$$\boldsymbol{e}[k]=\mathbf{x}[k]-\widehat{\mathbf{x}}[k]\ . \tag{19}$$

The ensemble timescale is defined as the mean between all corrected clocks. Usually, the implicit ensemble mean (IEM) is defined as a weighted average of the corrected clocks accounting for each clock stability in case of mixed ensembles.

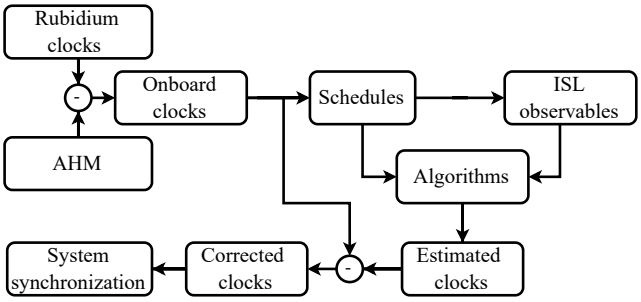

Fig. 3. Scheme of the system synchronization assessment in this paper.

In this case, the ensemble consists of the same clock type, so that the filter covariance has the same value for all clocks, and the weighted mean corresponds to the simple mean:

$$\bar{e}[k] = \frac{1}{N}\sum_{i=1}^{N} e_i[k] \ , \tag{20}$$

where e_i is the i-th element of $\mathbf{e}$.

We define the system desynchronization as the absolute error of each clock with respect to the ensemble mean:

$$\delta x_i[k] = |e_i[k] - \bar{e}[k]| \ , \tag{21}$$

while the maximum desynchronization is defined as the maximum offset between any two corrected clocks:

$$\begin{aligned}\delta x_{\max}[k] &= \max_{i\neq j}\left(|e_i[k] - e_j[k]|\right) \\ &= \max_i\left(e_i[k]\right) - \min_i\left(e_i[k]\right)\end{aligned} \tag{22}$$

These quantities have been chosen since they are representative of the contribution of the synchronization scheme to the signal-in-space error, a key quantity to assess GNSS navigation performance.

V. Assessment of System Synchronization Performance

The evaluation of the synchronization performance is conducted using the scheme shown in Fig. 3. The time evolution of the onboard clocks is obtained by measuring the time offset of three rubidium clocks with respect to an active hydrogen maser (AHM) located in DLR-KN clock laboratory. The AHM is not present in the simulated system, but it is used to obtain "true" onboard time offsets with respect to an external reference, since the AHM stability is better than that of the rubidium clocks. Only a limited number of rubidium clocks are available, thus the measurements are sliced in up to 24 pieces to simulate wider constellations. The slices are chosen as far apart as possible to reduce their temporal correlation.

The onboard clocks are combined according to the topology plan computed by the scheduling algorithm, to obtain differences between the onboard clocks. We add a zero-mean white measurement noise with $0.1\,\mathrm{ns}$ ($\sim$3 cm) standard deviation to all ISL observables, to represent the expected performance of TWTT over ISL in radio-frequency (RF) domain. For reference, analysis of the BeiDou system (BDS-3) ISLs showed time transfer performance of $0.4\,\mathrm{ns}$ root-mean-square [14]. We further assume topology slots of 60 seconds, with a variable data gap of 0, 20, and 40 seconds.

The ISL observables are the input to the synchronization algorithms. As described in Section IV, three different algorithms compute clock offset estimates, from which the corrected clocks are generated. The synchronization performance of the three algorithms is compared in terms of the parameters of Section IV-E.

A. Global vs local solutions

To evaluate the effect of reduced connectivity on the onboard clock offset estimates, an example with six satellites is used to compare the local and global solutions, with a data gap of $20\,\mathrm{s}$ and no backtracking. Fig. 4 compares the corrected clocks computed by the three algorithms. It can be noted that the local filters suffer from the limited onboard data availability, resulting in larger desynchronization. This effect is reduced in the global filter, which however is impacted by the pairwise connectivity. The batch solution is not impacted, since the effects of missing measurements during a slot are compensated by past and future observations. It can be noted that the mean of the corrected clocks is the same for the three solutions. The same effects can be seen in Fig. 5, showing the maximum desynchronization of the three solutions.

B. Local filters with backtracking

The desynchronization achieved by the local filters has been evaluated for a set of $N = 24$ satellites, using backtracking over an increasing number of slots, and the two scheduler approaches of Section III. Fig. 6 shows the time evolution of $\delta x_{\max}$ and empirical cumulative distribution function (eCDF) of δx using the *frozen* scheduling strategy, data gaps of $20\,\mathrm{s}$, and an increasing length of backtracking, for the local and global solutions. Backtracking reduces the desynchronization between onboard filters, and increasing backtracking length (number of slots over which the filters reprocess the estimates) drives the system synchronization to the level achieved by the global filter solution. In this configuration, backtracking over 5 slots is enough to achieve full data dissemination, so no additional benefit is obtained by backtracking over a longer period.

The effect of different data gap lengths and of the scheduler strategy is plotted in Fig. 7, showing the empirical CDF of the desynchronization for the *frozen* schedule (upper row), and *minimum age* schedule (bottom row). The *minimum age* scheduler improves the synchronization compared to the *frozen* scheduler, since it guarantees full data distribution at every slot. By comparing the three columns, it appears that longer data gaps do not significantly degrade the synchronization, compared to the effect of the scheduler. This is thanks to the stability of the rubidium clocks, which do not accumulate large time offsets during the data gaps. Nevertheless, for the global filter and when backtracking over several slots, the effect of

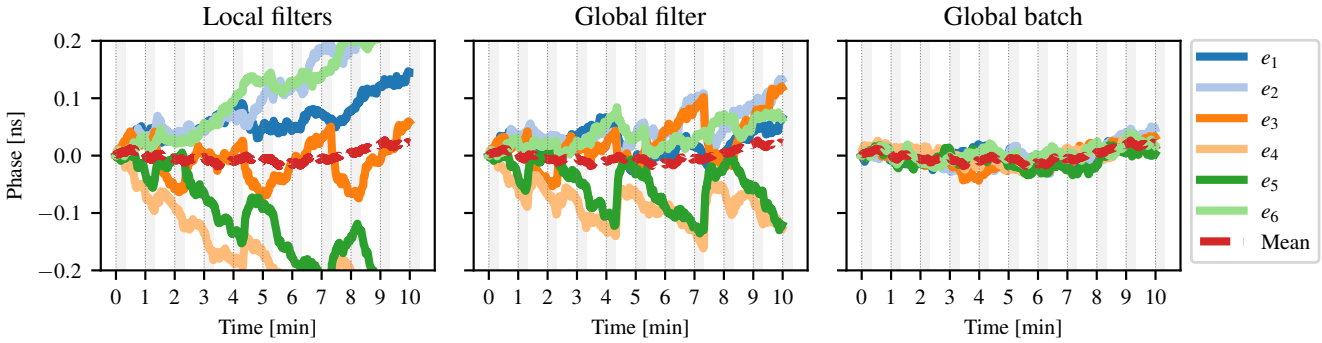

Fig. 4. Comparison of the corrected clocks computed by local filters, global filter and batch solution, for a scenario of six satellites. The vertical dashed lines represent topology switches, the gray areas are data gaps.

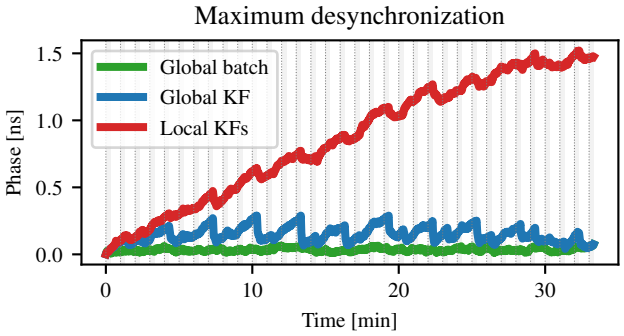

Fig. 5. Comparison of the maximum desynchronization of the three algorithms, for a scenario of six satellites (KF = Kalman filter).

data gaps is still visible, as depicted in Fig. 8, showing the 95-percentile of the system desynchronization for different data gap lengths, backtracking lengths, and the two scheduling strategies. For the *minimum age* scheduler, the values increase of about 30 ps comparing data gaps of 40 s and no data gaps, independently of backtracking. For the global filter and local filters with full backtrack (BT=5), this corresponds to an increase of about 40%.

VI. Conclusions

This paper described an onboard satellite synchronization scheme using a local Kalman filter processing time offset measurements obtained via TWTT over intersatellite links. Its performance was assessed in different configurations and compared to benchmark global (centralized) solutions, using measurements of real rubidium clocks. The synchronization achieved by local onboard filters is degraded by reduced connectivity and limited data distribution, compared to the global filter and batch solution. To counteract this effect, data broadcasting and filter backtracking can be used to improve the data availability across the constellation, and to update the local estimates using relayed data. Using this approach, the synchronization level of the local filters approaches that of the global filter solution for a backtrack length of five slots.

The scheduling strategy affects the synchronization algorithm, since it impacts data distribution and average age of the relayed data. The *minimum age* strategy, minimizing information age and guaranteeing complete data dissemination at every slot, improves the synchronization compared to the *frozen* strategy. This benefit is more significant when no backtracking takes place. When the backtracking is fully exploited, the difference between the performance for the two strategies is comparable to that for the global filter.

Data gaps at the beginning of each slot seem to have an impact on the system synchronization which is independent of backtracking, given by the clock divergence during the data gaps. This effect, impacting both global and local solutions alike, is masked by that of limited data availability, but it becomes apparent when backtracking over several slots. The usage of onboard clock with lower short-term stability might reduce the desynchronization caused by the presence of data gaps.

Future work shall address several open points, such as the effect of limited data availability in case of clock faults, or measurement biases. The impact of different relative magnitude of clock process noise and measurement noise must also be assessed. Finally, a real-time implementation of this scheme must consider link budget and onboard resources, potentially limiting the amount of data that can be broadcast.

References

[1] C. Günther, "Kepler – Satellite Navigation without Clocks and Ground Infrastructure," in *Proceedings of the 31st ION GNSS+ 2018*, ser. GNSS 2018, Institute of Navigation, Oct. 2018, pp. 849–856. DOI: 10.33012/2018.15997

[2] C. Günther, "Kepler – Satellite Navigation System Description and Validation," in *9th Workshop on Satellite Navigation Technologies (NAVITEC 2018)*, Noordwijk, The Netherlands, Dec. 2018.

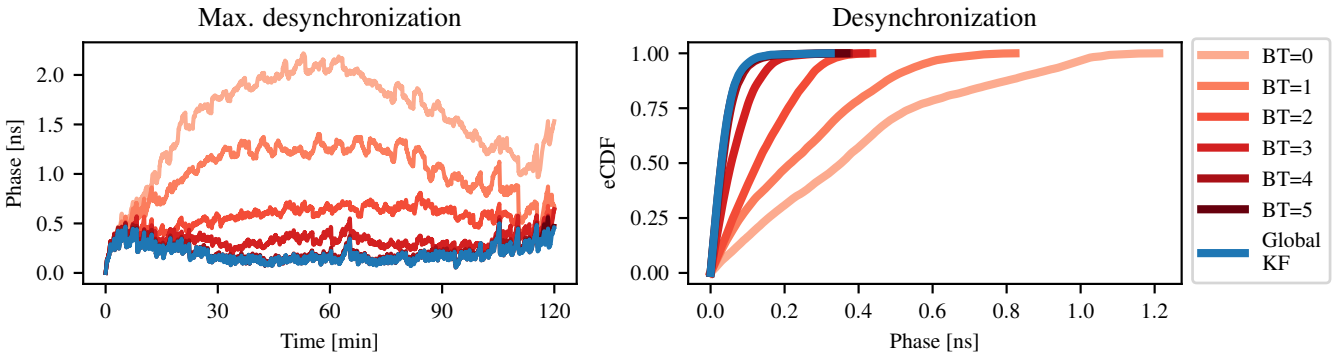

Fig. 6. Effect of backtracking on local filters compared to the global filter (KF), for a scenario of 24 satellites. Left: time evolution of the maximum desynchronization. Right: empirical CDF of the desynchronization. BT=i denotes backtracking over i slots in the past. The scenario assumes the *frozen* scheduler strategy and data gaps of 20 s.

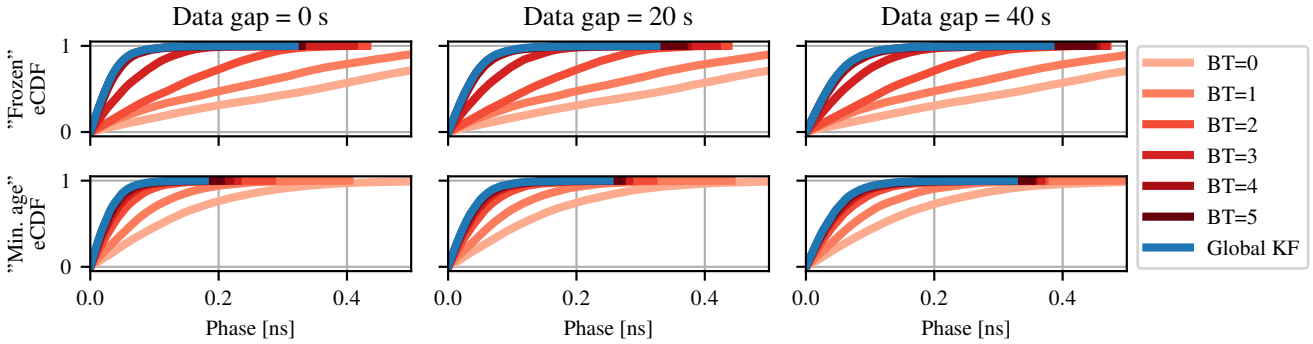

Fig. 7. Effect of backtracking, data gap length, and scheduler strategy on local filters compared to the global filter (KF), for a scenario of 24 satellites. The upper row shows the results using the *frozen* scheduler, the bottom row for the *minimum age* strategy.

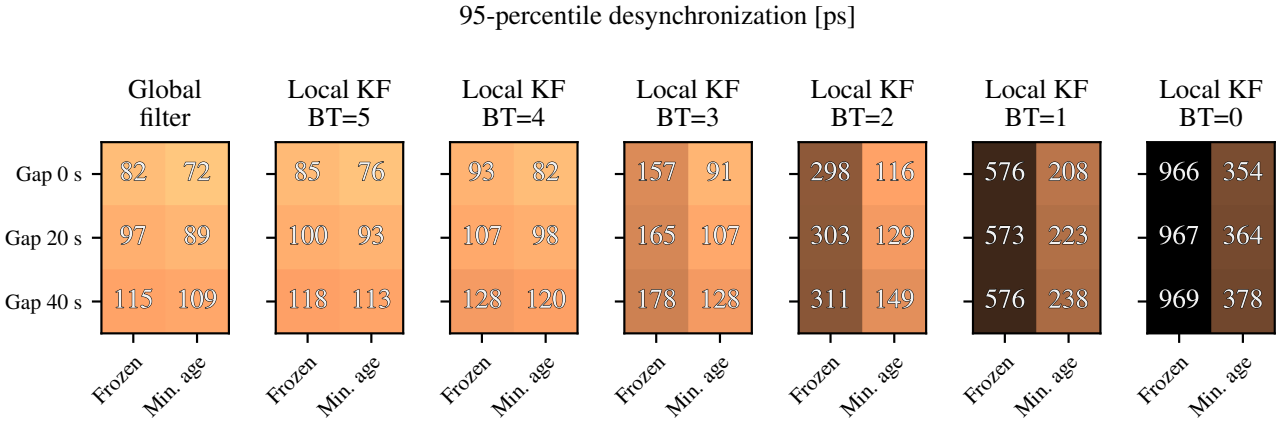

Fig. 8. 95-percentile of the desynchronization of local filters with increasing backtracking length, increasing data gap lengths, for the two scheduling strategies. Values in picoseconds.

[3] C. Trainotti, M. Dassié, G. Giorgi, A. Khodabandeh, and C. Günther, "Autonomous Satellite System Synchronization Schemes via Optical Two-Way Time Transfer and Distributed Composite Clock," in *Proceedings of the 35th ION GNSS+ 2022*, Denver, Colorado: Institute of Navigation, Sep. 2022, pp. 3646–3661. DOI: 10.33012/2022.18296

[4] C. Trainotti, M. Dassié, G. Giorgi, A. Khodabandeh, and C. Günther, "Autonomous Synchronization of Satellite Constellations via Optical Inter-Satellite Links," in *2023 Joint Conference of the EFTF and IFCS*, IEEE, May 2023, pp. 1–6. DOI: 10.1109/eftf/ifcs57587.2023.10272176

[5] Galileo Programme Department, *Galileo 2nd Generation Space-to-Space ICD*, GAL-ICD-ESA-SYST-X-4020, version 3.5, Mar. 2022.

[6] G. Schievano, G. Giorgi, and G. Michalak, "An analysis of inter-satellite link topologies in future gnss constellations: Operational constraints and figures of merit," in *Proceedings of the 2025 ION ITM*, 2025, pp. 604–614.

[7] B. Baker, "Gossips and telephones," *Discrete Mathematics*, vol. 2, pp. 191–193, 1972.

[8] D. W. Krumme, G. Cybenko, and K. Venkataraman, "Gossiping in minimal time," *SIAM Journal on Computing*, vol. 21, no. 1, pp. 111–139, 1992.

[9] M. C. Cooper, A. Herzig, F. Maffre, F. Maris, and P. Régnier, "The epistemic gossip problem," *Discrete Mathematics*, vol. 342, no. 3, pp. 654–663, 2019. DOI: j.disc.2018.10.041

[10] C. Zucca and P. Tavella, "The clock model and its relationship with the Allan and related variances," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 52, no. 2, pp. 289–296, Feb. 2005, ISSN: 0885-3010. DOI: 10.1109/tuffc.2005.1406554

[11] C. A. Greenhall, "A Kalman filter clock ensemble algorithm that admits measurement noise," *Metrologia*, vol. 43, no. 4, S311–S321, Aug. 2006, ISSN: 1681-7575. DOI: 10.1088/0026-1394/43/4/s19

[12] C. Trainotti, T. D. Schmidt, and J. Furthner, "Simulating the Realization of a Mixed Clock Ensemble," in *2019 Joint Conference of the IEEE IFC and EFTF*, IEEE, Apr. 2019. DOI: 10.1109/fcs.2019.8856103

[13] K. R. Brown Jr., "The theory of the GPS composite clock," in *Proceedings of the 4th ION ITM*, vol. 4, Albuquerque, NM, Sep. 1991, pp. 223–241.

[14] Y. Guo et al., "Time synchronization between satellites via inter-satellite link observations of bds-3 constellation: Method, experiment and analysis," *Measurement*, vol. 224, p. 113 855, Jan. 2024, ISSN: 0263-2241. DOI: 10.1016/j.measurement.2023.113855