Forensic Reconstruction of the 2016–2017 Galileo Clock Crisis via Temporal Principal Component Analysis

Marvin Hansen[1]

[1]The Center for Dynamic Causality, `marvin.hansen@causalcenter.com`

22 January 2026

Abstract

We reconstruct the 2016–2017 Galileo onboard-clock crisis from publicly available IGS final clock products, using a sliding-window principal component analysis (PCA) of the seven-satellite swarm {E08, E09, E11, E12, E14, E18, E19}. The first principal component (PC1) tracks the common-mode time-scale signal carried by the constellation; the second component (PC2) tracks differential drift, which spikes when one or more satellites diverge from the fleet. Year-over-year integration captures the macroscopic crisis: PC1 drops from 81% of total variance in 2016 to 61% in 2017, recovering partially to 67% in 2018. A 7-day sliding window resolves the event-level structure, and pinpoints the onset to E11 on 2016 November 30 at 11:05 UTC, with a PC2 share of 27.2%. This date is consistent with the European Space Agency's later disclosed timeline; the 11:05 UTC time and the satellite identification go beyond the public record. A second peak appears on 2017 September 30 at 19:00 UTC (E11, 23.5%), which, to our knowledge, has not been publicly characterized. We further detect, in the 2018 segment, a multi-week September anomaly cluster involving E08 and E12 with PC2 shares above 30%; the start date matches the 5 September 2018 anomaly in SVN E206 reported by Galluzzo et al. (2021), independently validating the method. The pipeline relies only on standard IGS products and a single hyperparameter (window length); we discuss its potential as a real-time GNSS integrity monitor.

Keywords: Galileo; GNSS integrity; clock anomaly; principal component analysis; passive hydrogen maser; time-series forensics

1 Introduction

The Galileo constellation experienced an extended period of onboard atomic-clock instability between 2016 and 2018, leading to public statements by the European Space Agency (ESA) on 18 January 2017 (ESA, 2017) and to a subsequent independent assessment by Galluzzo et al. (2021). The crisis affected both the rubidium atomic frequency standards (RAFS) and the passive hydrogen masers (PHMs) on multiple Full Operational Capability and In-Orbit Validation satellites, and resulted in a documented set of clock switchovers across the fleet. The exact onset, evolution, and per-satellite role have only been partially disclosed.

We treat the IGS final clock product as a multi-channel time series and ask a forensic question: given the seven-satellite Galileo swarm visible in the IGS combined product, when did the fleet first lose coherence, and which satellites drove that loss? The answer is constructed without ESA telemetry. The only inputs are publicly available IGS clock corrections at 30-second cadence and the IGS final orbits. The method is principal component analysis applied to the swarm residuals after removal of a constellation-mean drift, in two cadences: one annual integration that gives the macroscopic crisis envelope, and one sliding-window scan that resolves individual events.

2 Method

2.1 Signal model

For each satellite $s \in \mathcal{S} = \{E08, E09, E11, E12, E14, E18, E19\}$, we form the clock residual

$$r_s(t) \;=\; c_s(t) - \frac{1}{|\mathcal{S}|} \sum_{s' \in \mathcal{S}} c_{s'}(t), \tag{1}$$

where $c_s(t)$ is the IGS final clock correction for satellite s at epoch t, after removal of a per-satellite linear drift (the trend rate is removed by a robust Theil–Sen fit on the previous 24 hours, and the constant is subtracted). The constellation-mean term in Eq. (1) suppresses the global time-scale drift; what remains is the part of each satellite's clock that does not look like the fleet.

In matrix form, stacking residuals across the swarm gives a data matrix $\mathbf{R}(W) \in \mathbb{R}^{N \times |\mathcal{S}|}$ over a window W of N epochs. Centring each column to zero mean and applying the singular value decomposition,

$$\mathbf{R}(W) \;=\; \mathbf{U}\,\mathbf{\Sigma}\,\mathbf{V}^{\top}, \tag{2}$$

yields principal components $\mathbf{U}_k$, loadings $\mathbf{V}_k$, and singular values σ_k. The fraction of variance captured by component k is

$$f_k(W) \;=\; \frac{\sigma_k^2}{\sum_j \sigma_j^2}. \tag{3}$$

2.2 Two cadences of analysis

We use two analysis cadences against the same data.

The *annual* cadence uses $W =$ one calendar year. The first component f_1 measures how much of the total residual variance lives in a single fleet-wide mode. A healthy fleet has nearly all of its residual variance in PC1; differential modes are small. A fleet undergoing simultaneous individual failures redistributes variance into PC2, PC3, and so on, and f_1 falls.

The *event* cadence uses $W = 7$ days, stepped by 1 day. For each window we report $f_2(W)$, the share of variance carried by the differential mode. Spikes in f_2 identify epochs at which one or more satellites depart from the fleet. The associated loading vector $\mathbf{V}_2$ identifies which satellites are responsible: a single dominant entry in $\mathbf{V}_2$ implicates a single satellite, and a multi-entry profile implicates multiple satellites simultaneously. We use a 75% relative-loading threshold to attribute an event to the dominant entry, with multi-attribution when the second entry exceeds the same threshold.

2.3 Anomaly thresholds

We do not assume Gaussian residuals. We instead estimate the null distribution of $f_2(W)$ from the first 40% of the analysis period (2016 Q1–Q2), which precedes the documented onset, and define three tiers from the empirical quantiles of that null:

- Elevated: $f_2(W) >$ 90th percentile.
- Significant: $f_2(W) >$ 99th percentile.
- Extreme: $f_2(W) >$ 99.9th percentile.

The 99.9th-percentile threshold for the present dataset corresponds to $f_2(W) \approx 14\%$. Any window above that threshold is flagged as an event; consecutive flagged windows are merged.

2.4 Data

The analysis uses IGS Multi-GNSS Experiment (MGEX) final clock corrections at 30-second cadence and MGEX final orbits at 15-minute cadence, distributed by NASA's Crustal Dynamics Data Information System (CDDIS)[1] (CDDIS, 2026; Noll, 2010), for the seven-satellite swarm $\mathcal{S}$, over calendar years 2016–2018. We exclude GPS Weeks 1930–1937 (Jan 1 to Feb 25, 2017) due to the IGS08→IGS14 reference-frame transition, which introduces station-coordinate discontinuities into the orbit products and contaminates the differential clock signal. No satellite-internal telemetry is used. The processed swarm-residual dataset and the analysis code that generates the tables in this paper are openly available (Hansen, 2026a,b).

3 Annual structure: a single-number signature of the crisis

Table 1: Annual PCA decomposition of the seven-satellite Galileo swarm residuals.

Year	f_1 (PC1)	f_2 (PC2)	$f_1 - f_2$	Interpretation
2016	81.15%	18.85%	62.30%	Coherent fleet
2017	61.28%	38.72%	22.56%	20-point coherence loss
2018	66.85%	33.15%	33.70%	Partial recovery

Table 1 gives the macroscopic structure. The PC1 share drops by roughly 20 percentage points between 2016 and 2017, with PC2 absorbing the displaced variance. The 2018 row shows a partial return toward the 2016 baseline, consistent with the publicly reported clock-switchover programme and with new operational protocols introduced during 2017–2018. The single number f_1, computed across an entire year, is sufficient to flag 2017 as anomalous against the surrounding years; no per-event analysis is needed for that conclusion.

A complementary view comes from the raw spectral power. Computing the variance of the constellation-mean drift (the common mode) and of the residuals (the differential mode) separately gives

Table 2: Common-mode and differential-mode power, integrated over the calendar year.

Year	Common-mode power	Differential power	Ratio
2016	5.28×10^{11}	8.67×10^{10}	6.1
2017	1.20×10^{10}	2.19×10^{8}	55.0
2018	7.31×10^{11}	3.38×10^{8}	2162.0

The 2017 row shows the common-mode power dropping by more than an order of magnitude relative to 2016 and 2018. This is the signature of an aggressive constellation-mean correction applied at the IGS analysis-centre level: the IGS combined product is computed as a weighted average across analysis centres, and during 2017 the time-scale-stabilization step appears to have absorbed much of the affected satellites' drift into the common mode itself. We discuss the implications of this in Section 5.1.

---

[1] IGS MGEX final products at CDDIS: `https://www.earthdata.nasa.gov/data/catalog/cddis-gnss-igs-mgex-prod-1`

4 Event-level reconstruction

4.1 The November 30, 2016 onset

Stepping a 7-day window in 1-day increments through 2016 Q4 produces the timeline in Table 3. The first window with f_2 above the 99.9th-percentile threshold has its right edge at 2016 November 30 at 11:05 UTC, with $f_2 = 27.1\%$ and a dominant loading on E11 (with secondary loading on E19). Inspection of the per-satellite residuals during this window confirms that E11 carried the largest absolute departure from the constellation mean.

Table 3: Pre-crisis 7-day windows in 2016 Q4. Each entry is the right-edge timestamp of the window. The Nov 30 entry crosses the extreme threshold.

Date (UTC)	Dominant sat	Co-loading	f_2	Tier
2016-11-28 11:05	E11	—	15.1%	Significant
2016-11-29 11:05	E19	—	22.9%	Significant
2016-11-30 11:05	E11	E19	**27.1%**	**Extreme**
2016-12-01 15:00	E11	—	17.1%	Elevated

The European Space Agency's public statement on 18 January 2017 (ESA, 2017) attributed the suspicions to “a number of clocks” and dated the first internal observations to “some weeks earlier.” The Nov 30 11:05 UTC timestamp is consistent with that range and goes beyond it in two ways: it gives an exact UTC instant, and it identifies E11 as the dominant satellite at the onset (with E19 as a co-loading partner). Seven weeks separate the Nov 30 timestamp from the ESA announcement.

4.2 The September 30, 2017 secondary peak

A second extreme event appears at 2017 September 30 at 19:00 UTC, with E11 as the dominant satellite and $f_2 = 23.5\%$. The September 2017 cluster more broadly is summarized in Table 4.

Table 4: Extreme and significant events in September 2017.

Date (UTC)	Dominant sat	Co-loading	f_2	Tier
2017-09-13 13:15	E18	E19, E08	20.4%	Significant
2017-09-14 13:15	E12	—	17.8%	Significant
2017-09-17 13:15	E11	E19	21.8%	Significant
2017-09-30 19:00	E11	—	**23.5%**	**Extreme**

The September 30 event is a singleton, dominated by E11 alone, with no significant co-loading. Public ESA reporting through this date does not, to our knowledge, characterize the event. Its f_2 share is 4 percentage points below the November 30, 2016 onset, but the dominant-loading concentration on a single satellite is sharper.

4.3 The September 2018 cluster

The 7-day scan over 2018 detects a sustained multi-week cluster of extreme events centred on September 2018, summarized in Table 5.

The 2018-09-05 event aligns with the documented anomaly in SVN E206 reported in Galluzzo et al. (2021), Table 1, which records the start at 02:20 UTC and a duration of approximately 10 minutes. Our 1-day cadence flags the event on the same calendar day; the 09:30 UTC right-edge

Table 5: Extreme events in September 2018, $f_2 > 30\%$.

Date (UTC)	Dominant sat	Co-loading	f_2
2018-09-05 09:30	E08	—	39.6%
2018-09-06 09:30	E12	—	32.4%
2018-09-08 10:10	E12	—	36.9%
2018-09-15 10:10	E12	—	31.7%
2018-09-18 10:10	E12	—	35.8%
2018-09-30 10:45	E12	E08, E09	42.0%

of the detection window is the first 7-day window in which the anomaly is fully contained. The independent literature match validates the pipeline at the event level.

4.4 Per-satellite roles

Aggregating across all detected events, the seven satellites take on distinguishable roles. Table 6 reports the count of extreme events attributed to each satellite, and a one-line characterisation.

Table 6: Per-satellite event count and role assignment, 2016–2018.

Satellite	Extreme events	Role
E11	5	Sentinel: acute, high-magnitude spikes
E19	8	Chronic: sustained low-grade instability
E12	9	Early warning (2016) and 2018 dominant
E08	4	2017 onset (Mar 2017) and 2018 cluster
E14	1	Affected only at 2017 peak
E18	2	Mid-2017 escalation
E09	2	Stable reference, perturbed at peak

E11 acts as a high-amplitude indicator: its events coincide with the two crisis maxima. E19 carries the longest tail, with the largest event count, and disappears from the top components in 2018, consistent with the publicly known clock switchover from PHM to RAFS (Wang et al., 2020). E12 dominates the 2018 cluster, suggesting a distinct failure mode from the 2016–2017 onset.

5 Discussion

5.1 IGS post-processing artifacts

A complication appears in the eccentric satellites E14 and E18 specifically. Their 7-day-window cross-correlations with the constellation-mean residual are negative, often near -1, throughout 2017–2018. A naive reading would attribute this to "anti-correlated drift," but the physical mechanism is different.

The IGS combined clock product enforces an internal time-scale constraint: the weighted mean of all clocks across all analysis centres is held to a long-term zero. When the bulk of the constellation drifts in one direction, the constraint pulls the mean back, and the residuals of the satellites that did not drift are pushed in the opposite direction by exactly the amount needed to satisfy the constraint. The eccentric satellites are particularly susceptible to this effect, because their orbital geometry produces small but persistent gravitational-potential-induced rate offsets that the analysis-centre fitting does not always absorb cleanly. The result is an artificial -1 correlation between the eccentric satellites' residuals and the common-mode correction itself.

Two consequences follow. First, any PCA on a constellation that includes eccentric satellites must distinguish between physical anti-correlation (multiple physical clocks drifting in opposite directions) and correction-induced anti-correlation (a single time-scale constraint applied across the fleet). The 2017 PC1/PC2 split discussed above is robust against this distinction because the affected satellites in 2017 are not the eccentric ones; they are E11, E12, E19, and E08, all on circular orbits. Second, the artifact is detectable: a clean physical anti-correlation has a continuous distribution of correlation values, while the correction-induced one clusters near ± 1 with a structural cause. Correctly identifying it tells you about the IGS post-processing pipeline rather than about clock physics.

5.2 Comparison to ESA's public timeline

Table 7 compares the ESA public record against the present detection.

Table 7: Comparison of ESA public statements and the present analysis.

ESA public	This work	Match
Nov 2016 first suspicions	2016-11-30 11:05 UTC, E11 (27.1%)	Yes
Jan 18, 2017 announcement	7 weeks after detection	Plausible delay
Ongoing investigation 2017	2017-09-30 19:00 UTC, E11 (23.5%)	Not in public record
Sep 5, 2018 anomaly (Galluzzo et al.)	2018-09-05 09:30 UTC, E08 (39.6%)	Yes

The November 2016 onset and the September 2018 event are independent confirmations of the public and literature record. The September 2017 secondary peak is, as far as we have been able to determine, not characterized in any of the available ESA statements or published analyses. We document it here as a forensic finding rather than as an exclusive claim; the public record may simply be incomplete.

5.3 Method limitations

The pipeline depends on a clean common-mode estimate, which is provided by the IGS analysis-centre processing. In segments where the IGS combined product itself is unreliable (such as the IGS08→IGS14 transition window, GPS Weeks 1930–1937), the residuals carry processing noise rather than satellite-clock noise, and the PCA identifies that processing noise as if it were a physical event. The exclusion of those weeks is therefore necessary, not optional.

The 7-day window length is a tradeoff. Shorter windows (1–2 days) give better temporal resolution at the cost of larger statistical fluctuations and more false positives. Longer windows (14–30 days) integrate over multiple events. We selected 7 days as a compromise that resolves single-day events while keeping the empirical false-positive rate near 0.1%.

Single-satellite attribution is limited when the dominant loading is split between two or more satellites with similar PC2 weight. In those cases the algorithm reports both (multi-attribution), at the cost of leaving the precise causal sequence between the two underspecified.

5.4 Real-time integrity monitoring

The same algorithm runs at near real-time cadence with publicly available data products. The IGS rapid clock product is available within 17 hours of observation; the ultra-rapid product is available with a 3–9 hour lag and is sufficient for the residual-PCA pipeline at the cost of slightly higher noise. A real-time variant would replace the per-window full SVD with an incremental SVD update, yielding $f_2(W)$ as a continuous diagnostic rather than a per-window summary. The 99.9th-percentile threshold from the empirical null then provides a fleet-level integrity flag with a structurally bounded false-positive rate. We do not pursue this implementation in the present work; we note it as a direct extension.

6 Conclusions

We reconstruct the 2016–2017 Galileo onboard-clock crisis from publicly available IGS clock products using sliding-window principal component analysis on a seven-satellite swarm. The annual decomposition shows a 20 percentage-point drop in PC1 (common-mode) variance between 2016 and 2017. The 7-day event cadence pinpoints the onset to E11 on 2016 November 30 at 11:05 UTC, consistent with ESA's publicly disclosed timeline. A second extreme event, on 2017 September 30 at 19:00 UTC and dominated by E11 at 23.5%, is documented here as a forensic finding that is not present in the public record. The September 2018 event detected on 5 September matches the independent characterisation by Galluzzo et al. (2021), which validates the method against literature. The IGS post-processing pipeline introduces a detectable correction artifact on the eccentric satellites E14 and E18, which is itself a useful forensic signal once correctly identified.

Acknowledgments

This work uses the IGS final clock and orbit products, made available by the International GNSS Service. The Galileo telemetry for affected satellites is the property of the European Space Agency; this analysis uses no internal telemetry and reaches its conclusions from publicly available data only.

Statement on Responsible AI Usage

The Rust analysis pipeline that supports this paper was developed with the assistance of generally available AI coding agents, under human review at every step. AI tools were also used for proof-reading and editorial polishing of the manuscript text. The scientific design, the methodology, the data-quality decisions, the interpretation of the results, and final responsibility for the content rest with the corresponding author.

Data and code availability

The pipeline is implemented in Rust and operates on standard IGS MGEX `.clk` and `.sp3` products. The full source tree, including the exact configuration used to reproduce the tables in this paper, is published as open source on GitHub (Hansen, 2026b). The processed dataset, including the per-window PCA outputs, the event catalogue, and the GPS-week exclusion lists, is archived on Zenodo (Hansen, 2026a). The upstream IGS MGEX final clock and orbit products are distributed by NASA's Crustal Dynamics Data Information System (CDDIS) (CDDIS, 2026; Noll, 2010).

References

NASA EOSDIS Crustal Dynamics Data Information System (CDDIS), "GNSS IGS MGEX Final Products," NASA Earth Observing System Data and Information System, Greenbelt, MD, 2026. `https://www.earthdata.nasa.gov/data/catalog/cddis-gnss-igs-mgex-prod-1`.

C. E. Noll, "The Crustal Dynamics Data Information System: A Resource to Support Scientific Analysis Using Space Geodesy," *Advances in Space Research*, vol. 45(12), pp. 1421–1440, 2010. doi:10.1016/j.asr.2010.01.018.

M. Hansen, "Canonical dataset for chronometric analysis," [Data set], Zenodo, 2026. doi:10.5281/zenodo.20020236.

M. Hansen, "Chronodynamics: source code for chronometric experiments," [Source code], GitHub, 2026. `https://github.com/causalcenter/chronodynamics`.

European Space Agency, "Galileo Clock Anomalies under Investigation," ESA Statement, 18 January 2017. Available: `https://www.esa.int/Applications/Satellite_navigation/Galileo_clock_anomalies_under_investigation`.

G. Galluzzo et al., "Galileo Broadcast Ephemeris and Clock Errors Analysis: 1 January 2017 to 31 July 2020," *Sensors*, vol. 20(23), 6832, 2021. `https://doi.org/10.3390/s20236832`

L. Wang et al., "Detection and Analysis of Galileo Signal in Space Anomalies from 2017 to 2018," *IET Radar, Sonar & Navigation*, vol. 14, 2020. `https://doi.org/10.1049/iet-rsn.2019.0468`

F. Droz et al., "Space Passive Hydrogen Maser: Performances and Lifetime Data," in *Proc. Joint IEEE Frequency Control Symposium and EFTF*, 2009.

G. Petit and B. Luzum (eds.), *IERS Conventions (2010)*, IERS Technical Note 36, Bundesamt für Kartographie und Geodäsie, Frankfurt am Main, 2010.