

THE AUTONOMOUS DETECTION OF CLOCK PROBLEMS IN SATELLITE TIMEKEEPING SYSTEMS

Y. C. Chan, J. C. Camparo, and R. P. Frueholz
Mail Stop: M2-253, Electronics Technology Center,
The Aerospace Corporation, P.O. Box 92957, Los Angeles, CA 90009, USA

Abstract

Military satellite communications (milsatcom) systems require precise timekeeping in order to take advantage of spread-spectrum communication techniques. Though the level of precise timekeeping in milsatcom is typically not as stringent as that in satellite navigation, milsatcom nevertheless poses its own unique timekeeping problems. Specifically, milsatcom timekeeping must be robust, with the ability to autonomously detect and correct timekeeping problems during protracted periods when the ground control station is either not available or burdened with other pressing tasks. Here, we consider the ability of three different space-segment timekeeping subsystems to autonomously detect the failure of a single satellite clock and to then take appropriate action to remedy the situation. These systems include a Master/Slave system, an Ensembling system, and a Kalman-Filter system. Employing Monte Carlo simulation, we consider four types of "soft" clock failure: a time-jump failure, a frequency-jump failure, a failure arising from a sudden change in the clock's frequency aging rate, and a failure arising from an abrupt increase in the clock's random-walk frequency noise. Our results demonstrate that the performance of the three space-segment timekeeping subsystems can be enhanced by adding general "clock failure rules" to the basic algorithms that are associated with each system. Once in place, these rules provide for robust, autonomous space-segment timekeeping, even in the presence of satellite clock failures.

INTRODUCTION

Military satellite communications (milsatcom) systems require precise timekeeping in order to take advantage of spread-spectrum communication techniques. The level of precise timekeeping in milsatcom, however, is typically not as stringent as that in satellite navigation (satnav) systems. While satnav typically requires nanosecond timekeeping, in very general terms milsatcom lives in a world of microsecond timekeeping. Nevertheless, milsatcom poses its own unique timekeeping problems. Specifically, milsatcom timekeeping must be robust, with the ability to autonomously detect and correct timekeeping problems during protracted periods when the ground control station is either not available or burdened with other pressing tasks.

Here, we consider the ability of three different space-segment timekeeping subsystems at geosynchronous altitude to autonomously detect the failure of a single satellite clock and to then take appropriate action to remedy the situation. These systems include a Master/Slave system, similar to what is presently employed in the Milstar communications system [1,2], an Ensembling system based in part on NIST's AT1 algorithm [3,4], and a Kalman-Filter system similar to what GPS will employ when it takes advantage of crosslink ranging [5,6]. We employ Monte Carlo simulation of timekeeping, and consider the space-segment's response to four distinct types of clock failure, as will be discussed subsequently. The Monte Carlo

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE DEC 1999	2. REPORT TYPE	3. DATES COVERED 00-00-1999 to 00-00-1999			
4. TITLE AND SUBTITLE The Autonomous Detection of Clock Problems in Satellite Timekeeping Systems		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The Aerospace Corporation, PO Box 92957, Los Angeles, CA, 90009		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM001481. 31st Annual Precise Time and Time Interval (PTTI) Planning Meeting, 7-9 December 1999, Dana Point, CA					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 10	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

simulation of the three space-segment subsystems has been described previously [7], and will not be repeated in detail here. Briefly, a sequence of five satellite-constellation update intervals comprises a single Monte Carlo realization. The first four are conducted with the ground-segment monitoring the satellite clocks and correcting their times and frequencies as needed to maintain a somewhat arbitrary 2- μ sec synchronization level. During these “normal” intervals, the ground measures spacecraft time-offsets every 8 hours. Our studies show that by the fourth update interval any numerical transients associated with the initiation of a simulation have died out. The fifth update interval simulates system operation in the absence of ground-segment control, and we record the rate at which time error accumulates during this period. The five-update-interval scenario is repeated 10,000 times for each type of failure that we investigate. In all simulations, we assume the space segment is composed of four satellites, each carrying a complement of essentially identical rubidium (Rb) atomic clocks. For the Master/Slave simulations, the constellation is composed of a master satellite, two monitor satellites, and a single slave satellite. In the other space-segment subsystems, all of the satellites are equivalent. The full set of simulation parameters is provided in Ref. 7.

As noted above, we consider four different “soft” clock failures: a time-jump failure, a frequency-jump failure, a failure arising from a sudden change in the clock’s frequency aging rate (i.e., “aging failure”), and a failure arising from an abrupt increase in the clock’s random-walk frequency noise (i.e., “Allan variance failure”). These failures are especially pernicious during periods of autonomous operation, since we assume that there is no signature, independent of time comparisons among the spacecraft clocks, that indicates a failure. Moreover, we restrict our investigations to specific magnitudes for these failures. We believe that this is fair, since much smaller magnitude failures would have little influence on system performance, while much larger magnitude failures would be indicative of a “hard” failure (i.e., something truly broken within the clock). Presumably, hard failures could be detected by means other than timing comparisons (i.e., a specified voltage falls below a database value).

SPACE-SEGMENT TIMEKEEPING SUBSYSTEMS

MASTER/SLAVE

Perhaps the conceptually simplest space-segment subsystem we consider is a Master/Slave system. Here, the space segment is composed of a Master Satellite Reference (MSR) whose time and frequency are controlled by the ground segment. There is also a constellation of slave satellites that derive time and frequency information from the MSR via intersatellite crosslinks. Additionally, to guard against possible failure of the MSR clock, the constellation contains two independent Monitor satellites (MON1 and MON2) also directly controlled by the ground. Together, the MSR, MON1, and MON2 are referred to as the Triplet. The monitor satellites along with the MSR continuously assess each other’s timekeeping performance via time transfer over the satellite crosslinks.

Should the time difference between two members of the Triplet exceed a database value, Δt_{fail} , an “alarm” is sounded. The Triplet then enters an identification phase (nominally one hour) to determine which of the Triplet members has “failed.” Failure may be defined in one of two ways: 1) a Triplet member’s time-offset to any other Triplet member exceeds Δt_{fail} for the entire identification phase, or 2) a Triplet member’s fractional frequency-offset to any other Triplet member exceeds a database value, Δy_{fail} , over the identification phase. Following the identification phase, the failed Triplet member is demoted to a SLV_{np} role (i.e., non-promotable Slave), and the highest-ranked SLV_{p} (i.e., promotable Slave) promotes itself into a vacated Triplet role. Depending on the failed Triplet member’s identity, several promotional sequences are possible:

For a MON2 failure: $SLV_{\text{p}} \rightarrow \text{MON2}$

For a MON1 failure: MON2 → MON1 and SLV_p → MON2

For an MSR failure: MON1 → MSR, MON2 → MON1 and SLV_p → MON2

The process of: (1) determining that one of the Triplet member clocks has failed, (2) identifying *which* Triplet member has failed, (3) demoting the failed Triplet member, and (4) reconstructing the Triplet with a promotable Slave is referred to as Succession. In the present simulations, we somewhat arbitrarily set $\Delta t_{\text{fail}} = 5 \mu\text{sec}$ and $\Delta y_{\text{fail}} = 1 \times 10^{-10}$ based on our model milsatcom system's 2 μsec (normal operations) synchronization level [8].

ENSEMBLING

With the Ensembling space-segment subsystem, the spacecraft clocks essentially pool their timekeeping information to create an average time and frequency to which each satellite clock steers itself. We employ the NIST AT1 ensembling algorithm [9,10] as this algorithm produces time and frequency offset information for the spacecraft clocks in real time, which is a particularly attractive feature for a milsatcom application. The construction of a timescale via this algorithm requires the periodic determination of time-offsets among all contributing clocks, and these are obtained via time comparisons performed through the satellite crosslinks every τ (i.e., one hour). Using this information, the ensembling algorithm determines time and frequency offsets for each clock with respect to the Ensemble timescale, and makes a prediction for the time and frequency offset of each clock at the next update (i.e., τ in the future). The closer the actual time-offset, δt_i^a , for the i^{th} clock is to its prediction, the greater the weight given to that spacecraft clock in the formulation of the Ensemble timescale. So as to keep any single clock from dominating the Ensemble, we limit the weight that any one clock may have. Specifically, as a rule of thumb we assume that 2/3 of all possible ensemble members, N , are "good," and therefore should contribute to the timescale with roughly equal weights. This would give a nominal weight of $3/2N$ for any one clock, and we limit the clock weights to 110% of this value. Additionally, while we allow the Ensemble to update the time of the various satellites every hour, we limit the interval of frequency updates to once a day so that orbital diurnal temperature effects don't introduce oscillations into the timescale.

Detecting a general clock failure in an autonomous fashion is a nontrivial problem. Though a time-jump failure may be detected in a straightforward manner by the AT1 algorithm, as time-offset readings are directly measurable, a frequency-jump or aging failure is a different matter. In our modification of AT1, a clock is presumed to have failed if over the course of the time interval between Ensemble measurements, τ , the clock's estimated fractional frequency change is greater than Δy_c^E . Additionally, once a clock has been operating for more than $T=24$ hours, we estimate the clock's fractional frequency aging coefficient, \hat{D} , using a three-point estimator:

$$\hat{D} = \frac{2}{T} \left[\delta t_i^a(n\tau) + \delta t_i^a(n\tau - T) - 2\delta t_i^a\left(n\tau - \frac{T}{2}\right) \right] \quad (1)$$

If this (one-day-averaged) estimated aging coefficient is greater than ΔD_c^E , then again the clock is presumed to have failed. In the present work we set $\Delta y_c^E = 5 \times 10^{-11}$ and $\Delta D_c^E = 8 \times 10^{-12}/\text{day}$. Empirically, we found that these values were able to detect the failures we programmed into the simulations, but did not predict failures otherwise.

KALMAN FILTER

Similar to the Ensembling system, the Kalman system pools the timekeeping information of the satellite clocks in order to create a “Kalman-Filter” timescale to which the spacecraft clocks steer themselves. The difference, of course, is that the Kalman-Filter timescale is not a direct average of clock readings at some instant, but rather a filtered estimate of an aggregate constellation state. Basically, the Kalman assumes a model for the i^{th} clock’s time-error at some time τ , $x_i(\tau)$, and for Rb atomic clocks this is best expressed as:

$$x_i(t) = x_i(0) + y_i t + \frac{1}{2} D_i t^2 + \epsilon_i(t) . \quad (2)$$

Here, $x_i(0)$ is the i^{th} clock’s initial time-offset, y_i is the clock’s fractional frequency offset, D_i is the fractional frequency aging rate, and $\epsilon_i(t)$ corresponds to the i^{th} clock’s noise. The triplet $(x_i(0), y_i, D_i)$ corresponds to the i^{th} clock’s state. Again, the construction of the timescale requires periodic time-offset measurements among all contributing clocks, and these are obtained via time comparisons performed through the satellite crosslinks every τ (i.e., one hour).

The discrete time Kalman filter is characterized by a simple set of recursion relations that allow re-estimation of any satellite clock’s state when a new time-offset measurement is obtained [11]. The effects of ground-commanded time and frequency corrections are included in the Kalman filter, so that information on the clock states is carried across ground station update intervals; this results in good estimates of the D_i . Stein and Filler [12] have developed the formalism of the Kalman filter as it relates to precise timekeeping, and have shown how noise (as it relates to the Allan variance) should be incorporated into the Kalman filter; we have employed their results in our work.

As with the Ensembling system, we augment the basic Kalman system with a few simple rules in order to deal with clock failures during an autonomy period. Specifically, similar to the Ensembling system, a clock is presumed to have failed if over the course of the time interval between measurements, τ , the clock’s estimated fractional frequency change is greater than Δy_c^K . Additionally, once a clock has been operating for more than $T=24$ hours, we estimate the clock’s fractional frequency aging coefficient, \hat{D} , using a three-point estimator:

$$\hat{D} = \frac{2}{T} \left[\delta t_i^a(n\tau) + \delta t_i^a(n\tau - T) - 2\delta t_i^a\left(n\tau - \frac{T}{2}\right) \right] . \quad (3)$$

If this (one-day-averaged) estimated aging coefficient is greater than ΔD_c^K , then again the clock is presumed to have failed. In the present work we set $\Delta y_c^K = 5 \times 10^{-11}$ and $\Delta D_c^K = 8 \times 10^{-12}/\text{day}$. These are the same values used by the Ensembling system, and again they were chosen on an empirical basis.

SPACECRAFT ATOMIC CLOCKS

In all cases we assume that the spacecraft carry relatively high quality rubidium (Rb) atomic clocks defined by an Allan standard deviation, $\sigma_y(\tau)$, of

$$\sigma_y(\tau) = 1.5 \times 10^{-11} / \sqrt{\tau} + 1.0 \times 10^{-15} \sqrt{\tau} , \quad (4)$$

and an average fractional frequency aging rate, $\langle D \rangle$, of $7 \times 10^{-14}/\text{day}$. Moreover, we assume that the actual aging rate varies from clock to clock about this average by $\pm \frac{1}{2} \langle D \rangle$. The frequency of the Rb atomic clock

is digitally controlled as described in Ref. [7]. The one exception to the parameters given in Ref. [7] concerns the threshold for satellite fractional frequency corrections: in the Master/Slave and Kalman systems this threshold is set at 3.3×10^{-12} , while in the Ensemble system we obtain slightly better performance by setting this threshold level at 9.0×10^{-12} .

RESULTS

TIME-JUMP FAILURE

Figure 1 shows rms time-error, δt_{rms} , as a function of days into an autonomy period. The solid line corresponds to the time-error buildup for the failed satellite clock, *in the absence of a space-segment timekeeping subsystem*. In this particular failure scenario, a satellite clock suffers a jump in its time reading of 10 μsec at day 2 during an autonomy period. As no other problem occurs for the clock, time-error buildup proceeds nominally after day 2. In the case of the Master/Slave system, the failed clock is the MSR, while in the case of the Ensemble and Kalman systems the failed clock is simply a contributing member to the aggregate timescale.

In the figure, circles correspond to δt_{rms} for the Master/Slave system, squares correspond to δt_{rms} for the Ensemble system and diamonds correspond to δt_{rms} for the Kalman system. In each case, δt_{rms} for the failed clock is shown. Thus, in the case of the Master/Slave system, the Master is quickly identified by the monitors as having a failed satellite clock, and it is demoted to the role of unpromotable slave.* In the case of the Ensemble and Kalman systems, the failed clock is also demoted to a slave role (i.e., a non-contributing member to the aggregate timescale with its time and frequency tied to the aggregate timescale). Though it may be noticed that the Kalman system displays somewhat smaller time-errors than the Master/Slave and Ensemble systems, it is important to note that none of these systems has been fully optimized for the present investigations. Consequently, the main point to draw from the figure is not the superiority of one system compared to another, but rather the fact that all three systems efficiently detect the failure and take corrective action.†

FREQUENCY-JUMP FAILURE

Figure 2 shows rms time-error, δt_{rms} , as a function of days into an autonomy period in the case of a frequency-jump failure. The failure again occurs at day 2 in an autonomy period, but now the failed clock suffers a fractional frequency jump of 1×10^{-10} . Similar to Fig. 1, the solid line corresponds to the time error build-up for the failed satellite clock, *in the absence of a space-segment timekeeping subsystem*, while circles correspond to the same failed clock in the Master/Slave system, squares the Ensemble system and diamonds the Kalman system. As is clear from the figure, all three systems efficiently detect the failure of the satellite clock and take corrective action so as to ensure stability of constellation timekeeping.

Figures 3a and 3b illustrate the importance of augmenting the Ensemble and Kalman-Filter systems with rules to guard against clock failure. In Fig. 3a, we compare δt_{rms} for the Ensemble system with and without the rules described above in the case of a time-jump failure. Without the additional rules, the Ensemble

* Note that we only show the time-error at two-day intervals, so that time-error propagation during demotion is not apparent. Time-error dynamics during demotion is a separate problem from the one considered here, which may be defined as a "timescale stability" problem. We plan to treat the dynamic problem in the near future.

† We also wish to point out that in our simulations the failed clock simply has its role changed to that of Slave. However, in a real situation algorithms onboard the spacecraft would be required to turn off the failed clock and turn on a healthy clock. Nevertheless, as the new clock would have no timekeeping history, prudence would detect that it be assigned a slave role until the ground could properly assess its timekeeping performance.

algorithm allows the failed clock to remain a contributing member to the aggregate timescale (albeit a weakly contributing member), and converts the time-jump failure into a frequency “correction” for all healthy spacecraft clocks. Consequently, the entire constellation rapidly builds up a large time error. With the additional rules, however, this is not allowed to happen. Similar effects occur for the Kalman-Filter system as shown in Fig. 3b for the case of a frequency-jump failure. Here, the frequency-jump of the failed clock works its way into the Kalman filter, corrupting the entire timescale. Specifically, the frequency jump is simply reduced by a factor between one third and a quarter (i.e., only one out of four constellation clocks displays the large frequency offset).

AGING FAILURE AND ALLAN VARIANCE FAILURE

Figures 4 and 5 show the response of the three space-segment timekeeping subsystems to an aging failure and an Allan variance failure, respectively. In the case of the aging failure, the failed clock suffers a sudden increase in its frequency aging rate (i.e., $\langle D \rangle = 7 \times 10^{-14}/\text{day}$ becomes $D = 1 \times 10^{-11}/\text{day}$). In the case of the Allan variance failure, the clock suffers a sudden increase in the random-walk term of its frequency noise (i.e., $\sigma_y(\tau) = 1 \times 10^{-15} \sqrt{\tau} \rightarrow \sigma_y(\tau) = 3 \times 10^{-14} \sqrt{\tau}$). Note that in the case of the aging failure, the Master/Slave system shows a small time-error increase around the day of the failure (i.e., for clocks with no random-walk frequency noise the failure would occur at day 5.4). This is simply a consequence of the fact that the Triplet exceeds Δt_{fail} at some random time around day 5.4 as a consequence of (normal) timekeeping fluctuations. Notwithstanding this apparent “bump” in time-error, Figs. 4 and 5 again show that all three space-segment timekeeping subsystems are able to protect the constellation against clock failures, so that over the long term the constellation timescale is not corrupted by the clock failure.

SUMMARY

In this work we have considered various types of “soft” clock failure, and the ability of three different space-segment timekeeping subsystems to detect the failure and take appropriate action. By soft failure we mean a clock failure that results in poor timekeeping performance, yet all other indications of clock operation (i.e., various critical voltage levels) are within nominal bounds. Basically, the problem we have considered here deals with the creation of a stable, autonomously operating timescale. Our results show that a Master/Slave system, an Ensemble system, or a Kalman-Filter system are all able to create a stable timescale in the presence of clock failures, so long as the basic systems are enhanced by adding general “clock failure rules.”

REFERENCES

- [1] J. C. Camparo, R. P. Frueholz, and A. P. Dubin, “Precise time synchronization of two Milstar communication satellites without ground intervention,” *Int. J. Satellite Communications* 15, 135-139 (1997).
- [2] J. C. Camparo and R. P. Frueholz, “Monte Carlo simulations of precise timekeeping in the Milstar communication satellite system,” in *Proc. 26th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting - NASA Conference Publication 3302* (NASA, Greenbelt MD, 1995), pp. 291-304.
- [3] M. A. Weiss, D. W. Allan, and Trudi K. Peppler, “A study of the NBS time scale algorithm,” *IEEE Trans. Instrum. Meas.* 38(2), 631-635 (1989).

- [4] J. C. Camparo and R. P. Frueholz, "Space-segment timekeeping for next-generation milsatcom," Aerospace Corporation Report No. TOR-97(1453)-3 (Aerospace Corporation, El Segundo, CA, 1997).
- [5] M. P. Ananda, H. Bernstein, K. E. Cunningham, W. E. Fees, and E. G. Stroud, "Global Positioning System (GPS) autonomous navigation," in *Proc. of IEEE PLANS '90 Position, Location, and Navigation Symposium* (IEEE, Piscataway, NJ, 1990), pp. 497-508.
- [6] K. R. Brown, "The theory of the GPS composite clock," in *Proc. of ION GPS-91* (Institute of Navigation, Washington, DC, 1992), pp. 223-241.
- [7] J. C. Camparo, Y. C. Chan and R. P. Frueholz, "Space segment timekeeping for next generation milsatcom," in *Proc. of the 31st Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting* (U.S. Naval Observatory, Washington D. C., 1999), to be published.
- [8] This value of Δy_{fail} was chosen because it amounts to a little less than a 10 μsec time-offset in one day: large enough to be significantly different from the normal 2 μsec synchronization level, but not so large as to grossly exceed this level.
- [9] D. Allan, J. E. Gray, and H. E. Machlan, "The National Bureau of Standards atomic time scale: Generation, stability, accuracy and accessibility," in *Time and Frequency: Theory and Fundamentals*, NBS Monograph 140, 1974, pp. 205-231.
- [10] D. W. Allan, D. J. Glaze, J. E. Gray, R. H. Jones, J. Levine, and S. R. Stein, "Recent improvements in the atomic time scales of the National Bureau of Standards," in *Proc. of the 15th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting* (U.S. Naval Observatory, Washington DC, 1983), pp. 29-39.
- [11] H. W. Sorenson, "Kalman Filtering Techniques," in *Advances in Control Systems, Theory and Applications*, ed. C. T. Leondes (Academic Press, New York, 1966), pp. 219-292.
- [12] S. R. Stein and R. L. Filler, "Kalman filter analysis for real time applications of clocks and oscillators," in *Proc. 42nd Annual Symposium on Frequency Control* (IEEE Press, Piscataway, NJ, 1988), pp. 447-452.

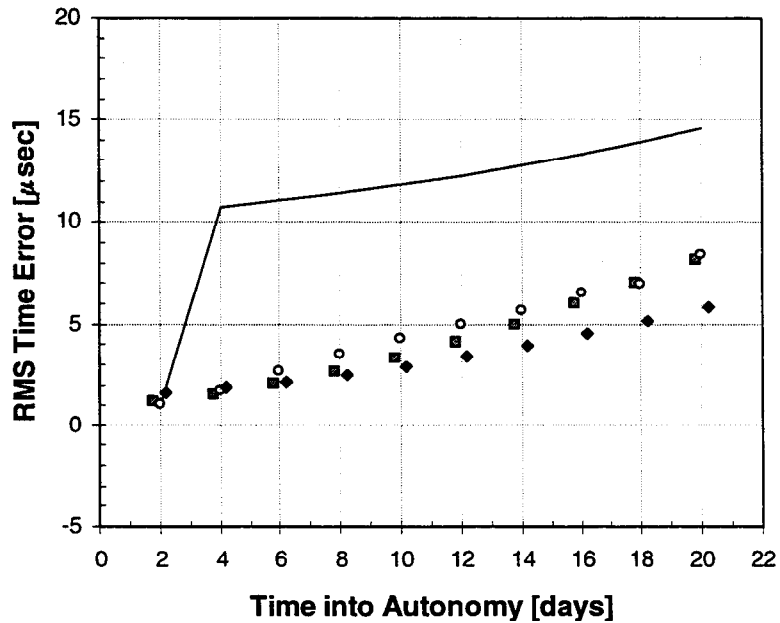


Figure 1: Rms time error as a function of days into an autonomy period for a time-jump Rb clock failure. The solid line corresponds to the time-error buildup that would occur for the failed satellite clock, *in the absence of a space-segment timekeeping subsystem*. Circles correspond to the Master/Slave system, squares to the Ensemble system, and diamonds to the Kalman system. As parameters for the Ensemble and Kalman systems have not been fully optimized, comparisons among the three systems should be treated with caution.

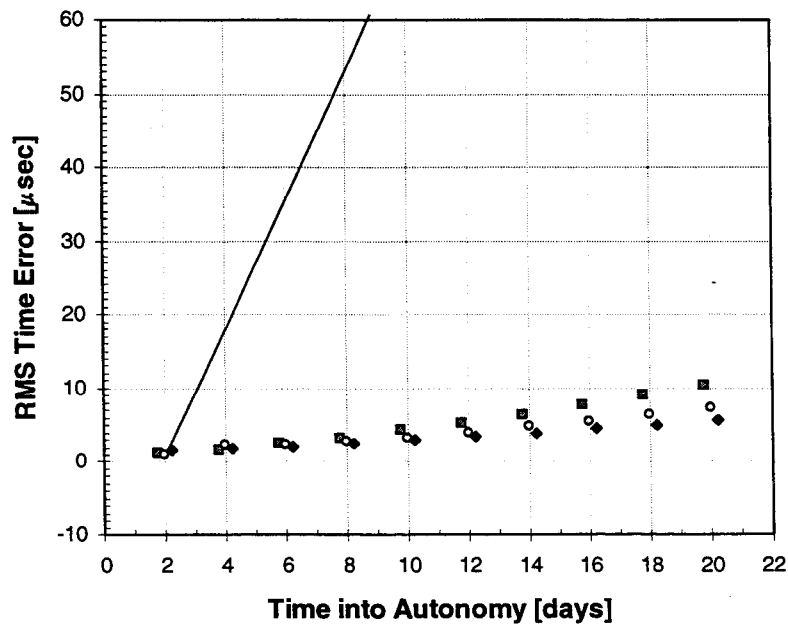


Figure 2: Rms time error as a function of days into an autonomy period for a frequency-jump Rb clock failure. The solid line corresponds to the time-error buildup that would occur for the failed satellite clock, *in the absence of a space-segment timekeeping subsystem*. Circles correspond to the Master/Slave system, squares to the Ensemble system, and diamonds to the Kalman system. As parameters for the Ensemble and Kalman systems have not been fully optimized, comparisons among the three systems should be treated with caution.

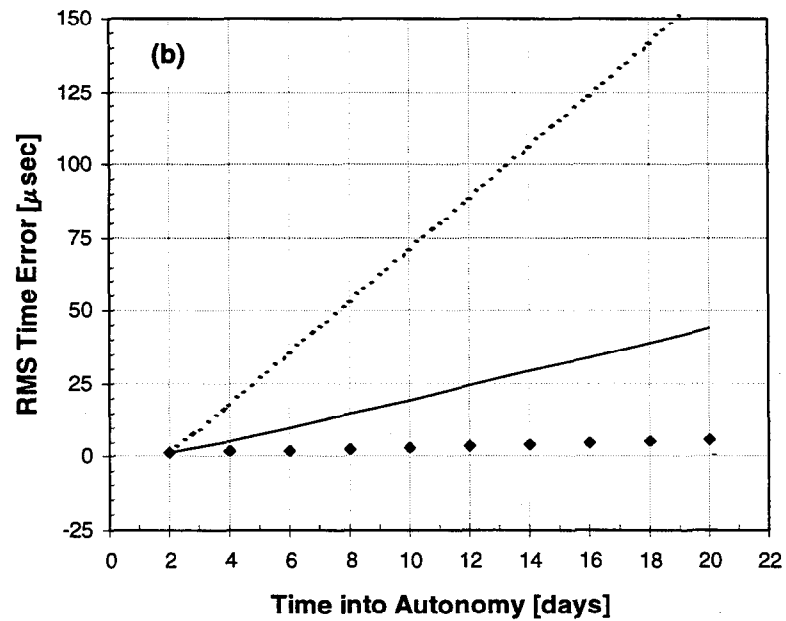



Figure 3: (a) Rms time error buildup in the case of a time-jump failure for the Ensemble system. The dashed line corresponds to the failed clock, *in the absence of a space-segment timekeeping subsystem*; the squares correspond to the failed clock with the Ensemble subsystem and the rules discussed in the text. The solid line is the time-error buildup for the failed clock in the Ensemble *without* the additional rules. (b) Rms time-error buildup in the case of a frequency-jump failure for the Kalman system. Again, the dashed line is the failed clock alone; the diamonds correspond to the Kalman with rules and the solid line the Kalman without rules.




Figure 4: Rms time-error buildup for an aging failure of a Rb clock. The solid line corresponds to the time-error buildup that would occur for the failed satellite clock, *in the absence of a space-segment timekeeping subsystem*. Circles correspond to the Master/Slave system, squares to the Ensemble system, and diamonds to the Kalman system. The dashed line is simply an aid to guide the eye for the Master/Slave system, in order to highlight the slight “bump” in time-error around the time of the clock failure.

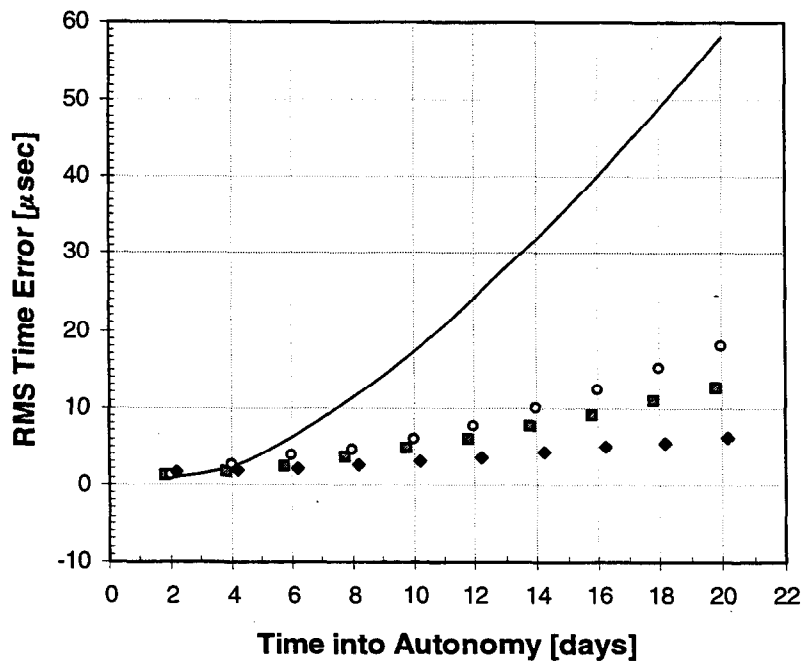


Figure 5: Rms time-error buildup for an Allan variance failure. The solid line corresponds to the time-error buildup that would occur for the failed satellite clock, *in the absence of a space-segment timekeeping subsystem*. Circles correspond to the Master/Slave system, squares to the Ensemble system, and diamonds to the Kalman system. As parameters for the Ensemble and Kalman systems have not been fully optimized, comparisons among the three systems should be treated with caution.