

NIST Technical Note 2187

**A Resilient Architecture for the
Realization and Distribution of
Coordinated Universal Time to Critical
Infrastructure Systems in
the United States**

*Methodologies and Recommendations from the National Institute of
Standards and Technology (NIST)*

Jeffrey A. Sherman
Ladan Arissian
Roger C. Brown
Matthew J. Deutch
Elizabeth A. Donley
Vladislav Gerginov
Judah Levine
Glenn K. Nelson
Andrew N. Novick
Bijunath R. Patla
Thomas E. Parker
Benjamin K. Stuhl
Douglas D. Sutton
Jian Yao
William C. Yates
Victor Zhang
Michael A. Lombardi

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.TN.2187>

NIST
**National Institute of
Standards and Technology**
U.S. Department of Commerce

NIST Technical Note 2187

A Resilient Architecture for the Realization and Distribution of Coordinated Universal Time to Critical Infrastructure Systems in the United States

*Methodologies and Recommendations from the National Institute of
Standards and Technology (NIST)*

Jeffrey A. Sherman, Ladan Arissian, Roger C. Brown, Matthew J. Deutch, Elizabeth A. Donley,
Vladislav Gerginov, Judah Levine, Glenn K. Nelson, Andrew N. Novick, Bijunath R. Patla,
Thomas E. Parker, Benjamin K. Stuhl, Douglas D. Sutton, Jian Yao, William C. Yates,
Victor Zhang, and Michael A. Lombardi
*Time and Frequency Division
Physical Measurement Laboratory*

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.TN.2187>

November 2021



U.S. Department of Commerce
Gina M. Raimondo, Secretary

National Institute of Standards and Technology
*James K. Olthoff, Performing the Non-Exclusive Functions and Duties of the Under Secretary of Commerce
for Standards and Technology & Director, National Institute of Standards and Technology*

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

National Institute of Standards and Technology Technical Note 2187
Natl. Inst. Stand. Technol. Tech. Note 2187, 189 pages (November 2021)
CODEN: NTNOEF

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.TN.2187>

Abstract

The Time and Frequency Division of the National Institute of Standards and Technology (NIST), an agency of the United States Department of Commerce (DOC), was tasked with fulfilling Section 4, Part (i) of the Position, Navigation and Timing (PNT) Executive Order 13905 that was signed by the President on February 12, 2020. The section reads as follows:

“Within 180 days of the date of this order, the Secretary of Commerce shall make available a GNSS-independent source of Coordinated Universal Time, to support the needs of critical infrastructure owners and operators, for the public and private sectors to access.”

This report was written to describe, propose, and recommend technical methods that NIST has already implemented, or can potentially implement, to meet the requirements of the Executive Order. By doing so, it explains how NIST can provide a resilient architecture for the realization and distribution of Coordinated Universal Time (UTC) in the United States.

Keywords

Coordinated Universal Time; critical infrastructure systems; satellites; synchronization; time; time scales; time transfer.

Table of Contents

Executive Summary and NIST Recommendations.....	1
Introduction.....	6
I. The Primary and Secondary UTC(NIST) Time Scales.....	7
I.A. Realization of the Primary NIST Time Scale in Boulder, Colorado.....	8
I.A.1 Technical Description of the Primary Time Scale Architecture	9
I.A.1.1 Generic components of a local UTC time scale	9
I.A.1.2 Principal oscillator: TA(NIST), an ensemble of atomic clocks.....	10
I.A.1.3 Atomic clocks and signal distribution	15
I.A.1.4 Phase measurement system.....	17
I.A.1.5 Data acquisition	21
I.A.1.6 Free-running time scale algorithm AT1, time scale output control.....	23
I.A.1.7 Generation of the UTC(NIST) reference signal and subsequent distribution	26
I.A.2 Methods Ensuring Continuous Operation of Primary Time Scale Architecture.....	29
I.A.2.1 Alarm systems and on-call staff	29
I.A.2.2 Failsafe by default.....	30
I.A.2.3 Redundancy and reliability	30
Primary and Alternate systems overview	31
Contingency systems overview.....	31
Emergency systems overview.....	34
I.A.2.4 Clock health telemetry	34
5071A cesium beam clocks	34
MHM-2010 hydrogen masers	34
AOG-110 synthesizers	35
I.A.2.5 Environmental control	35
I.A.2.6 Redundancy in power supply.....	37
I.A.2.7 Preventative maintenance	38
Room environmental control	38
Chamber environmental control	38
Backup power generators.....	38
Uninterruptible power supplies (UPS) and standby batteries	39
Multichannel measurement systems and data acquisition systems.....	39
I.A.2.8 Analysis of selected failure modes	39
I.A.2.9 Spare parts, component end-of-production risk, modes of improvement	41
Commodity, commercial-off-the-shelf parts	41
Semi-custom, high-performance, no foreseen end-of-production risk	41
Highly-custom, end-of-production risk high	41
I.A.3 Physical and Cyber Security of Primary Time Scale	42
I.A.3.1 Physical security	42
I.A.3.2 Cyber security	42
Confidentiality	42
Integrity.....	43

Availability	44
I.A.4 Validation of the Primary Time Scale as a Trusted Time Reference Source	45
I.A.4.1 Coarse tests of UTC(NIST) pulse-per-second distribution	45
I.A.4.2 Comparisons of Primary, Alternate, and Contingency time scale outputs	46
I.A.5 Accuracy and Stability of Primary Time Scale with Respect to UTC	49
I.A.5.1 UTC(NIST) accuracy and stability measures	49
I.A.5.2 The synthesis of EAL, TAI, and UTC	50
I.B. Realization of Secondary NIST Time Scales in Ft. Collins, Colorado and Gaithersburg, Maryland	54
I.B.1 Technical Description of Secondary Time Scale Architectures	55
I.B.1.1 Time scale at Fort Collins, Colorado	55
I.B.1.2 Time scale at Gaithersburg, Maryland	56
I.B.2 Methods Ensuring Continuous Operation of Secondary Time Scales	57
I.B.2.1 Alarm system and on-call staff	57
I.B.2.2 Failsafe by default	57
I.B.2.3 Redundancy and reliability	57
I.B.2.4 Clock health telemetry	58
I.B.2.5 Environmental control	58
I.B.2.6 Redundancy in power supply	58
I.B.2.7 Preventative maintenance	58
I.B.3 Physical and Cyber Security of Secondary Time Scales	59
I.B.3.1 Physical security	59
I.B.3.2 Cyber security	59
I.B.4 Validation of Secondary Time Scales as Trusted Time Reference Sources	60
I.B.5 Accuracy and Stability of Secondary Time Scales with Respect to UTC	61
II. Methods of Distributing UTC(NIST) at Sub-Microsecond Accuracies	62
II.A. Distribution of UTC(NIST) via Optical Fiber	63
II.A.1 Technical Description of Architecture for Time Over Fiber Distribution System	64
II.A.1.1 Types of optical fibers	64
II.A.1.2 Transport method and message format	64
II.A.1.3 Technical description of the OPNT architecture	66
II.A.2 Reliability and Redundancy of Time over Fiber Distribution System	71
II.A.3 Physical and Cyber Security of Time over Fiber Distribution System	72
II.A.4 Accuracy and Stability of Time over Fiber Distribution	73
II.A.4.1 Metrics utilized to evaluate the stability of a time over fiber system	76
II.A.4.2 Calibration of the delay in the channel from NIST to the remote location	82
The asymmetry in the two-way method due to dispersion	83
II.A.4.3 Results obtained by similar time distribution systems at other locations	84
II.A.5 Feasibility of Distributing UTC(NIST) via Fiber to a Large Number of Users	85
II.B. Distribution of UTC(NIST) via Two-Way Satellite Time and Frequency Transfer	86
II.B.1 Technical Description of Architecture for TWSTFT Distribution System	87
II.B.1.1 Star configuration network	91
II.B.2 Reliability and Redundancy of TWSTFT Distribution System	93

II.B.3 Physical and Cyber Security of TWSTFT Distribution System	94
II.B.4 Accuracy and Stability of TWSTFT with respect to UTC(NIST)	95
II.B.5 Feasibility of Establishing a TWSTFT Distribution Service	96
II.C. Distribution of UTC(NIST) via Multi-Source Common-View Disciplined Clocks ...	97
II.C.1 Technical Description of Architecture for MSCVDC Distribution System	98
II.C.1.1 Common-view disciplined clock (CVDC)	100
II.C.1.2 Multi-source common-view disciplined clock (MSCVDC).....	101
Current NIST Implementation of MSCVDC	103
II.C.1.3 Reference UTC time scales for MSCVDC.....	105
II.C.1.4 Common-view signal sources for MSCVDC	107
Inclusion of a GNSS CVS in a resilient timing architecture	108
Geostationary satellites	109
LF ground-based transmitters	112
II.C.1.5 Data transfer and processing for MSCVDC distribution system	113
Data transfer and processing methods employed by a NISTDC	113
Data transfer and processing methods for widespread deployment.....	113
Frequency of data transfer	114
Methods for providing data transfer redundancy	116
II.C.2 Reliability and Redundancy of MSCVDC Distribution System.....	117
II.C.3 Physical and Cyber Security of MSCVDC Distribution System.....	119
II.C.4 Validation of MSCVDC as a Trusted Time Reference.....	120
II.C.5 Accuracy and Stability of MSCVDC Distribution System.....	123
II.C.6 Feasibility of a Commercially-Available MSCVDC Product	125
III. Methods of Distributing UTC(NIST) as a Holdover Frequency for GPSDCs	126
III.A. Distribution of UTC(NIST) Frequency to GPSDCs via WWVB	127
III.A.1 Technical Description of WWVB Architecture	128
III.A.1.1 WWVB antenna array	129
III.A.1.2 WWVB transmitters.....	132
III.A.1.3 Helix houses	133
III.A.1.4 Broadcast control elements and station efficiency.....	136
III.A.1.5 Time code generators and time codes (legacy and phase-modulated)	137
III.A.2 Reliability, Redundancy, and Safeguards Ensuring Continuous Operation	140
III.A.2.1 Alarm systems and staff response time	140
III.A.2.2 Power failures and backup generator	141
III.A.2.3 Age of WWVB transmitters.....	142
III.A.3 Physical and Cyber Security of WWVB	143
III.A.3.1 Physical security	143
III.A.3.2 Cyber security	144
III.A.4 Frequency and Time Stability of WWVB with respect to UTC(NIST)	145
III.A.5 Feasibility of a Commercial WWVB Disciplined Oscillator	148
III.B. Distribution of UTC(NIST) Frequency to GPSDCs via the PTP Ring Method.....	151
III.B.1 Conceptual and Experiment Description of PTP Ring Architecture	152

III.B.2 Reliability and Redundancy of PTP Ring Architecture.....	153
III.B.3 Physical and Cyber Security of PTP Ring Architecture.....	153
III.B.4 Frequency and Time Stability of PTP Ring with Respect to UTC(NIST)	154
III.B.5 Feasibility of a Nationwide PTP Ring Architecture	156
IV. Indirect Distribution of UTC(NIST) via non-GNSS, non-NIST Signals	157
IV.A. Indirect Distribution of UTC(NIST) via the STL service.....	158
IV.A.1 Frequency and Time Stability of STL with respect to UTC(NIST).....	159
IV.A.2 Possible Establishment of a STL Ground Monitoring Station at NIST	161
IV.B. Indirect Distribution of UTC(NIST) via eLoran	162
IV.B.1 Possible Architecture of a Future eLoran System	162
IV.B.2 Potential Methods for Indirectly Distributing UTC(NIST) via eLoran.....	164
V. Reception of UTC(NIST) Time Codes for GNSS Time Validation	166
V.A. Reception of UTC(NIST) Time Codes via the Internet	167
V.B. Reception of UTC(NIST) Time Codes via WWVB, WWV, and WWVH	169
VI. Acknowledgements.....	170
VII. References.....	171

List of Tables

Table I.A.1. Atomic clocks and synthesizers in primary NIST time scale in Boulder	16
Table I.A.2. List of monitored time scale conditions that generate alarms to on-call staff	29
Table I.A.3. List of monitored hydrogen maser health telemetry signals	35
Table I.A.4. List of primary, alternate, and contingency power plans for time scale assets	37
Table I.A.5. Summary of component failure modes in the primary time scale system	40
Table I.B.1. Atomic clocks and synthesizers in secondary NIST time scale in Ft. Collins	55
Table I.B.2. List of secondary time scale conditions that generate alarms in Boulder.....	57
Table II.C.1. Steps performed by an MSCVDC distribution system.....	103
Table II.C.2. Signals received by multi-constellation GNSS timing receivers	107
Table II.C.3. List of direct broadcast satellites that could potentially provide a CVS	111
Table III.A.1. Percentage of WWVB time on air (2010 to 2019)	140
Table V.A.1. List of publicly accessible servers in NIST Internet Time Service (ITS)	167
Table V.A.2. List of authenticated NTP servers in NIST Internet Time Service (ITS)	168
Table V.B.1. Locations, frequencies, and power of the NIST time signal radio stations	169

List of Figures

Fig. I.A.1. Generic illustration of the major components of a UTC(<i>k</i>) time scale	10
Fig. I.A.2. Graph of holdover performance of various free-running atomic clocks	11
Fig. I.A.3. Diagrams of principal oscillators of varying complexity	12
Fig. I.A.4. Diagram illustrating how physical signals from UTC(NIST) are realized	14
Fig. I.A.5. Photograph of auxiliary output generator that produces time scale signals	14
Fig. I.A.6. Photograph of hydrogen maser inside an environmental chamber	15
Fig. I.A.7. Diagram of time scale signal routing for cesium clocks and hydrogen masers	18
Fig. I.A.8. Photograph of time scale multi-channel measurement system (MCMS).....	18
Fig. I.A.9. Diagram of a single measurement channel in a single MCMS module	19
Fig. I.A.10. Connectivity diagram of time scale data acquisition machine	21
Fig. I.A.11. Graph of UTC(NIST) versus UTC and UTCr (mid-2016 to mid-2020)	26
Fig. I.A.12. Historical graph showing rate offset of UTC(NIST) – TA(NIST)	27
Fig. I.A.13. Diagram of the redundant time scale arrangement in building B1	32
Fig. I.A.14. Diagram of the contingency time scale arrangement in building 81	33
Fig. I.A.15. Graph of temperature and humidity inside a maser environmental chamber.....	36
Fig. I.A.16. Photograph of three impulse clock displays driven by PPS from UTC(NIST) ...	45
Fig. I.A.17. Graphs of primary time scale versus alternate and contingency time scales	47
Fig. I.A.18. Graph of time differences between time scale outputs in buildings B1 and 81 ...	48
Fig. I.A.19. Graph of UTC(NIST) time stability vs. UTC (mid-2016 to mid-2020)	49
Fig. I.A.20. Graph of UTC(NIST) frequency stability vs. UTC (mid-2016 to mid-2020)	50
Fig. I.A.21. Historical graph of the number of atomic clocks that contribute to UTC	51
Fig. I.A.22. Diagram drawing analogy between EAL/UTC and TA(NIST)/UTC(NIST)	51
Fig. I.B.1. Diagram of transfer links between NIST primary and secondary time scales	54
Fig. I.B.2. Graph comparing secondary time scale (Ft. Collins) to primary time scale	61
Fig. II.A.1. Diagram of hardware configuration used in time over fiber calibration test	67
Fig. II.A.2. Graph of time over fiber comparison between Maryland and Virginia	67
Fig. II.A.3. Diagram of alternate configuration for time over fiber calibration test	68
Fig. II.A.4. Graphs of time over fiber loop back tests using different averaging periods	73
Fig. II.A.5. Diagram of portable rubidium clock calibration method	77
Fig. II.A.5. Diagram of portable GPS receiver calibration method	78
Fig. II.A.5. Graph of variations in the calibration constant of a portable GPS receiver.....	79
Fig. II.A.5. Graph showing differences in remote clock calibration methods	79
Fig. II.A.5. Graph of primary vs. secondary time scale calibrations (three methods)	81

Fig. II.B.1. Diagram of two TWSTFT stations engaged in a clock comparison	87
Fig. II.B.2. Spectrum of psuedorandom noise codes for TWSTFT at 70 MHz IF	88
Fig. II.B.3. Photographs of components of a NIST TWSTFT earth station	89
Fig. II.B.4. Graphs of a TWSTFT comparison between NIST and PTB (Germany)	90
Fig. II.B.5. Star configuration network for distribution of UTC(NIST) via TWSTFT	91
Fig. II.C.1. Diagram of a common-view time transfer system	98
Fig. II.C.2. Diagram of a common-view disciplined clock system	100
Fig. II.C.3. Diagram of a multi-source common-view disciplined clock (MSCVDC)	102
Fig. II.C.4. Photograph of NIST’s MSCVDC implementation, the NISTDC instrument.....	104
Fig. II.C.5. Diagram of a NISTDC	104
Fig. II.C.6. Photograph, with parts labelled, of the interior of a NISTDC chassis	105
Fig. II.C.7. Display of MSCVDC software that controls the NISTDC	106
Fig. II.C.8. Map of the coverage area of a direct broadcast geostationary satellite.....	110
Fig. II.C.9. Graph showing how MSCVDC improves long-term rubidium clock stability...115	
Fig. II.C.10. Graph comparing frequency stability of four free running quartz oscillators...116	
Fig. II.C.11. Diagram of a GPS disciplined clock (GPSDC)	117
Fig. II.C.12. Diagram of a common-clock MSCVDC calibration at NIST	120
Fig. II.C.13. Graph comparing direct vs. common-view comparisons of a NISTDC	121
Fig. II.C.14. Graph of time accuracy of MSCVDC deployed at a U. S. stock exchange	123
Fig. II.C.15. Graph of time stability of MSCVDC deployed at a U. S. stock exchange	124
Fig. III.A.1. Diagram of WWVB radio station architecture	128
Fig. III.A.2. Site map of NIST radio stations WWV and WWVB	129
Fig. III.A.3. Diagram of the WWVB antenna array	130
Fig. III.A.4. Aerial photograph of the WWVB antenna array	130
Fig. III.A.5. Photograph of connection of the WWVB antenna to one of the towers	131
Fig. III.A.6. Diagram of a single WWVB transmitter	132
Fig. III.A.7. Photograph of exterior of a WWVB helix house.....	133
Fig. III.A.8. Photograph of roof of helix house showing connection to antenna array	134
Fig. III.A.9. Photograph of interior of helix house showing variable inductor	134
Fig. III.A.10. Diagram of helix house, showing connections to antenna and transmitter	135
Fig. III.A.11. Diagram of WWVB RF switch matrix	136
Fig. III.A.12. Photograph of a WWVB time code generator.....	137
Fig. III.A.13. Diagram of the legacy WWVB time code format	138
Fig. III.A.14. Diagram comparing frame structure of PM time code to legacy time code....	139

Fig. III.A.15. Photograph of the motorized entrance gate to the WWVB site.....143

Fig. III.A.16. Photograph of the perimeter fence at the WWVB site144

Fig. III.A.17. Phase graph comparing two WWVB disciplined oscillators to UTC(NIST) .146

Fig. III.A.18. Frequency stability of two WWVB disciplined oscillators vs. UTC(NIST) ..146

Fig. III.A.19. Time stability of two WWVB disciplined oscillators vs. UTC(NIST)147

Fig. III.A.20. Photographs of a WWVB radio controlled wall clock and wristwatch.....148

Fig. III.A.21. Oscilloscope display comparing WWVB on-time marker to UTC(NIST)149

Fig. III.A.22. Graph of model of received phase of WWVB (~2000 km path).....150

Fig. III.B.1. Diagram of a proposed PTP ring network151

Fig. III.B.2. Diagram of an experimental setup of a PTP ring on a local network.....154

Fig. III.B.3. Graph comparing two PTP servers (one with GPS, one without)155

Fig. III.B.4. Diagram of experimental PTP ring with a local and remote network155

Fig. IV.A.1. Photograph of a prototype STL disciplined clock159

Fig. IV.A.2. Graph comparing the time accuracy of a STL disciplined clock to a GPSDC..160

Fig. IV.A.3. Graph comparing the time stability of a STL disciplined clock to a GPSDC...160

Fig. IV.A.4. Graph of average one-day differences between STL clock and UTC(NIST) ..161

Fig. IV.B.1. Map of coverage area of proposed eLoran system with four transmitters163

Fig. IV.B.2. Map of proposed eLoran timing service, showing differential stations163

Fig. IV.B.3. Diagram of differential eLoran station sending corrections to transmitter.....165

Executive Summary and NIST Recommendations

For many years prior to the issuance of *Executive Order 13905*ⁱ, the National Institute of Standards and Technology (NIST) has faithfully served the time and frequency needs of the private and public sectors in the United States. This effort has included the continuous realization of the UTC(NIST) time scale and the continuous distribution of UTC(NIST) time signals through various broadcast and calibration services. In recent years, NIST has demonstrated a strong commitment to study the extent of the dependency that U. S. critical infrastructure has on GPS time, both in technicalⁱⁱ and economic terms. For example, a 2019 study conducted by RTI International on behalf of NIST provided estimates of the economic consequences of 30-day GPS outages in various critical infrastructure sectors. The estimated economic cost of a loss of GPS timing was severe in the telecommunications sector, with loss estimates ranging between \$5.5 and \$14.2 billion. In the electric power sector, the estimated loss ranged from \$211.6 to \$338 millionⁱⁱⁱ. Overall, the report estimated that the combined loss of GPS position, navigation, and timing (PNT) services would have an adverse impact on the U. S. economy of at least \$1 billion per day^{iv}. Earlier, through funding offered by its Small Business Innovation Research (SBIR) program in 2016, NIST had encouraged industry to develop a single-chip eLoran receiver, noting that “the GPS signal is exceedingly weak, and it is vulnerable to interference, both accidental and deliberate.”^v

This report goes well beyond previous NIST efforts by presenting an overview of a resilient architecture for the realization and distribution of Coordinated Universal Time (UTC). This architecture combines existing systems that are already operational, with new technologies that are either proposed for development, or are already under development. The full development of a resilient architecture will, of course, require many additional commitments. These commitments must be made not only by NIST, but also by industrial partners that develop and manufacture commercially available products, and by public and private sector entities that can help provide necessary funding. In addition, before any change in the current culture of time distribution and usage is possible, critical infrastructure users must first develop an awareness of the risks associated with sole reliance on GPS and other global navigation satellite system (GNSS) timing signals, and then must also demonstrate a willingness to make the investments necessary to build resiliency into their timing systems.

Our key findings and recommendations for the five sections of the report are summarized as follows.

Section I. The Primary and Secondary UTC(NIST) Time Scales – This section describes the primary and secondary time scales that realize UTC at NIST. These time scales, known collectively as UTC(NIST), are the foundation of the resilient architecture presented in this report.

- **Finding I.A** - The UTC(NIST) time scales are a tremendous national resource, and NIST has a well-established track record of delivering UTC at high levels of accuracy with outstanding reliability. The technical expertise gained by NIST from many years of maintaining a national frequency or time standard would be difficult for another organization to either emulate or reproduce. In addition, the optical clock research program now in place at NIST makes the agency the national, and perhaps world leader, in the effort to redefine the SI second from the existing microwave definition based on an energy transition of cesium atoms to a new definition based on optical clocks^{vi}. Plans are in place and to eventually integrate optical clocks into the UTC(NIST) time scale^{vii}, and if history is any guide, the resulting improvements in frequency and time accuracy will eventually enable new technologies that will greatly benefit the citizens of the United States.

Recommendation I.A – As a result of this finding, we recommend full utilization of the UTC(NIST) time scale in all future plans to develop a resilient architecture to support critical

infrastructure timing systems and feel that continued and increased investment in the NIST time scale should become a national priority.

- **Finding I.B** - The UTC(NIST) time scales are engineered to be highly robust using a Primary, Alternate, Contingency, Emergency (PACE) framework and the primary time scale in Boulder, Colorado has operated continuously for more than 50 years through power outages, floods, HVAC failures, and many individual component failures. However, this resilience requires continuous “on-call” vigilance by a limited staff of highly trained personnel and imposes a substantial preventative maintenance burden. While most components of the Primary and Alternate UTC(NIST) time scales have reasonable commercial availability, the Multichannel Measurement System (MCMS; used to integrate the set of individual atomic clocks into an ensemble average) faces substantial obsolescence / end-of-production risk.

Recommendation I.B – As a result of this finding, we recommend that ongoing efforts to evolve the UTC(NIST) time scales to reduce maintenance should continue to be supported. Efforts to qualify an alternative to the legacy MCMS should be increased so that an alternative is available before operational components fail.

- **Finding I.C** - The secondary NIST time scales (in Fort Collins, Colorado and Gaithersburg, Maryland) provide off-site contingency-level backup capability to the primary realization of UTC(NIST), in addition to supporting radio stations WWV and WWVB, Internet time services, and the first implementation of NIST’s fiber-based time transfer calibration service (in Gaithersburg, Maryland). These secondary time scales are regularly steered to keep their offset within ± 25 ns of UTC(NIST). In the event of a total failure of the primary time scales, the secondary time scales will accumulate approximately ± 30 ns of error in the first month.

Recommendation I.C – As a result of this finding, we recommend that the secondary time scales be given more attention and receive further investments – for example, by increasing the number of atomic clocks and by improving their remote and local measurement systems – with the overall goal of increasing their distribution capabilities and their performance. In addition, we recommend better integration of the secondary time scales into the overall time scale program at NIST.

Section II. Methods of Distributing UTC(NIST) at Sub-Microsecond Accuracies – This section presents three methods that NIST employs to distribute time at sub-microsecond accuracies: time over optical fiber, two-way satellite time and frequency transfer (TWSTFT), and the Multi-Source Common-View Disciplined Clock (MSCVDC) method.

- **Finding II.A** - Time over optical fiber has, up until now, not received sufficient attention from NIST, largely due to the large size of the North American continent and the difficulties of establishing long fiber routes to connect timing systems. The use of dedicated fiber can provide a level of security that wireless signals often cannot, as it is immune from RF interference and other disturbances from outside sources, and it is entirely free from any reliance on satellite signals. Using White Rabbit or similar protocols, the time over optical fiber method is capable of providing calibrated, sub-microsecond accuracy links to UTC(NIST) over long-distance optical fiber links. Ultimate accuracies well below 100 ns should be achievable with suitable calibration of end-user equipment and cables. An initial test that demonstrates this accuracy is underway between the NIST facility in Gaithersburg, Maryland and commercial facilities in McLean, Virginia and Atlanta, Georgia.

Recommendation II.A - It seems clear that time over fiber systems will play a large role in future resilient time architecture efforts, and we recommend additional research to determine an optimal path for integration of wide-area time distribution into commercial fiber networks. The cost of the time distribution hardware is not prohibitive. However, dedicated optical fiber links do incur substantial ongoing costs, and the method's ability to "scale up" and reach enough users to meet critical infrastructure needs requires further study.

- **Finding II.B** - Two-way satellite time and frequency transfer (TWSTFT) provides point-to-point synchronization at the nanosecond level using communications satellites. Typical applications use geostationary relay satellites, with substantial airtime costs. Airtime requirements decrease as the end-user's local clock improves in quality, as fewer measurements are required to maintain a given level of accuracy. TWSTFT equipment is commercially available but requires a suitable satellite dish and estimated capital costs are near \$150 000 (2020 USD) to construct a TWSTFT ground station. NIST currently utilizes TWSTFT to contribute to UTC, which is indicative of its high level of performance. However, because TWSTFT requires dedicated airtime both on the satellite and from NIST's ground station, scaling TWSTFT services beyond a few tens of clients would be challenging. TWSTFT time delivery to a set of regional distribution networks would likely have better scaling behavior.

Recommendation II.B - Due to its nanosecond-level accuracy and its independence from GNSS, we recommend that other uses of TWSTFT be explored, including the possibility of a TWSTFT service at NIST, the use of TWSTFT to synchronize the NIST secondary time scales, and additional research and development that could lead to lower-cost Earth stations, more wide spread deployment, and more users.

- **Finding II.C** - The Multi-Source Common-View Disciplined Clock (MSCVDC) method is exceptionally versatile, as it can accommodate a diverse array of both reference time scales and transmitted signals. A primary advantage of the MSCVDC method is that it separates the time scale source from the signal transmission source. For example, it can potentially derive its reference from either the primary or secondary NIST time scales, or from any other UTC time scale. It does so by utilizing observations of multiple transmitters (GNSS satellites, geostationary communications satellites, or potentially terrestrial sources) to synchronize between remote sites. Specifically, the same sources are observed simultaneously at both sites, and the time difference of reception is then relayed over any available communications channel. NIST currently offers a partial implementation of an MSCVDC via its NIST disciplined clock (NISTDC) service; currently over 20 NISTDC units are in operation. The NISTDC implementation is not GNSS independent, but the inclusion and use of non-GNSS common-view signals, originating from geostationary satellites or from terrestrial transmitters, is entirely feasible.

Recommendation II.C - The MSCVDC method has considerable appeal because it is essentially a software-based solution that integrates other technologies that already exist. Multiple, non-overlapping layers of synchronization are desired characteristics of any resilient time distribution architecture, and the MSCVDC method allows adding as many layers as necessary. Thus, we recommend increases in the attention and investment given to the MSCVDC method, perhaps leading to the development of commercially available products.

Section III. Methods of Distributing UTC(NIST) as a Holdover Frequency for GPS Disciplined Clocks –

This section discusses ways that UTC(NIST) can assist existing GPS disciplined clocks (GPSDCs) by providing a holdover frequency source that clocks can rely upon during GPS signal outages. For example, a GPSDC could maintain time accuracy (at the $\leq 1 \mu\text{s}$ level) throughout long GPS disruption events if

they are modified to accept additional frequency references stable at the 10^{-12} level, such as WWVB-disciplined oscillators (VBDOs) or IEEE 1588 Precision Time Protocol (PTP) network connections to NIST time scale sites, or to sites where GNSS, TWSTFT, or MSCVDC time information is available. Because the GPSDC was already synchronized prior to a signal outage, only frequency information will be required to maintain synchronization during the outage.

- **Finding III.A** - Radio Station WWVB (Section III.A) has long been a national resource that delivers time-of-day information to many millions of low-cost radio controlled clocks, but the use of its stable groundwave signals as a redundant source of frequency for GNSS clocks has perhaps been overlooked.

Recommendation III.A - We recommend that manufacturers of GNSS clocks investigate the use of integrating WWVB receivers, and if necessary, work with NIST to improve the signal for that purpose. We recommended future investments in WWVB to modernize the station, to strengthen its long-term reliability, and to encourage more widespread usage of the phase modulated (PM) time code.

- **Finding III.B** - The proposed PTP ring method can allow networked GNSS clocks, located in areas where the GNSS signal is currently disrupted, to access a redundant source of frequency from another clock located elsewhere on the network.

Recommendation III.B - We recommend that manufacturers of GNSS clocks explore the possibility of having clocks “fall over” to a PTP frequency source, and to consider the inclusion of “PTP Ring” capability into their future products.

Section IV. Indirect Distribution of UTC(NIST) via non-GNSS, non-NIST Signals – This section discusses how systems not controlled by NIST or another government agency can become essential parts of a resilient time distribution architecture by indirectly distributing UTC(NIST).

- **Finding IV.A** - Many time distribution systems, both publicly and privately funded, do not originate from NIST but offer the potential of GNSS independence. These systems can benefit, in terms of performance, redundancy, and in the establishment of traceability to UTC, from a connection to UTC(NIST).

Recommendation IV.A - We recommend that NIST become more involved in the monitoring programs of non-NIST, non-GNSS systems, perhaps leading to UTC(NIST) corrections being included in some broadcasts. We highlight two systems, the Satellite Time and Location (STL) service, and the ground-based eLoran service, but note that other systems could also serve as indirect distribution sources for UTC(NIST).

Section V. Reception of UTC(NIST) Time Codes for GNSS Time Validation – This section discusses how systems referenced to UTC(NIST) that are not capable of microsecond-level accuracy can still contribute to a resilient time distribution architecture, by providing time-of-day validation checks for GNSS clocks.

- **Finding V.A** – Although the time codes broadcast via radio (WWVB, WWV, and WWVH) or via the Internet Time Service (ITS) are not capable of synchronization to UTC(NIST) at the $\leq 1 \mu\text{s}$ accuracy level, the integration of one or more of these services into the internal health and performance monitoring logic of GPSDCs could substantially improve those clocks’ ability to detect and reject spoofed GNSS signals.

Recommendation V.A – We recommend that GNSS clock manufacturers also consider the use of UTC(NIST) time codes broadcast by WWVB, by the NIST Internet Time Service (ITS), and by the shortwave radio stations WWV and WWVH, as verification sources for GNSS time-of-day information.

In conclusion, the authors wish to emphasize that NIST always welcomes new ideas about how to improve our systems and methodologies, as we continuously strive to make them more accessible, usable, and valuable to critical infrastructure owners and operators.

ⁱ Title 3, The President, “Executive Order on Strengthening National Resilience through Responsible Use of Positioning, Navigation, and Timing Services,” *Executive Order 13905*, Federal Register, vol. 85, no. 3, pp. 9359-9361, February 18, 2020.

ⁱⁱ M. Lombardi, “An Evaluation of Dependencies of Critical Infrastructure Timing Systems on the Global Positioning System (GPS)” summary report submitted to DHS Science and Technology (S&T) Directorate, 57 p., February 28, 2020.

ⁱⁱⁱ A. O’Connor, M. Gallaher, K. Clark-Sutton, D. Lapidus, Z. Oliver, T. Scott, D. Wood, M. Gonzalez, E. Brown, and J. Fletcher, “Economic Benefits of the Global Positioning System (GPS),” *RTI Report Number 0215471*, sponsored by the National Institute of Standards and Technology, RTI International, 306 p., June 2019.

^{iv} *RTI Report*, p. ES-4.

^v “Single-Chip eLoran Receiver,” (<https://www.sbir.gov/sbirsearch/detail/1242391>), 2016.

^{vi} W. McGrew, X. Zhang, H. Leopardi, R. Fasano, D. Nicolodi, K. Beloy, J. Yao, J. Sherman, S. Schäffer, J. Savory, R. Brown, S. Römisch, C. Oates, T. Parker, T. Fortier, and A. Ludlow, “Towards the optical second: verifying optical clocks at the SI limit,” *Optica*, vol. 6, no. 4, pp. 448–454, 2019.

^{vii} J. Yao, J. Sherman, T. Fortier, H. Leopardi, T. Parker, W. McGrew, X. Zhang, D. Nicolodi, R. Fasano, S. Schäffer, K. Beloy, J. Savory, S. Romisch, C. Oates, S. Diddams, A. Ludlow, and J. Levine, “Optical-Clock-Based Time Scale,” *Physical Review Applied*, vol. 10, 044069/1-12, 2019.

Note that the manufacturer and model names found in this report, as well as the references made to private entities, were included for the sake of technical completeness only, and do not express or imply any endorsement by NIST. Other products may be found to work equally as well or better, and other private entities may be found that offer similar services or products.

Introduction

The Time and Frequency Division of the National Institute of Standards and Technology (NIST), an agency of the United States Department of Commerce (DOC), was tasked with fulfilling Section 4, Part (i) of the Position, Navigation and Timing (PNT) Executive Order that was signed by the President on February 12, 2020. The Executive Order states that

“Within 180 days of the date of this order, the Secretary of Commerce shall make available a GNSS-independent source of Coordinated Universal Time, to support the needs of critical infrastructure owners and operators, for the public and private sectors to access.”

[1]

Although the Executive Order does not specifically provide an accuracy requirement, this report assumes an accuracy requirement of $\pm 1 \mu\text{s}$ with respect to Coordinated Universal Time (UTC). This specification is assumed because numerous critical infrastructure timing systems in the United States are now heavily dependent on signals received from the Global Positioning System (GPS) satellites that routinely provide sub-microsecond accuracy. Therefore, other UTC sources that deliver time at comparable accuracy levels are needed to alleviate or reduce these dependencies. An overview of these dependencies was previously provided by NIST in the report “An Evaluation of Dependencies of Critical Infrastructure Timing Systems on the Global Positioning System (GPS),” which was submitted to the Department of Homeland Security (DHS) in February 2020 [2].

This report was written to describe, propose, and recommend technical methods that NIST has already implemented, or can potentially implement, to fulfill Section 4, Part (i) of the PNT Executive Order. It explains how NIST can provide a resilient architecture for the realization and distribution of UTC in the United States. The report consists of five parts:

- Part I describes the primary and secondary UTC time scales maintained by NIST, known as UTC(NIST). These time scales serve as the reference clocks for each of the time distribution systems currently provided, or proposed, by NIST, and form the foundation of the resilient architecture.
- Part II describes available and proposed methods for distributing time referenced to UTC(NIST) to any U. S. location with sub-microsecond accuracy.
- Part III describes signals referenced to UTC(NIST) that could potentially provide sufficient frequency holdover to allow GPS disciplined clocks (GPSDCs) to continue to provide sub-microsecond accuracy for multiple days in the absence of GPS signals.
- Part IV describes other signals that are independent of GPS and other global navigation satellite systems (GNSS) and that are not operated by NIST, but that could still potentially indirectly distribute UTC(NIST) at sub-microsecond accuracies if real-time UTC(NIST) corrections were applied to their broadcasts.
- Part V describes the potential use of existing UTC(NIST) time codes that are not capable of sub-microsecond accuracy, as verification sources for GNSS time.

I. The Primary and Secondary UTC(NIST) Time Scales

The America COMPETES Act of 2007 [3] specifies that the official time for the United States is Coordinated Universal Time (UTC), as “interpreted or modified by the Secretary of Commerce, in coordination with the Secretary of the Navy.” Because NIST is an agency of the U. S. Department of Commerce, the passage of this bill in 2007 formally established the UTC time scales realized at NIST, known as UTC(NIST), as official reference time sources for the United States.

Prior to 2007, NIST was already officially responsible for determining the time used by various sectors of U. S. industry, including the financial and electric power sectors, and had a long history of maintaining atomic time scales and distributing time signals to the general public. This section describes how NIST “keeps” time, by providing a technical description of the primary NIST time scale, located in Boulder, Colorado, as well as the secondary NIST time scales located in Fort Collins, Colorado and Gaithersburg, Maryland.

I.A Realization of the Primary NIST Time Scale in Boulder, Colorado

The primary NIST time scale, located in Boulder, Colorado, is an official source of time for the United States, and the heart of the resilient time realization and distribution architecture described in this report. While reading this report, note that the primary time scale serves as the reference clock and/or the comparison source for the secondary NIST time scales (Section I.B); and for all of the time distribution methodologies and systems that we describe in subsequent sections.

This section begins by providing a detailed technical description of the primary time scale architecture. That description is followed by discussions of the reliability and redundancy of the primary time scale, its physical and cyber security, the validation methods that ensure the primary time scale is a trusted time reference source, and the accuracy and stability of the primary time scale with respect to UTC.

I.A.1 – Technical Description of the Primary Time Scale Architecture

The most accurate and reliable method of keeping time is to accumulate a count of periodic intervals, the durations of which are related to a (quantum-mechanical) frequency reference. The apparatus responsible for this accumulation, and for the resulting signal, is called a time scale. Errors in a time scale can grow divergently due to biases and noise processes; these are measured by comparison with an international time scale regarded as a reference (i.e., UTC) and kept small by active adjustment. The reference time scale follows similar principles of construction but is generally agreed to exhibit less noise and bias, at least over long averaging intervals, due to the contributions of many independent metrological laboratories.

I.A.1.1. Generic components of a local UTC time scale

The fundamental interval of time measured by the international time scale UTC is the International System (SI) second. Inspired by the SI definition [4] of the second, "...taking the fixed numerical value of the... unperturbed ground-state hyperfine transition frequency of the ^{133}Cs atom [to] be 9,192,631,770 when expressed in the unit Hz, which is equal to s^{-1} ," we might imagine constructing a UTC time scale by directly detecting and counting such periods of radiation from cesium atoms indefinitely. While ideal, this concept is impractical and is never implemented. We reserve a detailed discussion for Section I.A.5 but summarize the technical problem here as the motivation for NIST's time scale design: electromagnetic radiation from an unperturbed population of atoms cannot be observed sufficiently reliably or indefinitely.

Nonetheless, it is possible to indirectly pursue the goal of basing a continuous time scale on the defined SI second with small, quantifiable uncertainties. A metrological laboratory k seeking to approximate the UTC time scale as $\text{UTC}(k)$ will generally follow the method illustrated in Fig. I, with the following important elements:

Principal oscillator – the laboratory obtains or produces an electronic oscillation signal, which has some *a priori* unknown fractional frequency offset $|y| \ll 1$ from a nominal value f_0 , but which operates and is observable continuously and indefinitely. The fractional frequency offset, y , can be defined as $(f - f_0) / f_0$.

Counter/divider – In the laboratory's time scale, $\text{UTC}(k)$, every observation of $f_0/(1 \text{ Hz})$ cycles of the principal oscillator denotes a time interval of 1 s. An electronic counter/divider marks each such accumulation by generating a signal, ideally a "pulse" waveform with a high voltage slew-rate for precise temporal resolution.

Reference point – Three local conventions fix the definition of a $\text{UTC}(k)$ time scale to its underlying periodic electronic signal: a physical location, an electronic measurement protocol, and an association to UTC. Because the speed of all signaling is finite, a physical location (for example, a specific connector in an electronic circuit path) is defined as the reference measurement point for $\text{UTC}(k)$. All other distributions and observations of $\text{UTC}(k)$ must account for additional signaling delay from the reference measurement point. Because electronic signals are emitted at finite impedance and with finite bandwidth, we must also adopt a protocol for indicating the moment in time associated with the periodic signal at the reference point. For example, measuring instruments are configured to "trigger" when the physical pulse voltage rises above a certain threshold under a particular measurement impedance, followed by an appropriate amount of trigger "dead-time" to protect against noise or signal imperfections that could falsely indicate the early arrival of the next pulse. Finally, we must adopt an interpretation that a particular physical pulse will be associated forever with a named integer second in the UTC time scale, regardless of whatever fractional-second time offset, x , exists or develops between $\text{UTC}(k)$ and UTC. A count of pulses from this reference pulse produces the named second in the $\text{UTC}(k)$ time scale.

Transfer – $UTC(k)$ is transmitted (see, for example, Section II.C and [5]) to the International Bureau of Weights and Measures (BIPM), compensating for all signal delays from the $UTC(k)$ reference location. BIPM produces the data series $UTC - UTC(k)$ approximately monthly in a publication called Circular-T, where results are given at five-day intervals, and also in a weekly series “UTC-Rapid”, where $UTC_r - UTC(k)$ is provided at one-day intervals [6]. From these measurements and possibly other observations (for example, direct comparisons with high-accuracy atomic frequency references [7]), knowledge about the principal oscillator’s frequency offset, y , and the time scale’s offset, x , is obtained.

Coordination – Laboratories wishing to minimize the $|UTC - UTC(k)|$ difference, and/or increase the stability of $UTC(k)$, do so by either adjusting the principal oscillator’s frequency offset y , or by compensating for the offset.

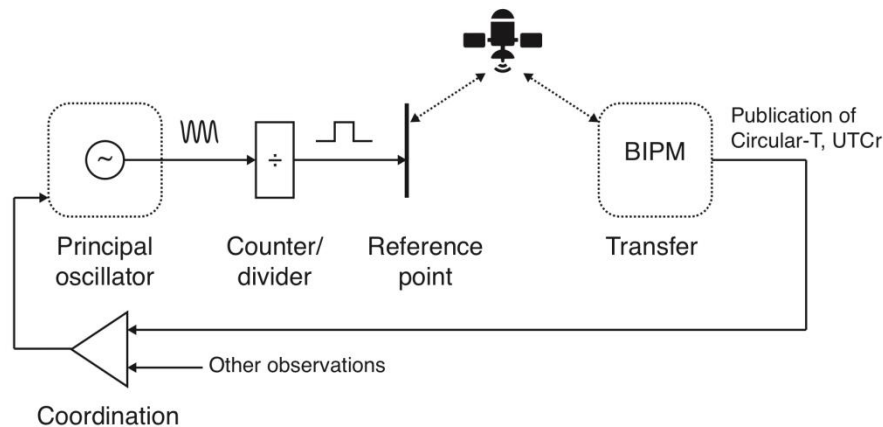


Fig. I.A.1. Generic illustration of the five major components of a coordinated laboratory time scale, $UTC(k)$. Briefly, the oscillations of a clock-derived “principal oscillator” are counted to produce a periodic signal marking integer seconds in $UTC(k)$ when received at a defined reference point location. A representation of this signal is transmitted to the BIPM (with all signaling delays measured or calibrated); BIPM publishes the data series $UTC - UTC(k)$ approximately monthly and the series $UTC_r - UTC(k)$ weekly. Practitioners seek to coordinate their time scale by minimizing $|UTC - UTC(k)|$ and/or the instability of $UTC(k)$; either is accomplished by making smooth phase and frequency tuning adjustments to the principal oscillator.

The remainder of this section details NIST’s implementation of each of the five elements in Fig. I.A.1, except for the “Transfer” component, which is discussed in Section II.C. Not present in the conceptual outline above are additional complications we introduce for ensuring physical-asset and virtual-asset security, high-availability, component redundancy, continuity across failures, and internal and external comparisons to estimate and reduce time scale noise and uncertainty. These issues are discussed in Sections I.A.2 through I.A.4.

I.A.1.2. Principal oscillator: TA(NIST), an ensemble of atomic clocks

NIST’s principal oscillator is a statistical ensemble of atomic clocks, a free-running atomic time scale called TA(NIST). Because the implementation of TA(NIST) is complex, we introduce several concepts first.

Besides operational reliability, a time scale’s principal oscillator is subject to two technical requirements: sufficient frequency stability, and programmability to achieve coordination with UTC. Frequency stability, quantified by metrics such as Allan deviation [8], can be translated into the expected holdover performance of the time scale between coordination/steering adjustments.

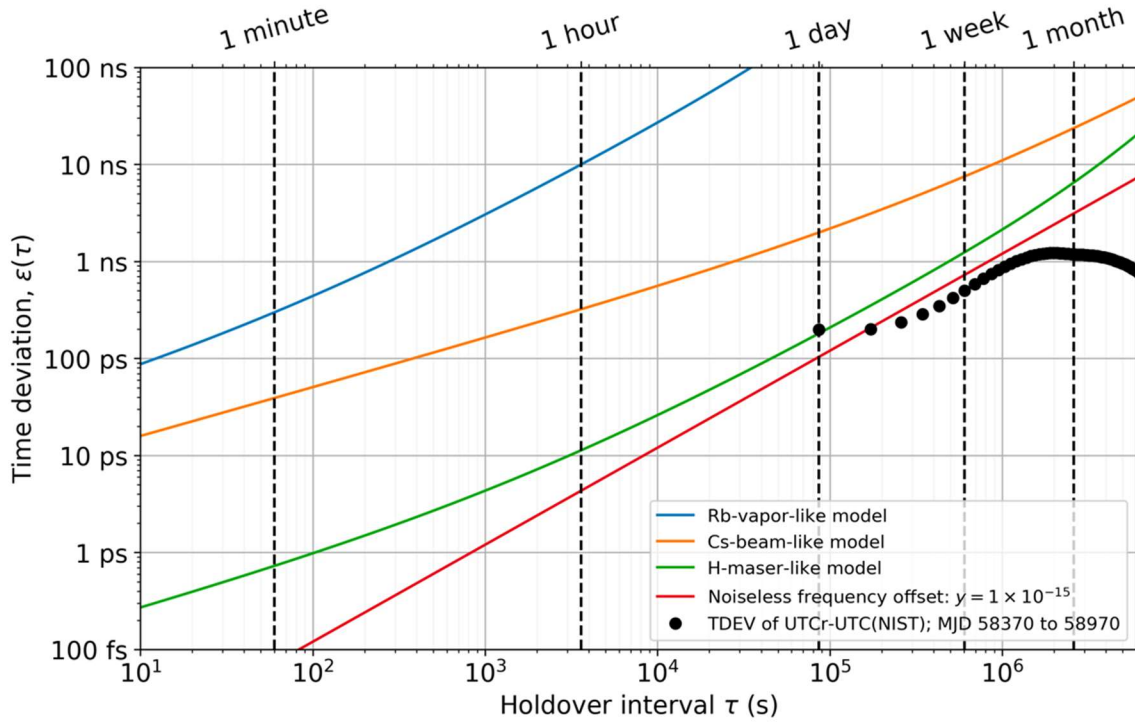


Fig. I.A.2. Here we explore the expected time dispersion (68% confidence level) of free-running atomic clocks, modeled with frequency instabilities $\sigma_y(\tau)$ derived from manufacturer datasheets, or a fixed frequency bias y (red curve). Black data points show a time deviation (TDEV) calculation of the daily data series UTCr – UTC(NIST) over a 600-day interval. At intervals longer than one day but less than one week, UTC(NIST) is less noisy than a single maser due to statistical averaging afforded by the free-running ensemble time scale, TA(NIST). At intervals longer than one week, the dispersion of UTC(NIST) is held low due to coordination actions performed approximately weekly. We note UTCr contributes noise to these data as well; over short intervals, the stability of UTC(NIST) is better than shown here.

Figure I.A.2 explores this relationship with models of three candidates for principal oscillators: a vapor-cell clock based on Rubidium, a commercial cesium-beam device (e.g. model 5071A, high-performance option) and a commercial hydrogen maser (e.g. model MHM-2010, low-phase-noise option). We highlight certain findings:

1. Typical holdover performance on the order of ≤ 10 ns over one-week intervals requires the use of high-performance cesium-beam devices or better. Performance on the order of 1 ns over one-week intervals requires the use of a hydrogen maser; no commercial alternative technologies with comparable performance and reliability currently exist.
2. In addition to low frequency instability, a principle oscillator’s particular frequency offset y must be characterized at the level of 10^{-15} or better to negligibly contribute to time dispersion at the desired level of 1 ns per week.
3. The time deviation (TDEV) [8] of historical data UTCr – UTC(NIST) shows that, over intervals of about one day, UTC(NIST) exhibits noise no worse than a single maser. Over longer intervals, UTC(NIST) performs somewhat better than a single maser. For intervals of about one week, this advantage is due to the ensemble algorithm used to create the principal oscillator TA(NIST). At longer intervals, $|\text{UTCr} - \text{UTC(NIST)}|$ remains small due to active coordination efforts: NIST

researchers use BIPM data and other observations to apply frequency steering corrections to UTC(NIST) approximately weekly.

Regarding the programmability requirement, the simplest principal oscillator is an atomic clock featuring a programmable frequency offset parameter (see Fig. I.A.3a). These devices are commonly available, because most atomic clocks include a synthesizer as part of their design. Atomic clock outputs are the fundamental tones or harmonics of high-stability, quartz crystal oscillators. A synthesizer, which often has a digital component, links the quartz oscillator frequency ($f_0 = 5$ MHz is nearly ubiquitous) to the device's atomic transition frequency (e.g., 1.420 GHz in a hydrogen maser, or 9.192 GHz in a cesium device) and allows for modulation techniques necessary for frequency locking.

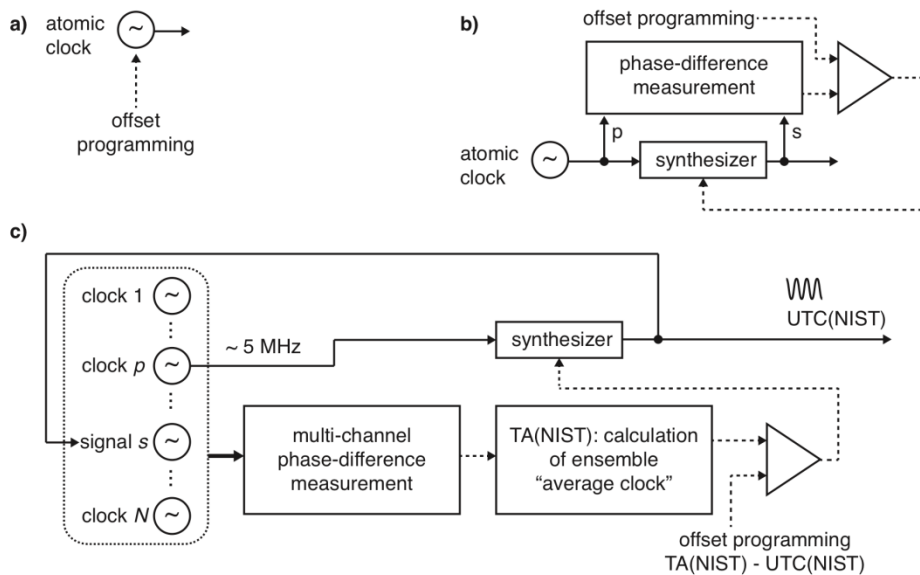


Fig. I.A.3. Pedagogical illustrations of principal oscillators of increasing complexity. a) Atomic clocks usually feature a programmable synthesizer as part of their implementation; external programming can adjust the phase and frequency of the clock's output signal. b) In order to observe and preserve the atomic clock's free-running behavior, a principal oscillator system can be designed by driving a low-noise, low-offset external synthesizer with the atomic clock's output, labeled p . Closed-loop control is achieved by phase-continuously measuring the phase evolution of synthesizer output s with respect to p . The difference $s - p$ can be driven to a programmable, time-dependent offset via the synthesizer's command interface. c) NIST implements the external-synthesizer approach with one additional complication: an ensemble "average clock" virtual signal TA(NIST) is computed such that TA(NIST) – clock i is defined for all clocks, including clock p . Together with an operator-defined relationship TA(NIST) – UTC(NIST), synthesizer commands are calculated to realize UTC(NIST) as a phase and frequency offset of clock p . Measurement of the synthesizer output as a pseudo-clock signal s allows for closed-loop control, though signal s receives no weight in the TA(NIST) ensemble.

An alternative to directly programming atomic clock offsets is to introduce an external synthesizer to the output of a free-running atomic clock; together the free-running clock and programmable synthesizer serve as the principal oscillator (see Fig. I.A.3b). This external synthesizer approach is favored by NIST for the following reasons:

1. A requirement of BIPM is that only observations of free-running (un-steered) atomic clocks may be submitted to contribute to the international atomic time scales EAL (Free Atomic Time) or TAI (International Atomic Time), and NIST wishes to make all of its atomic clocks available for such contributions (see Section I.A.5).

2. External synthesizers generally allow higher-resolution frequency adjustments, over a larger range, than the internal synthesizers found in atomic clocks. External synthesizers can also manage phase-continuous “phase slew” operations, where the output phase is smoothly adjusted by temporarily increasing or decreasing the programmed frequency offset for a fixed time interval, rather than requiring an immediate phase step.

Due to the additional signal distributions and phase-lock-loops involved, the addition of an external synthesizer does increase the principal oscillator’s noise level (noticeably only over short averaging intervals), the number of failure modes, and complexity. However, these disadvantages are outweighed by the additional operational flexibility and the preservation of the free-running atomic clocks’ behavior.

NIST implements this “external synthesizer” approach with one further complication: a statistical ensemble of clocks serves as the principal oscillator which is steered by the external synthesizer (see Fig. I.A.3c). The phases of several free-running atomic clocks are measured against a common, arbitrary oscillator by a multi-channel measurement system (MCMS). An algorithm known as AT1 (Section I.A.1.6) computes the principal oscillator TA(NIST) as data (not a physical signal), such that TA(NIST) – clock i is known for all clocks at all measurement epochs. Steering is achieved through programming a synthesizer whose reference oscillator is phase-locked to clock p . The signal, TA(NIST) – clock p , is combined with a defined relationship, TA(NIST) – UTC(NIST), to compute the appropriate synthesizer steering commands. The synthesizer output is measured in the MCMS as pseudo-clock s to allow for closed-loop control. Due to statistical averaging (and outlier mitigation) of the noise of its constituent clocks, TA(NIST) achieves better frequency stability and robustness than any single clock. A final advantage of this scheme: should clock p experience technical failure, repair of UTC(NIST) is accomplished by cabling a new clock p' to the synthesizer’s reference input and configuring software to substitute TA(NIST) – clock p' data in the coordination routine.

Figure I.A.4 illustrates NIST’s approach with more technical detail, though components related to redundancy and reliability are omitted for clarity (refer to Section I.A.2 instead). Commercially available, low-noise, low-offset synthesizers of the appropriate type are known as Auxiliary Offset Generators (AOG) or High-Resolution Offset Generators (HROG). The primary and redundant time scales in Boulder currently employ the model AOG-110 (Microchip), Fig. I.A.5, which offers a frequency offset range of 5×10^{-8} , frequency offset resolution of 1×10^{-19} , phase offset resolution of 1 ps, excess phase noise equivalent to $3 \times 10^{-13} (\tau / 1 \text{ s})^{-1}$, and a temperature-coefficient of phase offset $< 10 \text{ ps}/^\circ\text{C}$.

Figure I.A.4 also illustrates the components detailed in the next several sections. Section I.A.1.3 describes the atomic clocks and signal distribution elements. Further details about the MCMS follow in Section I.A.1.4. Data acquisition is performed by a Time Scale Computer (TSPC), covered in Section I.A.1.5. Details about TA(NIST) and its generating algorithm, called AT1, follow in Section I.A.1.6. Finally, the conversion of the AOG’s output at 5 MHz into a pulse-per-second signal, defined as UTC(NIST), is discussed in Section I.A.1.7.

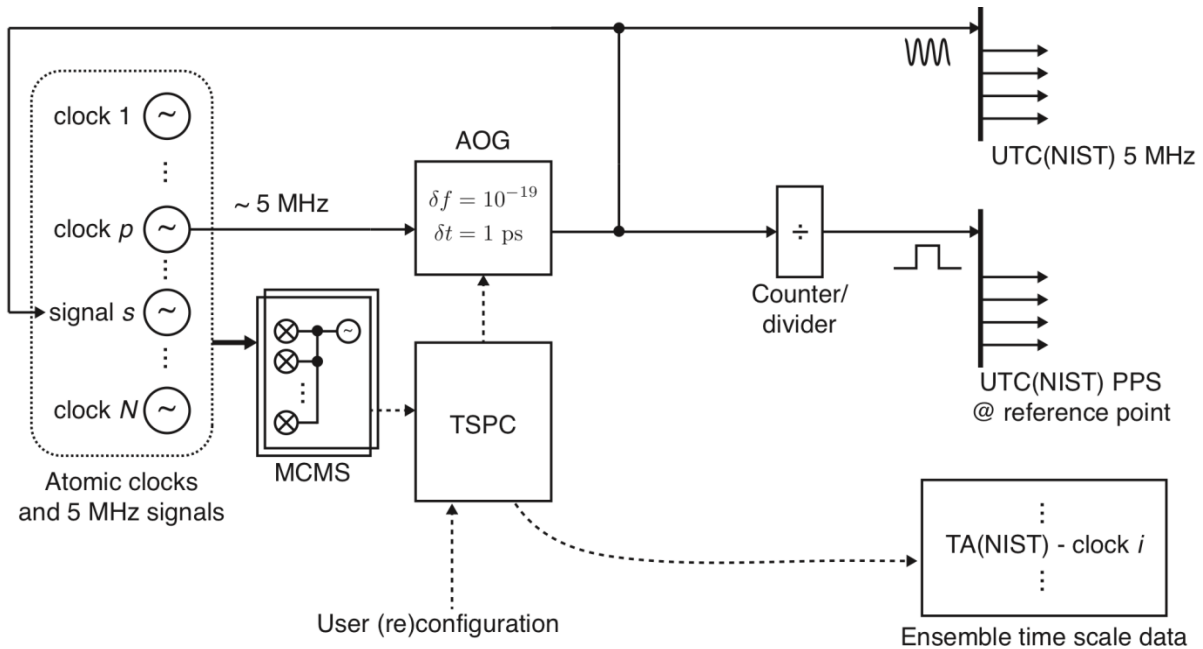


Fig. I.A.4. Illustration of the realization of the free-running atomic time scale TA(NIST) as data, sometimes called a real-time (not post-computed) “paper scale,” and the physical signal UTC(NIST), which is coordinated to approximate UTC. Components to ensure redundancy and reliability are omitted. For clarity, closed dots represent isolated distribution amplifier and cable routes. Solid lines indicate physical signal connections; dashed lines indicate data- or logical connections. Clock p is the source for the synthesizer (AOG), the output of which is measured (as pseudo-clock s) to allow closed-loop control but is not given weight in TA(NIST). One or more multichannel measurement systems (MCMS) produce phase-difference data of all clock signals to a data acquisition computer (TSPC). User input defines time scale parameters and the intended relationship between TA(NIST) and UTC(NIST); TSPC calculates the ensemble time scale TA(NIST) on a regular schedule and dynamically programs the AOG to realize UTC(NIST) as a phase and frequency offset from atomic clock p . The 5 MHz signal is distributed and divided to produce a pulse-per-second (PPS) signal, which is defined as UTC(NIST) at a particular distribution location.




Fig. I.A.5. An auxiliary output generator (AOG) is a high-resolution synthesizer that produces small phase and frequency offset signals derived from a reference atomic clock. The frequency resolution of this instrument is 10^{-19} . The phase adjustments have 1 ps resolution and are implemented as smooth “slews.”

I.A.1.3. Atomic clocks and signal distribution

At the time of writing, NIST operates 21 commercial atomic clock devices for the generation of the primary time scale UTC(NIST) in Boulder, Colorado. Of these, 14 are atomic hydrogen masers (Fig. I.A.6) and seven are high-performance cesium beam clocks. Table I.A.1 lists these 21 atomic clocks as well as six AOGs utilized by the NIST time scale. While the outputs of all atomic clocks are measured and tracked, currently 12 contribute to the primary ensemble TA(NIST). Others are withheld from participating for long periods following technical failure or major repair, and some contribute only to an independent, redundant time scale (see Section I.A.2). Because atomic clocks do occasionally experience technical failure, are repaired or replaced, or are physically moved, the roster contributing to TA(NIST) and its redundant backup implementations is always subject to change. However, due to the way TA(NIST) models and discounts every clock's time and frequency offsets, changes to the clock roster have minimal immediate impact on the frequency offset of TA(NIST).



Fig. I.A.6. Hydrogen maser in an environmental control chamber with its door open. The shelf above the maser holds a redundant DC supply (left) for the distribution amplifier and monitoring rack (right).

Table 1.A.1. Atomic clocks (ST: hydrogen maser, HP: 5071A cesium beam device) and time scale synthesizers (AOG) participating in NIST’s time scale (Primary) or redundant systems (Alternate, Contingency; see Section I.A.2). NIST clock identifiers are derived from serial numbers or order-of-acquisition by NIST. If clock data are submitted to the BIPM for contribution to the international time scales EAL/TAI, the related identifier is listed in the second column. Codes in the rightmost columns indicate the clock’s role, if any, in each time scale’s measurement and control system. M: measured in the MCMS. W: receives statistical weight in the AT1 calculation. S: serves as the time scale’s output synthesizer. P: serves as the synthesizer’s reference signal. I: serves as the MCMS NCO-tuning reference frequency (“channel 1”). C: mounted on a mobile cart. Information is accurate at time of writing but subject to change.

NIST identifier	BIPM identifier	NIST building	Roles in Primary time scale	Roles in Alternate time scale	Roles in Contingency time scale (“KGA”)
ST0022	1400222	B1	M, W	M, W	-
ST0003	-	B81	-	-	M, W
ST0004	1400004	B1	M, W	M, W	-
ST0005	1400205	B1	M, W	M, W, P	-
ST0006	-	B1	M	M	-
ST0007	1400207	B1	M, W, P	-	-
ST0008	-	B81	-	-	M, W
ST0009	-	B81	-	-	M
ST0010	1400210	B81	M, W	M, W	M, W, 1
ST0011	-	B81	-	-	M, W
ST0012	1400212	B1	M, W	M, W	-
ST0013	-	B81	-	-	M, W
ST0014	-	B81	M, W	M, W	M, W
ST0015	1412015	B81	M, W	M, W	M, W, P
HP1074	1351074	B1	M, W	M, W	-
HP2579	1352579	B1	M, W	M, W	-
HP2672	1352672	B1	M, W	M, W, 1	-
HP2792	-	B81	-	-	M, C
HP2857	-	B1	M	M	-
HP2935	1352935	B1	M, W, 1	M, W	-
HP2975	-	B81	-	-	M, W
AOG1	-	B1	M	M	-
AOG2	-	B81	-	-	-
AOG3	-	B1	M, S	M	M
AOG4	-	B81	-	-	M
AOG5	-	B81	-	-	M, S
AOG6	-	B1	M	M, S	-

Hydrogen masers benefit from careful environmental control of temperature and humidity. Therefore, masers participating in TA(NIST) are installed in independent environmental control chambers, decoupling

them from room/building fluctuations and each other. Similar closed-system environmental control is available for some of the cesium beam devices, including clocks sent to NIST for frequency calibration, though current 5071A units demonstrate low susceptibility to ambient environmental fluctuations. Further discussion of the environmental control continues in Section I.A.2.

The nominal output of all of NIST's atomic clocks is about +15 dBm $f_0 = 5$ MHz; frequency offsets are generally $|y| < 5 \times 10^{-12}$. A termination impedance of 50 Ω is standard for all clock outputs, amplifiers, and measurement system inputs.

All atomic clock signals are transported by high-isolation RF distribution amplifiers before measurement or transmission between rooms. The distribution amplifier (current model SpectraDynamics HPDA-5i) is designed for unity-gain, low compression below +18 dBm, galvanic isolation between input- and output-coaxial shield conductors, 140 dB isolation between all five outputs and the input, low distortion, and low temperature coefficient of delay (specification: 1.2 ps/°C, over 0 to 50 °C). The distribution amplifiers are housed in eight-slot chassis systems with redundant DC-powered backplanes.

Cesium-beam devices are generally located in the same room as the central measurement system. Given this fact, and the relative instability of these clocks, the cables transporting cesium clock signals to the measurement system are less critical: common, flexible coaxial cable types such as RG-58 and RG-142 are used. Since they are located in rooms adjacent to the measurement system, the cables transporting maser signals are longer. Given longer cables and the lower frequency instability of masers, double-shielded or solid-copper shielded cabling is used, such as RG-223 or RG-401.

All critical signals transported for distances ≥ 30 m demand low-loss cabling featuring solid-copper, or dual-layer shielding and expanded polyethylene foam dielectric for reduced temperature coefficient of delay: such as LMR-240, FSJ1-50A, or LDF2-50. Two maser signals are currently transported between buildings with amplitude-modulated fiber-optic devices (Menlo FiberLink).

I.A.1.4. Phase measurement system

All clock signals are transported to one of four electronics racks in the measurement room. Each clock signal is received by an isolating RF distribution amplifier; outputs from these are routed to at least two commercial multi-channel measurement systems (MCMS). Figure I.A.7 depicts two typical clock distribution route plans.

The aim of the MCMS (Fig. I.A.8) is to record a continuous record of relative (unwrapped) phase between several input clock signals. As configured by NIST, each MCMS (Microchip, model: TSC 12030; formerly manufactured by Microsemi, Symmetricom, and Timing Solutions) is capable of measuring 16 clock signals and consists of four elements:

1. Redundantly powered chassis (AC and temporary DC-backup power inputs; redundant internal regulated DC voltage supplies).
2. An industry standard architecture (ISA) bus and embedded industrial computer featuring limited solid-state storage (operating system: FreeBSD 4.1, 6.2, or 8.3 depending on firmware revision) and 10 Mbps or 100 Mbps Ethernet for configuration and data recording.
3. Chassis modules housing a free-running 32 MHz quartz oscillator, and a four-channel numerically-controlled oscillator (NCO, a synthesizer based on direct digital synthesis principles).

- Four chassis modules capable of measuring four clock signals each via the dual-mixer time-difference (DMTD) high resolution measurement scheme.

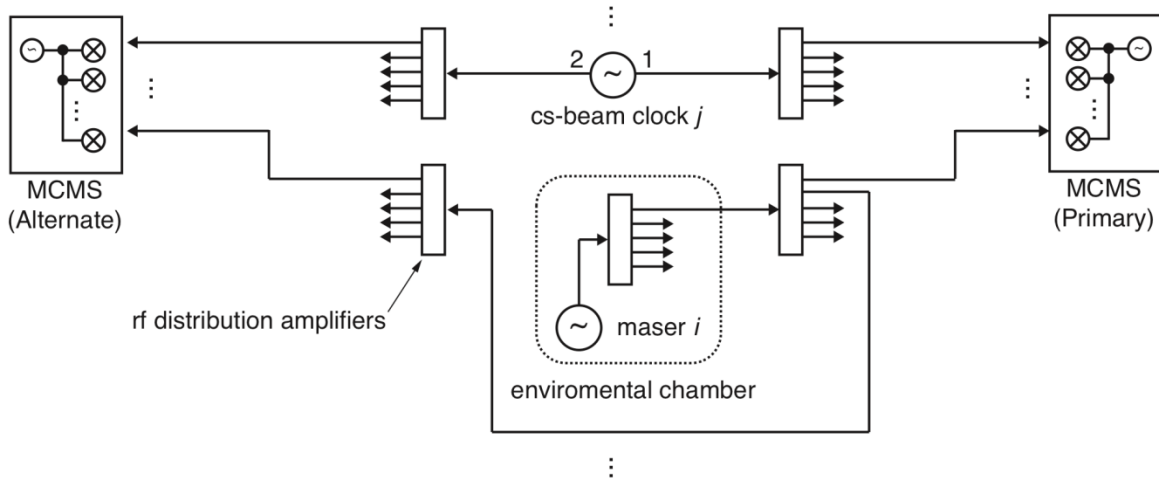


Fig. I.A.7. Typical signal routing for cesium beam devices and hydrogen masers. Both buffered 5 MHz outputs (labeled 1 and 2) of some cesium clocks are routed to Primary and Alternate phase measurement systems (MCMS) via independent RF distribution amplifiers. For all masers, and some cesium clocks, the Alternate signal path is derived from the Primary distribution amplifier as shown. Masers' outputs are buffered within their environmental chambers by ground-isolating distribution amplifiers before the (long) cable routes to the measurement room. The figures are drawn for clarity but not to scale. Primary and Alternate measurement racks are physically close; masers are distributed in adjacent rooms.




Fig. I.A.8. Photograph of a MCMS chassis with 16 clock measurement channels. Modules from right to left: industrial embedded computer, four measurement modules implementing DMTD, distribution amplifiers and sources for the 32 MHz onboard quartz oscillator counting time base and NCO synthesizer signals.

The principle of dual-mixer time-difference (DMTD) phase measurement is well described elsewhere [9]. Briefly, the technique aims to make use of the limited temporal resolution of a time interval counter (TIC)

by first translating phase variations of the $f_0 = 5$ MHz clock signals to a much lower heterodyne frequency, e.g. $f_h = 10$ Hz, in a coherent balanced mixer which preserves the phase-difference (in radians) between clock signals under test. In general, the resolution gain offered by DMTD is the ratio of $f_0 / f_h \approx 5 \times 10^5$. Practically, lower limits on f_h occur due to the presence of flicker-noise processes and the desire for a suitably rapid fundamental data rate, which is proportional to f_h .

Figure I.A.9 illustrates a simplified version of the DMTD technique as implemented in the MCMS modules. The on-board 32 MHz quartz oscillator drives the NCO synthesizer module to create a frequency, f_{nco} , roughly $f_h = 10$ Hz below the nominal atomic clock frequency $f_0 = 5$ MHz, near 4,999,990 Hz. The f_{nco} signal is mixed with two clock signal inputs using double-balanced RF mixer (DBM) circuit elements. Each mixer is followed by a low-pass filter to select the 10 Hz mixer product. These are amplified by zero-crossing detectors (ZCD) into waveforms with high-slew-rate “trigger” features near 10 Hz. Each of these triggers drives a “start” or “stop” input of a TIC, which uses the 32 MHz quartz oscillator as its counting reference time base.




Fig. I.A.9. Simplified diagram of a single channel pair in a single MCMS module, presented to illustrate the principle of measurement. LFP: low-pass filter. ZCD: zero-crossing detector and amplifier. NCO: numerically-controlled oscillator. XTAL: free-running quartz crystal oscillator. See text for full description.

In Fig. I.A.9, we consider the voltage waveforms of clocks 1 and 2 differing only in their amplitude and phase:

$$V_1(t) = V_1 \sin(2\pi f_0 t + \phi_1), \quad (\text{I.A.1})$$

$$V_2(t) = V_2 \sin(2\pi f_0 t + \phi_2). \quad (\text{I.A.2})$$

We define their instantaneous *time-difference*, x_{21} , as:

$$x_{21} = \frac{\phi_2 - \phi_1}{2\pi f_0}. \quad (\text{I.A.3})$$

By convention, the time-difference x_{21} uses “unwrapped” phases ϕ_2 and ϕ_1 (phases are not bound to remain in a single unambiguous angle domain $-\pi \leq \phi_i < \pi$). RF mixing preserves the (wrapped) phase difference $\phi_2 - \phi_1$ while exactly removing the common evolving phase $2\pi f_0 t$; phase due to the NCO signal is also common to both down-converted tones near f_h . The TIC measures an interval δt , which can be related to the clock phase difference as:

$$\delta t = \frac{\phi_2 - \phi_1}{2\pi f_h}. \quad (\text{I.A.4})$$

Combining the above equations, we find we can obtain the clock time-difference by scaling δt by a tiny ratio (which accomplishes the promised resolution enhancement):

$$x_{21} = \delta t \frac{f_h}{f_0}. \quad (\text{I.A.5})$$

Generalizing to clocks with relative frequency offsets (or frequency drift), the treatment is the same: the time-difference simply becomes a linearly (or quadratically) time-dependent quantity. In general, phase noise components with Fourier frequencies much slower than the rate f_h are well-sampled by the DMTD scheme; faster variations are subject to aliasing or attenuation by filter elements (the approximate bandwidth in the MCMS is 500 Hz).

The MCMS computer and firmware, consisting of field-programmable gate arrays (FPGA) within the modules communicating with a multithreaded computer process known as “utimer”, collate the results of the many TIC measurements into a phase estimate of each clock input with respect to a common oscillator (f_{nco} , a product of the 32 MHz quartz crystal). The MCMS reports each clock's monotonically increasing time-difference (with respect to the MCMS's common NCO) in units of f_0 periods ("cycles"). Because f_h is kept near 10 Hz, each channel's measurement appears to increase by approximately 10 cycles per second. Time-differences between the clocks, in units of seconds, are obtained by subtracting these reported cycle counts (removing the NCO's common influence) and multiplying the result by the quantity 200 ns (one period of $f_0 = 5$ MHz).

Channel 1 of each MCMS is treated somewhat specially. The MCMS assumes that channel 1's signal represents an accurate $f_0 = 5$ MHz frequency reference; the measurement of channel 1's phase evolution with respect to the quartz oscillator is used to retune the NCO programming every five minutes, keeping the frequency $f_0 - f_{nco}$ near 10.0 Hz. Because NCO retune events are phase-continuous and simultaneously applied to all the modules' DMTD circuits, they have no net influence on clock-difference measurement data produced by the MCMS. Owing to these NCO adjustments, no voltage or temperature tuning of the 32 MHz quartz crystal is required. To produce any data at all, the MCMS requires channels 1 and 2 to be occupied by signals near 5 MHz at sufficient amplitude. Noise and a small frequency offset y of channel 1 is of no consequence because all later processing steps are based on difference signals, clock 1 – clock i .

The data from two MCMS chassis are processed together to yield 32 channels of coherent clock phase comparisons, though the method to do so differs for different MCMS firmware versions. Briefly, in one implementation (currently used in the Primary time scale; see Section I.A.2.2), the MCMS can be programmed to act as a ‘supervisor’ or ‘client’ chassis. Identical clock signals are input into channel 1 of both units. The supervisor and client communicate over a private network; the supervisor collates (and indexes) all measurement data from both units. In this mode, the data acquisition system must communicate only with the supervisor chassis. In a second implementation (currently used in the Alternate time scale), two MCMS units operate independently, though both are also given identical clock signals to their respective channel 1 inputs. In this case the data acquisition system must communicate with both MCMS units, and collate their data streams in software. Both methods produce measurement records for up to 32 clocks, where channels 1 and 17 are understood to be linked. Additional considerations are discussed in the section on redundancy, I.A.2.2.

I.A.1.5. Data acquisition

A dedicated time scale personal-computer (here called a “TSPC”) maintains a persistent network connection to either one supervisor MCMS chassis, or two independent MCMS chassis, together accounting for 32 measurement channels. The connectivity of the TSPC machine is shown in Fig. I.A.10.

A TSPC has six responsibilities:

1. Serve as the “time-of-day” reference (labeling measurement epochs) for associated MCMS systems via the network time protocol (NTP).
2. Maintain and provide an operator interface for updating model parameters of the clock ensemble: clock names and their associated MCMS channel numbers, “cable delay” constants, initial estimates of clock frequency offsets, drift offsets, and so on.
3. Sample and record clock phase-difference data produced by the MCMS hardware.
4. Periodically (currently every 720 s, synchronized to UTC midnight and each 12 minutes thereafter) run an implementation of the AT1 algorithm to compute the ensemble time scale TA(NIST).
5. Accept programmatic adjustments which relate TA(NIST) into the coordinated time scale UTC(NIST). Realize UTC(NIST) as a physical signal by computing and sending programming to a synthesizer (AOG) whose reference is a particular ensemble member clock (clock p in Fig. I.A.4).
6. Send, promptly via a private network interface, all phase measurement data and AT1 computation outputs to a public-networked computer (here, “Hub PC”) for purposes of backup, general analysis, and testing for alarm conditions.

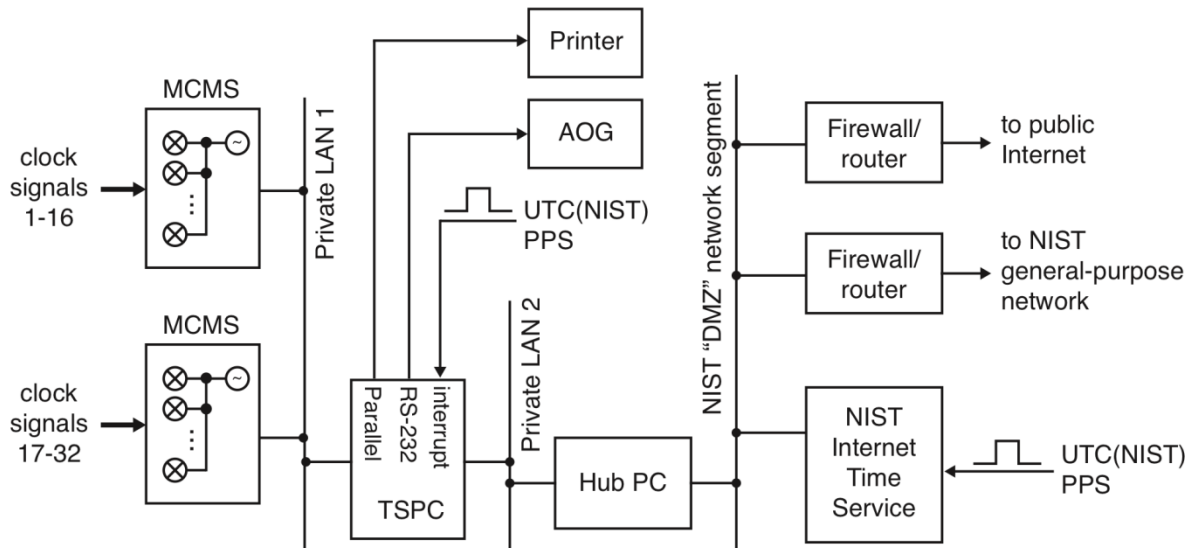


Fig. I.A.10. Simplified connectivity diagram of the data acquisition machine (TSPC) to the clock phase measurement systems (MCMS) via Ethernet, synthesizer (AOG) via RS-232 serial, a parallel-port line printer, and data repository Hub PC via Ethernet. Local area networks (LANs) are isolated and non-routed. At cold start, TSPC’s system time is obtained by network messages to Hub PC, which uses the NIST Internet Time

Service as an authoritative source. In steady-state operation, TSPC's system time is disciplined by the UTC(NIST) PPS input to a serial interrupt interface pin. The Hub PC shifts data onto NIST's general-purpose network for data analysis and archiving

TSPC machines are built from commodity PC hardware, generally replaced every four to six years, usually coinciding with operating system upgrades (currently FreeBSD 10.3, i386 architecture type). The absolute minimum TSPC hardware requirements are not known but are very modest; minimum practical storage is perhaps 30 GB (though more allows for less logfile/backup truncation), minimum practical RAM is perhaps 2 GB. Any modern x86 CPU is sufficient as interface latency is the practical bottleneck. The interface requirements are:

- two RS-232 serial ports
- one 25-pin parallel port
- two Ethernet network interfaces (typically satisfied by use of a PCI/PCIe expansion card)
- one compatible keyboard, mouse, and display

After formatting the TSPC storage device, the operating system (OS) and custom software are deployed from verified optical media or universal serial bus (USB) flash memory images without ever connecting the TSPC system to the Internet or any public network. No software packages or libraries are required that are not part of the default OS installation. All custom software is deployed as C code which is compiled and tested on each TSPC machine. Two important modifications are made to the FreeBSD kernel on TSPC machines: the applicable serial port driver is edited to allow for faster interrupt capture on the control pins used for PPS input, and the kernel code for controlling the system time is modified to disallow control by standard OS components like *ntpd*.

Upon cold start, custom software called *lockclock* exchanges network messages between the TSPC and Hub PC over private LAN 2 to establish the date and time-of-day (approximately 1 ms accuracy) in the standard UNIX time scale (defined as the number of seconds since January 1, 1970). A PPS signal derived from UTC(NIST) is input to a TSPC RS-232 control pin (e.g., DCD, CTS, or RI). The pulses trigger interrupt routines which capture the system time at microsecond-level accuracy. The *lockclock* process monitors the stability and offset of the PPS signal with respect to the UNIX time-of-day obtained via network exchanges with Hub PC. After its stability and accuracy is confirmed, *lockclock* changes mode and begins using the PPS signal to implement a phase-locked loop (PLL) so that TSPC system time is directly locked to UTC(NIST).

Lockclock corrects the apparent frequency offset of the TSPC's local quartz crystal oscillator; therefore, good holdover is achieved should either the network or UTC(NIST) PPS signal suffer brief interruption. Network messages to the Hub PC periodically confirm stable behavior of the PPS and TSPC system time; even small deviations (several milliseconds) trigger visual and printed alert messages. A logfile that records system time comparisons, transitions of *lockclock* PLL mode, and any anomalies, is kept indefinitely.

As shown in Fig. I.A.10, the MCMS units exist on a private network ("LAN 1") with no direct connectivity to the Hub PC; instead, the associated TSPC machine is assigned to be a MCMS "time-of-day" reference via standard NTP utilities.

The MCMS industrial computer relies on its system clock to label the acquired clock time-difference data. Noise or bias in these labeling timestamps can, in principle, add noise to the realization of UTC(NIST). Given that the observed measurement noise of the MCMS hardware averages as $\sigma_y(\tau) \leq 1 \text{ ps } (\tau / 1\text{s})$, and that the frequency offsets of NIST atomic clocks are on the order of 10^{-12} , the excess noise that would result from a poorly synchronized MCMS mislabeling data timestamps is kept negligible when the system time

dispersion is less than about 100 ms. When tested using additional NTP clients, MCMS units in steady-state operation typically demonstrate synchronization to UTC(NIST) within 2 ms. During cold start events, MCMS firmware tests for system time stability before producing any output data.

The TSPC provides an interactive user interface for status display, implemented in the X Windows environment. Input commands are entered as plain text with a strict syntax. Most commands must be executed by physical presence at the keyboard and general remote access to the system is not allowed.

The TSPC communicates with the assigned MCMS system(s) via a persistent Transport Control Protocol (TCP) line-oriented streaming text connection. Following each time scale computation (currently every 720 s), the TSPC briefly connects to the Hub PC to report recent results. In both connections, TCP guarantees data correctness through check-summing, correct received packet ordering, and automatic retransmissions when necessary, or at least produces a software error code when these conditions cannot be met. We discuss the consequences this has on reliability and holdover behavior in Section I.A.2.2.

I.A.1.6. Free-running time scale algorithm AT1, time scale output control

NIST's principal oscillator, TA(NIST), is a free-running ensemble time scale computed using the AT1 algorithm. The AT1 algorithm is described in detail elsewhere [10]; here we summarize its design goals and implementation.

AT1 seeks to synthesize TA(NIST) as a reliable and stable clock signal. Ideally, TA(NIST) performs better than its best constituent oscillators and is not made worse by including contributions from its worst performing constituent oscillators. Like many ensemble time scale approaches, AT1 takes advantage of favorable conditions for signal averaging, and the high stability of atomic clocks even when they lack absolute frequency accuracy. AT1 uses its independent constituent clocks as a basis for outlier detection and attenuation. It then rewards the most stable clocks with increased weight when computing TA(NIST).

Averaging reveals each clock's predictable behavior. Over sufficiently short averaging intervals, isolated atomic clocks tend to exhibit uncorrelated frequency noise which is approximately Gaussian ("white" in spectral density). The noise type is a direct consequence of the clocks' interrogation of independent quantum absorbers (atoms) as frequency discriminators. The quantized nature of the interrogation results in a Poisson-like random process which is transduced into white (random) frequency modulations. AT1 is designed to iteratively compute a moving average, y_i , of each clock's apparent frequency offset, where the averaging interval is made large enough to benefit from Gaussian noise averaging, but not so large that non-stationary noise processes become significant. While the averaging interval is determined by staff for each clock, generally the averaging interval is approximately 30 hours for hydrogen masers and five months for cesium beam clocks. By manual analysis, operators also periodically update AT1's model of each clock's frequency drift rate, d_i , most relevant for hydrogen masers. Readers should refer to [10] for a full description of how the clock models evolve as data is acquired.

Each clock is modeled to discover its predictability. Along with the time difference determination, x_i , AT1 models each clock's frequency y_i , and linear drift rates, d_i , with respect to TA(NIST); in fact, the computed record of x_i values define TA(NIST) with respect to each clock i . Following Eq. (I.A.3) above, MCMS measurements of clock time differences (with respect to the clock input to channel 1), $x_{i1} = x_i - x_1$, are sampled at regular intervals Δ (at present, $\Delta = 720$ s) aligned with 0 h UTC. Between some measurement epoch T and the next epoch $T + \Delta$, the clock model parameters $x_i(T)$, $y_i(T)$, and $d_i(T)$ are used to predict each clock's time difference with TA(NIST), $\hat{x}_i(T + \Delta)$, where

$$\hat{x}_i(T + \Delta) = x_i(T) + y_i(T) \Delta + \frac{1}{2} d_i(T) \Delta^2. \quad (\text{I.A.6})$$

When MCMS data arrives, and $x_i(T + \Delta)$ is computed (discussed later in this section for presentation clarity), the quantity

$$\epsilon_i(T + \Delta) = \hat{x}_i(T + \Delta) - x_i(T + \Delta), \quad (\text{I.A.7})$$

is recorded as clock i 's time-prediction error. A long moving-average (typically 7 to 15 days of averaging for masers; 31 days of averaging for cesium beam clocks) of the magnitude $|\epsilon_i|$ forms a model of expected time prediction error per measurement, per clock, σ_i . In terms of time-domain analysis functions, σ_i is similar to the time deviation (TDEV) [8] of clock i , sampled over an interval $\tau = \Delta$, discounting the deterministic effects of y_i and d_i .

Predictable clocks receive more weight in TA(NIST). One of the clocks is arbitrarily chosen as a “pivot” clock, r . MCMS data is linearly combined to express the measured time differences x_{ri} for all clocks i with respect to the pivot clock r . The time difference of the pivot clock r with respect to the ensemble is then computed as a weighted average:

$$x_r(T + \Delta) = \sum_j w_j(T + \Delta) [\hat{x}_i(T + \Delta) + x_{ri}(T + \Delta)], \quad (\text{I.A.8})$$

where each coefficient w_j is the fractional weight of clock j with the properties

$$w_j(T + \Delta) \propto 1/\sigma_j^2(T), \quad (\text{I.A.9})$$

and,

$$\sum_j w_j = 1. \quad (\text{I.A.10})$$

In other words, Eq. (I.A.9) ensures that clocks with records of low prediction error (small σ_j) contribute more to the determination of x_r . Equation (I.A.10) normalizes the weights so their sum is unity. The summation terms in Eq. (I.A.8) combine each clock's modeled behavior, \hat{x}_i , with the most recent measured data, x_{ri} , a concept common to all methods of filtering. Once x_r is determined following (I.A.8) and after completing the error handling procedures described below, all other ensemble-clock time differences can be calculated using the measured data x_{ri} :

$$x_i(T + \Delta) = x_r(T + \Delta) - x_{ri}(T + \Delta). \quad (\text{I.A.11})$$

Equation (I.A.11) even holds for the case $i = r$, where the measurement $x_{rr} = 0$. In this sense, the choice of pivot clock r is made arbitrary.

Administrative limits control nonlinear behavior. As Eq. (I.A.9) shows, predictable (stable) clocks are rewarded with increased weight coefficients w_j . As a consequence, the ensemble time TA(NIST) becomes more correlated with these high weight clocks. This positive feedback mechanism can lead to a superior clock gradually acquiring $w_i \approx 1$, leaving TA(NIST) vulnerable to sudden failure modes of a single clock. To prevent this outcome, we sacrifice some of the theoretical benefits of purely statistical averaging and enforce an administrative limit on all the weight coefficients:

$$w_i \leq 0.3. \quad (\text{I.A.12})$$

The calculations above, from Eq. (I.A.7) onwards, repeat recursively if and when this limit is applied.

Outlier discovery and suppression; the AT1 “reset” procedure. If a clock’s time prediction error ($|\epsilon_i|$) far exceeds its expected magnitude (σ_i), AT1 applies a variance test to decide the clock’s status as an outlier. If $|\epsilon_i| \geq 4\sigma_i$, clock i is subject to a “reset” procedure with the following consequences. For the contemporaneous ensemble calculation only, $w_i \rightarrow 0$, and modeled clock parameters like y_i are frozen; they do not receive updates. In the marginal case, $3\sigma_i < |\epsilon_i| < 4\sigma_i$, the clock’s effective weight is linearly decreased from the computed value in Eq. (I.A.9) such that $w_i \rightarrow 0$ as $|\epsilon_i| \rightarrow 4\sigma_i$. In either case, all calculations from Eq. (I.A.7) are repeated recursively with the new weight conditions applied. After each recurrence, a minimum of four clocks must remain with finite weight, otherwise TA(NIST) enters a holdover state.

The record of a clock in the “reset” condition is still subject to Eq. (I.A.11) following the time scale computation; in other words, outlier data is not “removed” or “corrected” automatically. Because AT1 is computed as a “real time” time scale, clock records are unalterable. However, at some future measurement epoch, NIST staff can introduce compensation corrections manually to “repair” the large time steps. Typically, this is only done when the root cause of the malfunction can be determined to be technical and not a random statistical fluctuation of the atomic clock. For example, a faulty cable will result in “reset” events with ϵ_i nearly equal to exact multiples of 5 MHz periods (multiples of 200 ns) due to “missing” oscillation cycles; these are easy to identify and “fix” by introducing compensating model adjustments to x_i data following a cable repair. Manual compensation corrections, from the perspective of AT1, also appear as large time prediction errors; they will cause a benign “reset” condition to occur for one measurement cycle. Therefore, should corrections become necessary for more than one clock, they are usually applied during distinct measurement intervals to avoid needlessly shrinking the population of clocks contributing to TA(NIST) in a given cycle.

TA(NIST) conclusion. In summary, TA(NIST) (a “virtual” signal or “paper clock”) is defined at all measurement epochs t in terms of its time difference $x_i(t)$ with respect to any clock i in the NIST ensemble. TA(NIST) is a statistical weighted average of the clocks, discounting their deterministic frequency offsets and linear frequency drift rates, where the more stable clocks receive increased weight. Due to the “reset” algorithm, large outlier deviations from malfunctioning clock signals do not affect TA(NIST) significantly. Should we wish to realize TA(NIST) as a physical signal, conceptually, the frequency offset $y_p(t)$ could be programmed into a synthesizer (AOG) referenced to clock p , while the time-prediction error signal $\epsilon_p(t)$ provides phase offset programming. If the output of the AOG itself is measured in the MCMS (labeled as clock s in Fig. I.A.4), any excess deviations seen in clock s can be directly removed with additional phase offset programming. NIST implements these ideas, but we do not realize TA(NIST) as a physical signal. Instead, the programming described above is combined with additional frequency and phase offset information to realize UTC(NIST) in the AOG output.

I.A.1.7. Generation of the UTC(NIST) reference signal and subsequent distribution

We defer a discussion of how UTC is created to Section I.A.5. The BIPM monthly publication Circular-T informs NIST of the official determination of the time difference UTC – UTC(NIST), provided at 5-day intervals. The BIPM provides a weekly series UTC_r – UTC(NIST) which approximates UTC – UTC(NIST). UTC_r data are noisier but are provided at 1-day intervals and are available sooner after measurement. Figure I.A.11 illustrates some recent history of both signals.




Fig. I.A.11. Approximately four years of recent data (mid-2016 to mid-2020), of UTC – UTC(NIST) (available at 5-day intervals) and UTCr – UTC(NIST) (available at 1-day intervals). UTC – UTC(NIST) (blue) represent the realized time offset of UTC(NIST) with respect to UTC, as observed by BIPM.

Using BIPM data and other observations (for example, direct comparisons with other metrological laboratories, GPS carrier-phase observations, and local or remote comparisons with atomic frequency references), NIST staff decide by consensus what frequency offset ought to be applied between TA(NIST) and UTC(NIST) to satisfy two abstract goals: reduce the absolute magnitude of the future time error UTC – UTC(NIST), and preserve or increase the frequency stability of UTC(NIST). In general, due to the divergent nature of clock noise and time transfer noise, both goals cannot be simultaneously achieved optimally over all averaging intervals. Nonetheless, NIST staff demonstrate a routine of weekly frequency adjustments which keeps the root-mean-square deviation of UTC – UTC(NIST) < 3 ns and the frequency instability of UTC(NIST) generally below 2×10^{-15} . Over long averaging intervals, such as years, both deviation measures are much lower. See Section I.A.5 for further discussion.

Figure I.A.12 illustrates a long history of NIST’s coordination actions: the frequency adjustments made to TA(NIST) in order to realize UTC(NIST). Also plotted are BIPM’s independent rate observations of UTC – TA(NIST), made approximately one month after the coordination actions are implemented.

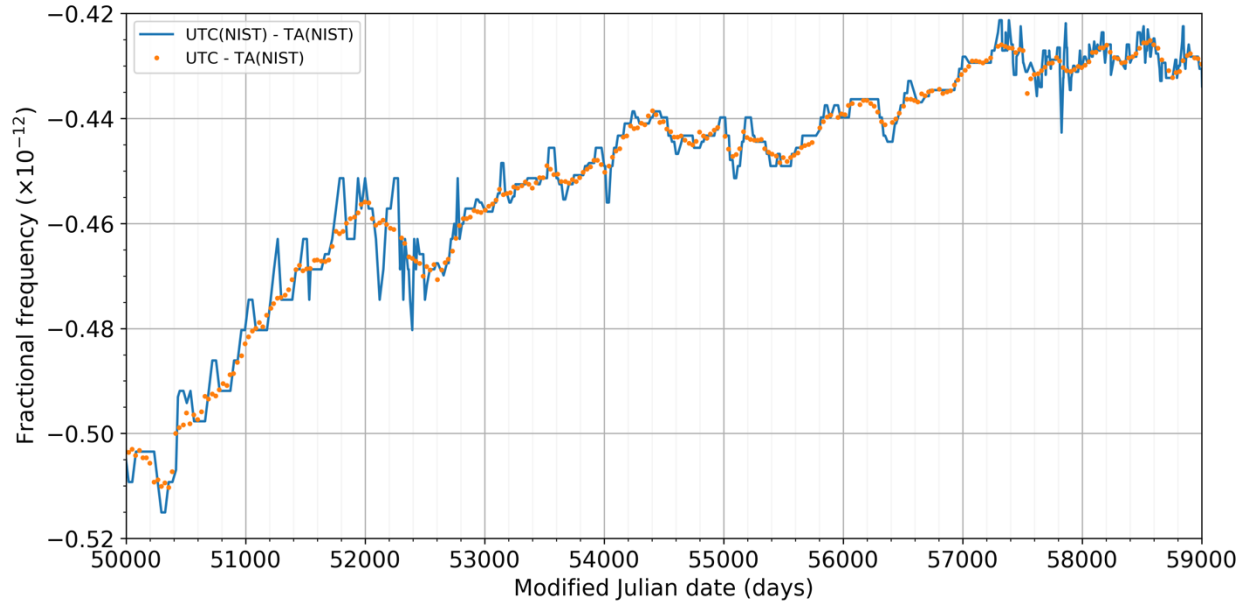


Fig. I.A.12. This series of the rate offset UTC(NIST) – TA(NIST) (blue) represents the frequency steering of TA(NIST) to create UTC(NIST), NIST’s realization of UTC. The series UTC – TA(NIST) (orange) is generated by BIPM roughly one month after measurement. Relative agreement between these two series validates NIST’s steering method. Note, BIPM time transfer recalibration occurred near MJD 57500, explaining a transient deviation in the UTC – TA(NIST) data.

In Fig. I.A.4, the output of the AOG synthesizer at 5 MHz is a physical expression of the rate of UTC(NIST): a count of 5 million cycles denotes a time interval of 1 s in UTC(NIST). However, by convention, the absolute phase of this 5 MHz signal does not directly define the time scale UTC(NIST) anywhere. Instead, the signal is routed to a counter/divider device also known as a PPS generator (currently, SpectraDynamics PPS-2RM-B1).

In normal operation, the PPS generator employs a zero-crossing detector/amplifier and digital circuit to count 5 000 000 input cycles before triggering an output pulse (specified rise/fall times of 1.5 ns, high and low levels of 4 V and 0.1 V, respectively, into 50 Ω loads) and resetting the counter. After a cold start, or when the output pulse needs (re)aligning to another reference pulse edge, a manually-interlocked “SYNC” pulse input is provided. When the interlock is engaged by the operator, the digital cycle count is set to 0 by the next observed SYNC PPS rising edge. This function is occasionally used to align redundant time scale implementations (see Section I.A.2.2), or to realign the Primary time scale following a failure. Discounting any differential cable delays, this procedure aligns the output and SYNC pulses to within ± 100 ns (the total range, 200 ns, is a single period of 5 MHz); closer alignment must be accomplished by appropriate phase steering of the 5 MHz AOG signal with TSPC commands.

Following the counter/divider device, the generated PPS pulse is received at the input of a pulse distribution amplifier (PDA). By convention, a particular output connector of this PDA is defined as the observation/reference location for UTC(NIST); rising voltage edges there represent the beginning of each SI second in UTC(NIST), and counting such pulses indefinitely defines the continuous time scale.

All continuous users of the UTC(NIST) PPS signal derive their signals from other channels of this reference PDA, or from channels in a hierarchy of cascaded PDAs, through well-shielded cabling with additional delays ranging from tens to thousands of nanoseconds. The cabling delays must be calibrated by any of several techniques: traveling clock comparisons, traveling time-transfer receiver comparisons, or by direct measurement of delay using a TIC with access to both ends of the cable-of-interest. In general, such

calibration techniques can achieve delay uncertainties ranging from 50 ps to about 200 ps depending on approach and link lengths. With differential techniques, and with ideal control of instrument and ambient conditions, these uncertainties could perhaps be reduced to as small as 10 ps.

I.A.2 – Methods Ensuring Continuous Operation of Primary Time Scale Architecture

NIST strives for (and has achieved) continuity of UTC(NIST) across equipment upgrades, prolonged power outages, room/building floods, and failures in environmental control, active and passive components, cables, computers, and networks. This is accomplished with a combination of monitoring/alarms, “failsafe-by-design” and “defense-in-depth” engineering guidelines, and operation of redundant systems. As an art, reliability engineering is one of balance. Reliability enhancements typically increase operational complexity, and NIST’s capability to supply trained operators is finite. In this section, we first outline these design elements, followed by an analysis of some failure modes in detail.

I.A.2.1. Alarm systems and on-call staff

Two computers, located in different buildings (B1 and B81), running two different operating systems (FreeBSD and Windows 10), and administered by two different staff members, monitor the environmental sensors, clock data, and time scale validation data for several defined alarm conditions. The implementation details of the two machines differ, but the primary alarm mechanism for each is a “dialout-only” modem (attached to an “analog” telephone line, not VoIP) that calls a dedicated phone number which is multiplexed to several of the staffs’ personal mobile devices, laboratory phones, and NIST-owned mobile devices. At least one trained staff member is scheduled to be “on-call” to respond to such alarms at all times, a duty that typically rotates biweekly. Secondary alarm mechanisms are also employed by both computers: short message service (SMS) text messages and email (both dependent on the availability of NIST’s Internet connection) are also available to staff, and generally these messages contain additional information about the alarm condition. Both alarm computers test their alarm daily: one does so in the mid-morning, the other in the afternoon. Staff members who are “on-call” are trained to notify the alarm computer operators if these calls (and secondary emails, etc.) are not received.

The B81 alarm computer implements an additional safeguard designed to detect a crashed or powered-off alarm computer: a peripheral interface controller (PIC) placed between the computer and modem will automatically dial the modem if its “watchdog timer” is not reset by a periodic message issued by the functioning alarm computer’s software. This feature has successfully alerted staff to automatic reboots caused by NIST’s forced-deployment of operating system upgrades. The B1 alarm computer is never subject to such reboots. Table I.A.2 lists some of the alarm conditions handled by each alarm computer system.

Table I.A.2. List of conditions (environmental sensors, data availability, time scale measurements) which are actively monitored and recorded (code: M) and alarmed to on-call staff (code: A) by the two independent alarm systems in B1 and B81. Refer to Fig. I.A.13 and Fig. I.A.14 for diagrams of the three monitoring TICs and their named PPS signals.

Conditions	B1 monitoring/alarm	B81 monitoring/alarm
B1 time scale rooms, temperature/humidity	M, A	M, A
B1 TWSTFT/GPS room temperature/humidity	-	M, A
B1 utility/generator power status	M, A	M, A
B1 cesium clock environmental chamber	M, A	-
B1 maser environmental chambers	M, A	M, A
B1 maser health telemetry	M (four B1 masers)	M, A (two B1 masers)
B81 time scale rooms, temperature/humidity	-	M, A
B81 TWSTFT/GPS room temperature/humidity	-	M, A
B81 utility power status	-	M, A
B81 maser environmental chambers	-	M, A
B81 maser health telemetry	-	M, A (all B81 masers)

B81 maser and monitor redundant power supplies	-	M, A
Primary MCMS or TSPC halt; holdover condition	M, A	M
Alternate MCMS or TSPC halt; holdover condition	M, A	M
Contingency (“KGA”) MCMS or TSPC halt; holdover condition	M, A	M
Primary, Alternate, Contingency: clock has successive “resets” in AT1 algorithm	M, A	-
Primary, Alternate, Contingency: large AOG steering commands by AT1 algorithm	M, A	-
TIC A measures large Primary-Alternate PPS offset	M, A	
TIC B measures large UTC_REGEN-KGA PPS offset	M, A	
TIC C measures large offsets of UTC_REGEN vs. any of: UTC_COAX, UTC_WR, Alternate PPS, KGA	-	M, A
Any single clock developing an apparent cycle offset (> 200 ns) between the Primary and Alternate time scale systems	M	M, A
Daily alarm test “dialout” function	M, A	M, A
B81 alarm computer (or client monitoring computers) crashes	-	A
B1 alarm computer crash	-	M, A (ping check)

I.A.2.2. Failsafe by default

A key feature of the TA(NIST) design is that if active control halts, the noise performance of UTC(NIST) merely degrades to the noise performance of a single, frequency-compensated hydrogen maser. Were this condition to persist for a whole day, the time scale’s dispersion would increase by about 100 ps (see Fig. I.A.2): a level probably undetectable to other observers or users of the time scale. Nonetheless, we aim to detect and alarm on such “holdover” conditions rapidly (≤ 24 minutes) and to fix the problem within two hours of notice.

This failsafe holdover property is relied upon by several parts of the system design. For example, if the network connection between the MCMS and the TSPC is interrupted for any reason (e.g. the MCMS experienced a software fault and restarted), the system enters a holdover state until manual action can restart it under careful observation. This safely prevents the system from automatically processing malformed data from a possibly-malfunctioning MCMS. Similarly, if a TSPC needs to be rebooted for any reason, UTC(NIST) safely exists in a holdover state until the custom TSPC software can be restarted by a specialist.

I.A.2.3. Redundancy and reliability

Redundancy provisions of the NIST time scale can be analyzed in the popular “PACE” framework originally used to describe redundant communication plans. PACE is an acronym for: Primary, Alternate, Contingency, Emergency.

In our usage, the Primary system is employed to continuously produce the best publicly available UTC(NIST) signals and data. An Alternate system is operated continuously to produce indefinitely indistinguishable alternative UTC(NIST) signals and data in case of failure of the Primary, either automatically or semi-automatically by use of physical or software-switches. Wherever possible, we attempt to design non-common points of failure between the Primary and Alternate systems.

A Contingency system is operated continuously and is designed to exhibit similar noise performance (at least for short holdover intervals) as the Primary/Alternate UTC(NIST) but without being completely indistinguishable to outside observers. Even fewer points-of-failure are common with the Primary/Alternate systems. Should a Contingency system need to be put into broadcast service, additional time-to-repair

(compared to a pre-engineered switch to the Alternate) will be required as some physical hardware/cabling must be re-arranged.

Finally, an Emergency system is one which will exhibit inferior noise performance compared to the others, but ought to have zero common points-of-failure including some amount of geographic diversity. Depending on the context of the failure, which is presumed to be catastrophic given elimination of Contingency systems, a much larger time-to-repair may be required to broadcast an Emergency system's output as UTC(NIST).

The remainder of this section describes our PACE approaches; refer to Fig. I.A.13 and Fig. I.A.14 for illustration.

Primary and Alternate systems overview – Table I.A.1 lists the atomic clocks with statistical weight (at time of writing) in the primary ensemble TA(NIST). A parallel implementation with duplicated hardware components exists to produce TA(NIST)_{alt} continuously. The atomic clocks comprising the Primary ensemble also comprise the Alternate ensemble, but independent MCMS hardware is interrogated by an independent TSPC over independent network devices for data acquisition. An independent AOG is programmed to realize a 5 MHz UTC(NIST)_{alt} signal; the Alternate AOG is referenced by a different clock than the Primary AOG. An independent counter/divider converts the Alternate AOG output into a PPS realization of UTC(NIST)_{alt}. Two hardware switches, one for the RF signals at 5 MHz and one for the PPS signals allow operators to route either the Primary or Alternate outputs to all broadcast channels and downstream users. During normal operation, the Alternate system receives additional, occasional frequency steering commands meant to maintain a time offset of < 50 ps between TA(NIST) and TA(NIST)_{alt}. When any large deviations are detected, this automatic steering is halted to preserve the independence of TA(NIST)_{alt} as a viable backup.

Contingency systems overview – As shown in Fig. I.A.13, a third AOG is deployed whose output is divided into a PPS signal. This Contingency AOG is sourced by a maser not used by either Primary or Alternate AOGs. It is also not automatically programmed by a TSPC, so without manual operator programming, linear frequency drift of its source maser will degrade it as a stand-in for UTC(NIST). The PPS switch can select this Contingency AOG as a source, but operators must first update the AOG programming with recently determined frequency and time offsets of the source maser. There is no fixed position for the RF switch to use the Contingency AOG; operators would have to connect new cable routes. Use of the Contingency AOG can provide UTC(NIST) broadcast signals of sufficient quality for most, but not all, downstream users. For example, it does not replace the functions of collecting or transmitting clock offset data to the BIPM which is necessary for NIST to contribute to EAL/TAI. Therefore, use of the Contingency AOG shall be temporary as Primary and Alternate systems are repaired/replaced.

A second Contingency system was recently created at NIST which can both propagate UTC(NIST) signals and continue NIST's contributions to the BIPM, albeit with a different set of atomic clocks. The system is referred to internally as KGA (as it exists in B81, also named the Katherine Gebbie building). Table I.A.1 lists its member clocks. Its construction is similar to that of the Primary and Alternate systems; see Fig. I.A.14 for an illustration.

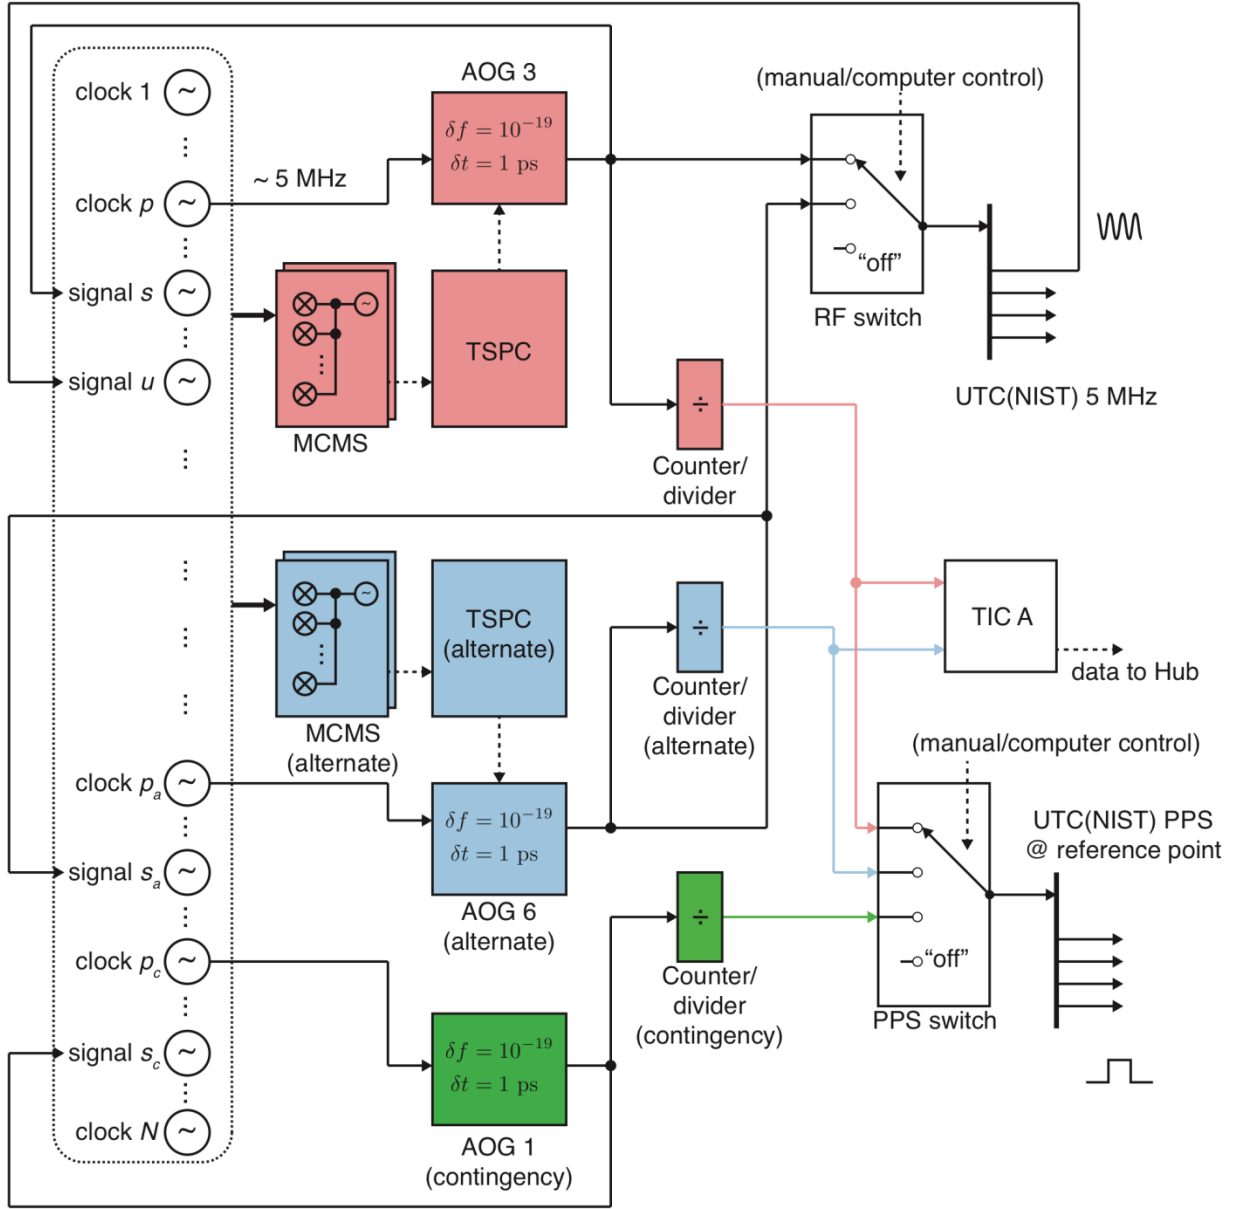


Fig. I.A.13. Partial diagram of the redundant time scale arrangement in building B1, see text for full description. For clarity, solid circles stand for isolated distribution amplifier routes. Solid lines indicate 5 MHz signals; up to the PPS switch, colored lines indicate pulse-per-second (PPS) signals (red: primary, blue: alternate, green: contingency). MCMS: multichannel measurement systems. TSPC: time scale personal computer. AOG: auxiliary offset generator (low-noise, low-offset synthesizer). Counter/divider units derive pulse-per-second signals from 5 MHz inputs. TIC: time-interval counter, used for measuring online agreement of the Primary and Alternate time scale outputs. Pseudo-clock signals labeled s , s_a , s_c , and u do not receive weight in the ensemble average but are measured to achieve and verify closed-loop lock operation. Clocks labeled p , p_a , and p_c serve as the inputs to AOG synthesizers and do generally receive statistical weight in the time scale.

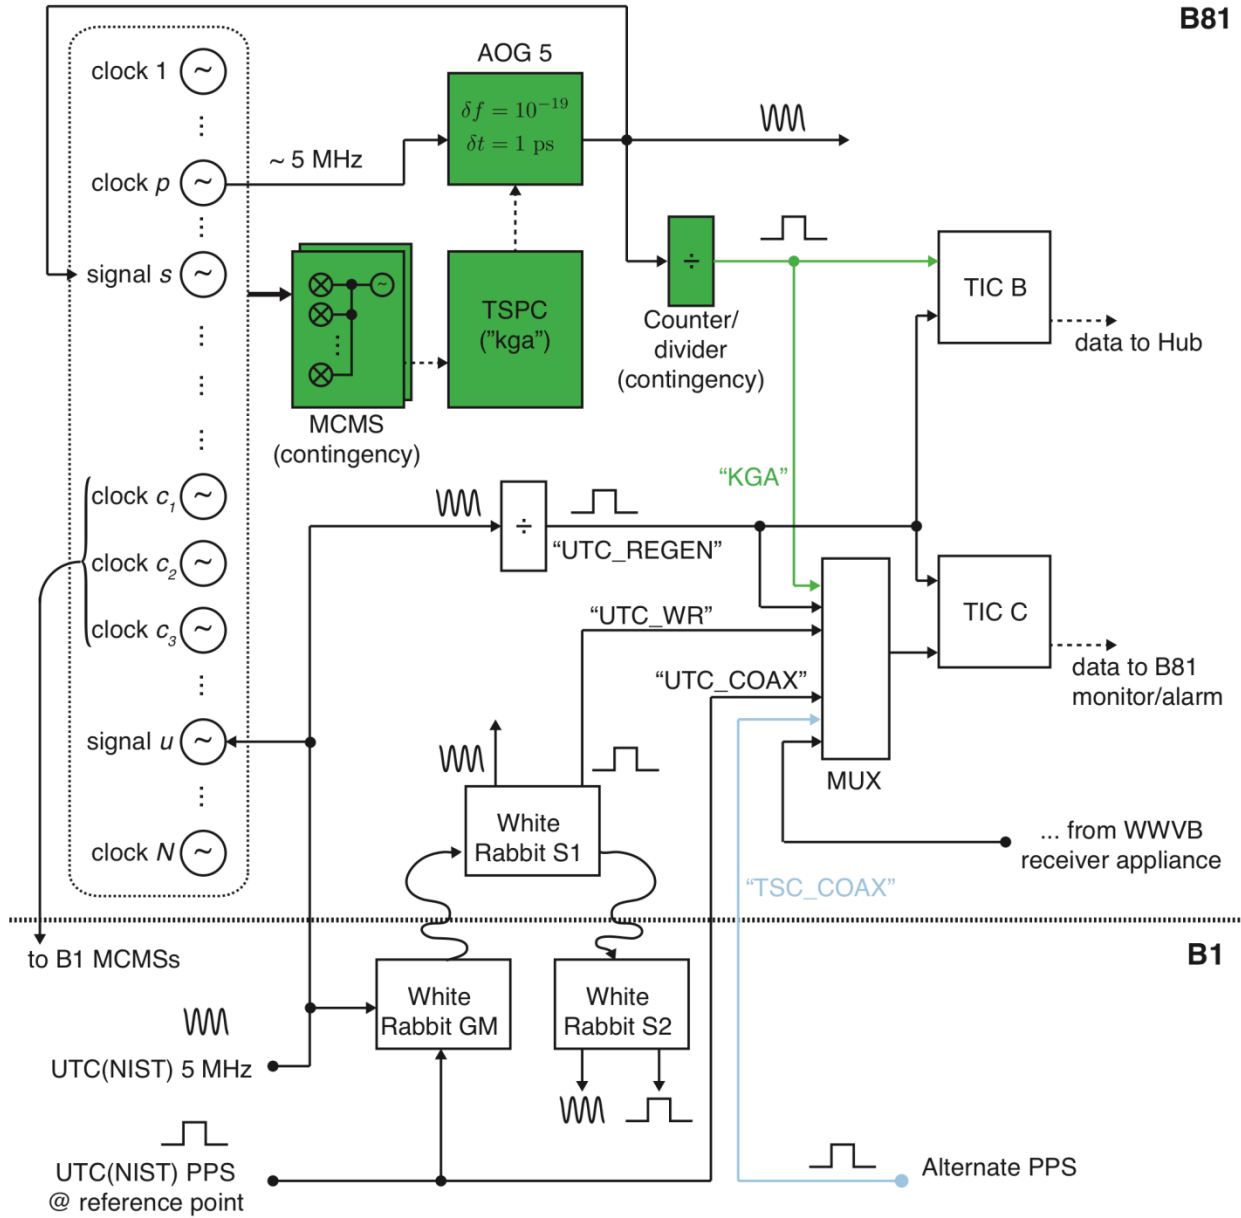


Fig. I.A.14. Partial illustration of the Contingency time scale (“kga”) system in B81. We note that three 5 MHz clock signals (currently NIST masers ST10, ST14, and ST15) are transported to B1 over coaxial cable and fiber optic links to be redundantly measured in the Primary and Alternate time scale systems. Likewise, various representations of UTC(NIST) are transported from B1 to B81 to be compared against KGA products. The simplest, but noisiest link simply transports the UTC(NIST) PPS over 330 m of coaxial cable (“UTC_COAX”). The lowest noise, most reliable link is known as UTC_REGEN; it is made by transporting 5 MHz UTC(NIST) and regenerating it into a PPS, synchronized (to the nearest 5 MHz period) to UTC_COAX. The Alternate time scale PPS is transported over coax (“TSC_COAX”), for redundant validation. White Rabbit devices, discussed in Section II.A.1, deliver delay-compensated UTC(NIST) signals (known as “UTC_WR”) over telecom fiber optics. TIC C compares several PPS signals to UTC_REGEN using a multiplexer switch (MUX); samples are acquired every minute. All UTC(NIST) signals besides “UTC_WR” experience about 1400 ns of cable delay.

Emergency systems overview – A byproduct of the AT1 algorithm is an evolving model and record of the time and frequency offsets of all clocks participating in the ensemble. If any single clock is able to survive some catastrophic failure that caused all other clocks to fail (e.g. a prolonged, unmitigated power outage with failure of backup power systems), TA(NIST) can in principle be restarted using the continuity of a single clock, with an expected penalty of holdover dispersion of the magnitude shown in Fig. I.A.2. Likewise, any clock whose time difference with respect to UTC(NIST) is continuously known through comparisons, could, in principle, be used to “restart” the time scale. These include high-performance cesium-beam clocks located at NIST WWV/WWVB radio station in Fort Collins, CO (Section I.B), high-performance cesium-beam clocks located at staff offices at JILA, University of Colorado, Boulder, high-performance cesium-beam clocks located at NIST, Gaithersburg, Maryland (Section I.B), or low-performance cesium beam clocks that serve as validation references at NIST in Boulder but are not part of the time scale.

I.A.2.4. Clock health telemetry

With the aim of increasing reliability, all 5071A cesium beam devices, MHM-2010 hydrogen masers, and AOG synthesizers measure several internal physical quantities, the data record of which can aid in device troubleshooting or in predicting an approaching end of service life. We obtain this “health telemetry” either by manually reading and recording data from a digital display on the device, or automatically via isolated serial data connections.

5071A cesium beam clocks – Experience has shown that generally, 5071A cesium beam clocks operate through their expected beam tube life (8 to 10 years) relatively free from operator intervention. Several quantities are casually monitored by staff for anticipating the end of service life, typically caused by extinguishing the supply of cesium in the tube’s thermal effusive oven.

While the operating principle of the 5071A is beyond the scope of this document, we briefly describe the physical basis for the critical telemetry data. Cesium atoms exiting the microwave-interaction region of the clock with the correct nuclear spin orientation are guided, via a strong magnetic field gradient, into a charge multiplier/detector. A slow control loop adjusts the multiplier’s high-voltage, known as the quantity *E-mult*, to optimize the strength of received and amplified atomic signal against “background” signal which rises with amplification as well. An *E-mult* driven to its maximum, 2553 V, forecasts the last few years of service life. Subsequently, an electronic gain parameter, *Gain*, slowly increases to maintain clock operation; when it nears its maximum of 58%, very little tube life remains. Other signals indicate the state of the tube’s ion pump current (indicative of residual gas pressure or severe vacuum faults), oven and voltage-controlled oscillator heater and control voltages, microwave signal level, magnitude of current responsible for the magnetic bias field, and a software latch signaling any loss of power, frequency lock, or other fatal condition.

MHM-2010 hydrogen masers – Though hydrogen masers boast a longer service life (> 20 years) than cesium beam devices, their construction is much more complicated. NIST’s experience is that various sub-components can experience occasional failures that greatly compromise the maser’s clock performance. Fortunately, several indicative signals are measured and reported by the maser. Without providing a complete explanation of the maser’s design, we provide Table I.A.3 to highlight several important signals and which failure modes each may indicate.

Table I.A.3. Partial list of hydrogen maser health telemetry signals studied by NIST staff to diagnose/predict unhealthy atomic clock operation. All data (32 channels) are recorded for permanent storage; erroneous values for some signals will directly trigger the B81 alarm system (code: A). Maser signal channel numbers are relevant for internal tracking purposes.

Maser telemetry signal	Telemetry channel(s)	Code	Notes
IF amplitude	0, 26	A	Proportional to hydrogen maser signal strength; low values could imply vacuum system fault, cavity mistuning, microwave electronics problems, or loss of excited hydrogen atoms. Very low values indicate the clock output has become an uncontrolled quartz oscillator.
Cavity register	1, 28	A	Tuning voltage on a varactor responsible for centering the cavity response on the atomic transition frequency. Sudden jumps can accompany moderate frequency steps.
VCO control voltage	2, 27	A	Value exceeding the range ± 1.0 V indicates increased risk of the maser losing frequency lock of the voltage-controlled quartz oscillator. Every few years, manual tuning of a potentiometer can return this value to its preferred range.
Hydrogen ion pump and upper ion pump currents	3,4	-	Indicates vacuum integrity. Old ion pumps can exhibit large pump currents due to field emission caused by electrode decay.
LO amplitude	5	-	Measures the signal strength through the microwave synthesis chain. Low values can indicate electronics failures.
Various heater voltages	6,7,8,9,12,13,25	-	Can indicate loss of temperature control, change in temperature gradients, or heater failure.
Source discharge current and collector voltage	10, 18		Analyzed together, can detect changes in the plasma discharge designed to prepare excited atomic hydrogen for injection into the maser bulb/cavity.
Magnetic field	14	-	Measurement of the applied magnetic field along the maser cavity axis.
Source pressure, control, and heater voltage	15,16,17	-	Voltages related to the palladium hydrogen leak source.
Battery voltage and charging current	19, 11		Monitors the float voltage and (normally low) trickle charging current of the onboard 24V (44 Ah) backup battery bank.
Main power bus	20	-	Indicates the main DC voltage bus powering all internal electronics.
Cavity average	21	-	During normal operation, this value varies randomly. If stationary, may indicate cavity control circuit is in "hold" state (normally only appropriate for testing or diagnostics).
AC-DC output 1, AC-DC output 2, external DC output	22,23,24		Respectively, these indicate the DC output voltage of the three redundant power supplies (other than the backup battery bank).
On battery	31	A	Indicates all redundant power supplies have failed and the maser is discharging its onboard backup batteries.

AOG-110 synthesizers – The core of the AOG synthesizer is a voltage-controlled quartz oscillator offset-phase-locked to an external reference signal. The status of this phase lock, including the value of the control voltage, is available on the AOG front panel or by serial communication. The quartz oscillator's environment is made highly stable through two layers of temperature control; voltages related to these oven heaters are similarly displayed. Finally, information about the RF signal level at various points in the synthesis chain, and about the redundant sources of power, is also periodically observed by staff.

I.A.2.5. Environmental control

In general, environmental variation affects clocks more than non-clock devices (such as cables, distribution amplifiers, and measurement systems). Non-clock devices do exhibit a phase or delay dependence on ambient conditions, but these phase and delay deviations are generally reversed when ambient fluctuations

revert to normal. In contrast, ambient conditions can introduce clock frequency perturbations that integrate into time errors and that persist after ambient control is reestablished.

Hydrogen masers are significantly perturbed by ambient variations, their frequency deviations are on the order of 10^{-14} per $^{\circ}\text{C}$ and 10^{-15} per %RH, though the exact response varies significantly per clock. Empirically, cesium beam devices appear to be less significantly perturbed by such variations. Therefore, NIST hydrogen masers are typically decoupled from room temperature and humidity variation by operating them inside of sealed “reach-in” chambers capable of precise temperature and humidity control. With one exception, the chambers are current or earlier versions of the model SP-Hotpack 33HR (Scientific Products, Inc.). Interior construction is stainless-steel with most shelves removed. The chambers are cooled by closed-cycle split refrigeration, exchanging heat into room air. Precise thermal control is achieved with a 1.2 kW ballast heater, driven by a process control servo (Watlow EZ-Zone or comparable). Interior temperature set points are chosen to be near nominal room temperature or at an observed minimum maser frequency sensitivity (20°C to 25°C); over long averaging intervals, observed stability is better than 0.25°C . In the current revision, humidity (near 50 % RH, long-term stability $< 1\%$ RH) is stabilized by a digital process controller driving a heater band clamped onto a steam-generation tank. Float-valves or float-relays regulate water level in the bath, and dual-stage filtration is installed on the back of each chamber to pre-treat the incoming potable water supply. In several cases, we have replaced a chamber’s failed humidity system with an ultrasonic humidification system driven by a digital dead-band controller; performance is similar with less risk of plumbing/overflow leaks, no need for water filtration, and much less power consumption.

All chamber temperatures and humidity levels are monitored by the B81 alarm computer; samples are taken every 10 minutes. The monitoring sensors are distinct from the sensors driving the chamber’s process controllers. Monitoring sensors are measured by dedicated microcontrollers, which digitize the results and produce them on command to the B81 alarm computer over an RS-485 serial bus. As an example, Fig. I.A.15 shows a few months of recorded maser chamber environmental variables.




Fig. I.A.15. Example of temperature and humidity monitoring of a maser environmental chamber (nickname “Donna”; houses NIST maser ST10 in B81). Large transients can be caused by periods where service is performed on the chamber, periodic “defrost” procedures (labeled “A0” on plot) or when room environmental fluctuations are too large for the chamber to fully compensate. By melting accumulated ice, the defrost procedure alters temperature gradients in the chamber, likely explaining some of the shift seen here. Note that the sensors responsible for this data are distinct from the sensors driving the chamber’s control loops.

I.A.2.6. Redundancy in power supply

All critical time scale components receive power from multiple sources; no time scale asset should lose power due to the loss of electric power on the NIST campus. Table I.A.4 lists various categories of time scale assets and their modes of Primary, Alternate, and Contingency power plans. Not listed in the table, a shared Alternate power plan is provided by NIST’s site-wide mains electricity switch, which ideally can rapidly actuate between two independent metropolitan utility feeds. Common emergency plans for all categories include manual reconfiguration of battery banks (where possible) and emergency selected use of NIST-operated life support generator systems that normally are not burdened with time scale loads. Hydrogen masers are governed by additional emergency triage provisions allowing most functionality to be disabled to preserve vacuum integrity under battery power for extended outages.

Table I.A.4. List of Primary, Alternate, and Contingency power plans for various categories of time scale assets in buildings B1 and B81. ATS: automatic transfer switch provides devices with AC power via relays which select from between two independent sources of AC power. UPS: uninterruptible power supply. SLA: sealed lead acid battery. FLA: flooded lead-calcium acid battery (chosen for long service life in standby capacity).

Asset type	Inputs	Primary power	Alternate power	Contingency power
B1 cesium clocks	1 AC, 1 DC	UPS provides AC power (with a few, legacy exceptions)	Generator/ATS powers UPS or AC input	Large FLA battery bank supplies DC input
B1 masers	2 AC, 2 DC	Generator/ATS powers AC input 1	UPS, backed by generator/ATS, powers AC input 2	Onboard SLA batteries specified for eight hours runtime
B1 chambers	1 AC (208 V)	Utility power	Generator/ATS	None; manually open door to cool clocks
B1 room control	AC, 3-phase	Utility power	Generator/ATS powers: maser room B’s primary, maser room A’s backup, measurement room’s backup	Generator/ATS powers portable air conditioners, room-fans.
B1 RF distribution amplifiers	2 DC	Online battery charger powers DC input 1	Different online battery charger powers DC input 2	Large FLA battery banks supply respective DC inputs
B1 MCMS chassis	1 AC, 1 DC	UPS provides AC power	Generator/ATS powers UPS	Large FLA battery banks supplies DC input
B1 TSPC, networking, monitoring	1 AC	UPS provides AC power	Generator/ATS powers UPS	Primary/Alternate use different UPSs; switch to working unit
B1 AOG and PPS generation, distribution	1 AC, 1 DC	UPS provides AC power	Generator/ATS powers UPS; large FLA battery banks supply DC inputs	Primary/Alternate time scales use different UPSs; switch to working UPS
Hub PC	2 AC	UPS provides power to AC input 1	Different UPS provides power to AC input 2	Generator/ATS powers UPSs
B81 cesium clocks	1 AC, 1 DC	UPS provides AC power	Rack-mounted ATS selects between two UPSs	Generator/ATS powers UPSs; (DC battery banks for DC input is planned)
B81 masers	2 AC, 2 DC	UPS powers AC input 1	Utility-grade UPS powers AC input 2	Generator/ATS powers all UPSs; onboard SLA batteries specified for eight hours runtime
B81 chambers	1 AC (208 V)	Utility power	Utility-grade UPS powers chamber (~ one hour)	Generator/ATS powers all UPSs

B81 room control	AC, 3-phase	Utility power	Utility-grade UPS powers backup room heat pumps	Generator/ATS powers UPS; can deploy portable air conditions if necessary
B81 RF distribution amplifiers	3 DC	Linear AC-DC converter 1 powered by UPS	Linear AC-DC converter 2 powered by UPS	Generator/ATS powers UPSs; can attach batteries or other supply to DC input 3.
B81 MCMS chassis	1 AC, 1 DC	UPS provides AC power	Rack-mounted ATS selects between two UPSs	Generator/ATS powers UPSs; (DC battery banks for DC input is planned)
B81 TSPC and networking, monitoring hardware	1 AC	UPS provides AC power	Rack-mounted ATS selects between two UPSs	Generator/ATS powers UPSs
B81 AOG and PPS generation, distribution	1 AC, 1 DC	UPS provides AC power	Rack-mounted ATS selects between two UPSs	Generator/ATS powers UPSs; (DC battery banks for DC input is planned)

I.A.2.7. Preventative maintenance

Here we detail provisions made for regular maintenance and upkeep of NIST time scale components and laboratory assets.

Room environmental control – The B1 rooms (maser room A and measurement room) were retrofitted with high-precision laboratory air conditioning units that rely on closed-loop chilled water for heat exchange. NIST’s campus Central Utility Plant (CUP) operates and maintains the chilled water loop. Should the CUP lose the ability to maintain chilled-water flow and temperature, thermostatic valves located adjacent to maser room A automatically switch over to city-serviced potable water as the heat exchange medium, which is drained instead of circulated. Maser room B is cooled by a refrigerated split-system air conditioner. Annually reevaluated contracts with the campus Office of Facilities and Property Management (OFPM) specifies maintenance for these air conditioning units including: filter changes, inspections, scheduling repairs when needed, and cleaning of outdoor components.

Compared to B1, B81 features more uniform environmental control. All time scale B81 rooms feature generator-backed heat pumps, thermostatically controlled to switch on if primary cooling fails. All B81 building system monitoring and maintenance is the responsibility of OFPM. Full shutdowns are scheduled approximately annually for equipment inspection and replacement. The B1 and B81 room temperature and humidity excursions trigger our alarm systems if conditions are severe.

Chamber environmental control – The hydrogen masers and some cesium beam clocks are decoupled from room temperature/humidity fluctuations by being installed in independently-controlled environmental chambers. Maser chambers are manually defrosted on a 90 or 120-day schedule. Water filter cartridges are replaced on an as-needed basis; the filters feature a color-changing band which indicates end-of-service life. Their refrigeration units are serviced only as needed. We have replaced some chamber’s humidification systems with tank-based ultrasonic-based humidifiers; the water levels are checked biweekly and filled as needed. The chamber enclosing cesium clocks is a NIST-built apparatus. Airflow is measured and monitored; fans are replaced as needed. Cooling is provided by a refrigerated water/alcohol reservoir, where heat is exchanged into room air; this fluid level is checked biweekly and refilled when necessary.

Backup power generators – Annual maintenance contracts with OFPM specify the upkeep of the B1 diesel (80 kW), B1 natural-gas (10 kW), and B81 diesel (180 kW) generators, including starter-battery testing/replacement, air and fuel filter replacement, fuel refilling, and arranging engine service when appropriate. All generators are tested weekly for 30 minutes (without load transfer). The time scale staff

checks the B1 generators' "hour meters" to verify that they increase by 0.5 hours weekly, and OFPM monitors the tests of the B81 generator. NIST schedules site-wide power shutdowns approximately yearly; these afford an opportunity to test and monitor the generators under loaded conditions.

Uninterruptable power supplies (UPS) and standby batteries – At the time of writing, the NIST time scale employs and tracks over 360 batteries, deployed among 68 uninterruptable power supplies (UPS) or as standby batteries attached to individual hydrogen masers. All are replaced on a staggered two or three-year preventative maintenance schedule, according to manufacturer's recommendations. Except for six UPS units located next to the B1 diesel generator, most batteries are kept in air-conditioned spaces that tend to prolong standby battery life. At least two spare batteries specified for use with hydrogen masers (SLA1161 type; 12V, 44Ah capacity) are kept in B81 on trickle chargers. As part of "on-call" duties, staff members examine all UPSs and battery chargers at least weekly for error/warning indicators.

A large array (combined energy capacity of about 60 kWh) of lead-calcium flooded battery cells are kept on continuous trickle chargers in B1; these serve as DC backup power for several cesium clocks, MCMS chassis, AOG synthesizers, and distribution amplifier rack chassis. In a prolonged emergency outage, the cell banks can be reconfigured as necessary to triage assets.

The B81 building provides generator-backed utility-scale UPS for several specially-marked outlets. At its present loading, this UPS can provide about one hour of run time before starting the B81 generator.

Multichannel measurement systems (MCMS) and data acquisition systems (TSPC) – Each MCMS chassis features redundant cooling fans, redundant ac-to-dc on-board power supplies, and an indicator that signals the health of the attached battery backup power. The status of each of these is checked by staff at least weekly. MCMS machines must be kept informed about any announced leap seconds so that correct UTC time stamps can be associated with all clock data across the leap second. This is implemented by a *leap-seconds.list* file following the NTP convention [11]. Every six months, typically in June or December (conventional months at the end of which leap seconds are adjusted), staff update the *leap-seconds.list* file stored on each MCMS. Even if no new leap seconds are announced, the so-called "expiry date" field must be kept "in the future" to avoid MCMS error messages and possible failure modes.

If either MCMS or TSPC systems halt, automatic alarming results in ≤ 24 minutes triggered by the lack of contemporaneous data from TSPC systems. Staff perform checks at least weekly for non-fatal TSPC warnings and errors, such as a loss of tight system time synchronization. These weekly checks also verify correct operation of two line-printers attached to the TSPC systems. The printers log certain anomalies and a daily record of clock parameters necessary to "restart" each time scale from 0 h UTC. This information is logged electronically and redundantly, so the printer output is considered a Contingency/Emergency backup.

I.A.2.8. Analysis of selected failure modes

In Section I.A.2.3, we described the Alternate and Contingency systems which can quickly substitute for the Primary time scale system at least in producing useful UTC(NIST) broadcast signals. In Table I.A.5, we examine several potential points of failure within each time scale system. We note that many of these failures are not fatal to the time scale or even observable to outside users of UTC(NIST), but instead induce either a state of holdover or marginally increased noise performance.

We note that the only common single-points-of-failure shared by the Primary and Alternate time scale systems, besides shared location, are the RF and PPS switches that select UTC(NIST) reference signals for broadcast. NIST judges that the increased component risk by introducing the switch is balanced by the

decreased time-to-repair should the operators require switching to the Alternate time scale system. Given a switch failure, spares for the switches are available, or they may be bypassed with cabling. Either repair mechanism will introduce new signaling delays that must be measured and compensated. Continuous (relative) measurements are made of the switch output signals against Primary and Alternate signals before the switch (both at 5 MHz and PPS), to facilitate realignment with a new switch or cabling.

Table I.A.5. Summary of several component failure modes in the Primary time scale system, ordered from fundamental inputs (i.e., clock signals) through outputs. In general, only failures in the output synthesis chain are observable to downstream users of UTC(NIST).

Failure mode in Primary time scale	Visible to UTC(NIST) observers?	Typical results, recovery mechanisms, notes
Outlier measurement in a single clock	No	Automatic handling by “reset” procedure in AT1; clock temporarily given no weight in TA(NIST). Two consecutive resets trigger an alarm.
Clock cable to MCMS unplugged temporarily or jostled, failure of 5 MHz distribution amplifier	No	Automatic handling by “reset” procedure in AT1. Such problems detectable by measured deviations being near exact multiples of 200 ns (i.e., periods of 5 MHz). BIPM reporting may have to be reviewed before submission (estimated time-to-repair: two days).
Frequency step in a single clock, automatic maser cavity retuning, or failure of environmental control for a single clock	No	If large (for masers, on the order of 2×10^{-14} , though this threshold varies dynamically by clock), clock will receive consecutive “reset” handling in AT1. If small, step can influence the rate of TA(NIST) in proportion to clock’s statistical weight. Operators can manually de-weight such clocks to avoid impact on time scale.
Failure of a single clock	No	Clock receives consecutive “reset” handling in AT1. Operators manually “drop” clock from scale, remove from BIPM file processing.
Failure of MCMS channel	No	If severe, this mimics outlier measurements for a single clock. Distinguishable from clock failure by examining Alternate measurement channel of same clock.
Failure of MCMS	No	TSPC software halts; time scale enters a holdover state.
Failure of TSPC	No	Time scale enters a holdover state.
Failure of Hub PC	No	BIPM requires data uploads two days after measurement, affording time-to-repair. Spare Hub is kept in “warm standby.”
Failure of monitoring or alarm systems	No	No direct coupling between monitoring/alarm systems and time scale.
Actuation of the PPS switch	No	Switching is accomplished between pulses leading to no missing pulse counts.
Actuation of the RF switch	Yes	Switching contacts are necessarily “break-before-make,” therefore several microseconds of dead output results from any use of the RF switch. Further, RF phase continuity is not preserved; downstream instruments using UTC(NIST) at 5 MHz as PLL references may come unlocked.
Failure or interruption of Primary AOG synthesizer, or clock serving as its reference, or the closed-loop measurement channels in MCMS	Yes	Can result in large offsets in UTC(NIST) or loss of signal. Time scale steering is clamped by software controls to 7.2 ns per 12 minute interval, but this mitigation is unhelpful if a reference clock or AOG loses frequency lock.
Failure of PPS generator	Yes	Can result in large offsets in UTC(NIST), usually by exact multiples of 5 MHz periods (200 ns), or duplicate/missing pulses, and loss of the phase relationship between 5 MHz and PPS signals.
Failure of UTC(NIST) RF or PPS switches	Yes	Even if healthy, Primary time scale signals may be corrupted arbitrarily by switches. Cold spares for switches or cabling must be deployed to repair or to switch to Alternate/Contingency time scales. 5 MHz switch failure distinguished by monitoring of pseudo-clock channel <i>u</i> (see Fig. I.A.14).

I.A.2.9. Spare parts, component end-of-production risk, modes of improvement

The NIST time scale has operated in a state of gradual evolution for decades. Technological improvements are incorporated as appropriate. However, a strong preference is given for long-lived technology that can be maintained or adapted by current and future staff with minimal specialized knowledge and preventative maintenance. In general, time scale components can be sorted into three categories of availability:

Commodity, commercial-off-the-shelf, largely interchangeable, no foreseen end-of-production risk – Includes atomic clocks, environmental control hardware, data acquisition computers (TSPCs), TICs, monitoring and alarming computers, digital storage, networking hardware, cabling, batteries, and UPS units. NIST staff members replace cesium beam tubes as they expire in order to return 5071A cesium clocks to service, and repair hydrogen masers when practical. NIST keeps spare components in this category unless they can be obtained commercially with little or no lead-time.

Semi-custom, high-performance, no foreseen end-of-production risk – Includes RF and PPS distribution amplifiers, low-offset synthesizers (AOG), microcontrollers and software employed in the monitoring and alarming system, and the RF and PPS switches for UTC(NIST). In these cases, commercial alternatives have been identified, but may feature different performance characteristics, interface formats, and power configurations. They are not trivially interchangeable; their substitution would require planning and testing.

Highly-custom, end-of-production risk high – Unfortunately, at present, this category includes the MCMS necessary for clock phase inter-comparison. NIST owns enough chassis and module components for four complete 32-channel time scale systems, but the risk is non-negligible that additional parts and repair service becomes unavailable in the future, with no interchangeable upgrade path.

In the next two paragraphs, we consider alternative approaches and technologies should the MCMS platform require full replacement. Pursuing either is non-trivial, as each requires significant reworking of physical and software components of the time scale system. However, outlining and analyzing these options forms a part of our long-term contingency planning.

The MCMS hardware was originally chosen because its measurement noise of $\sigma_y(\tau) \leq 1 \times 10^{-12} (\tau / 1 \text{ s})^{-1}$ was then superior to commercial alternatives such as TICs. Lacking high-precision, high-bandwidth components, the MCMS accomplished this by use of DMTD (see Section I.A.1.4), which down converts the clock frequency prior to phase detection. However, recent advances in high-bandwidth electronics and measurement techniques have led to commercial TICs capable of $\sigma_y(\tau) \leq 2 \times 10^{-12} (\tau / 1 \text{ s})^{-1}$ [12], which is nearly as good.

Similarly, digital signal processing and high-speed sampling electronics have enabled software-defined radio (SDR) techniques, which NIST has tested as an alternative for high-resolution clock phase measurement [13], capable of $\sigma_y(\tau) \leq 1 \times 10^{-13} (\tau / 1 \text{ s})^{-1}$. In addition to better performance, the SDR offers the advantage of frequency agility. Except for the SDR's analog converter's bandwidth, there are no filtering RF circuit elements. Instead, frequency down-conversion and filtering is done digitally, programmatically, in a FPGA pipeline processed synchronously with sample acquisition. For instance, a numerically-controlled oscillator (NCO) computation takes the place of a tuned local-oscillator in the SDR analogy to radio reception; digital multiplication (or the CORDIC [14] algorithm) takes the place of RF frequency mixing. Therefore, diverse clock frequencies (i.e. not just 5 MHz but, in principle, any clock frequency) could be simultaneously compared with appropriate NCO programming.

I.A.3 - Physical and Cyber Security of Primary Time Scale

In this section, we outline the security controls meant to ensure only authorized NIST staff can exert control of operating time scale assets.

I.A.3.1. Physical security

NIST time scale assets are divided between two adjacent buildings on the Boulder campus: B1 and B81. The NIST campus controls site-access for vehicles (single guarded/gated entrance and exit) and building exterior door access by electronic badges (US government common access card, CAC, with RF identification). Outside of weekday work hours, building access is subject to additional controls: part-time staff and non-US citizens require special permission and building exterior doors require staff to key in a personal identification number (PIN), in addition to scanning their electronic badge. Laboratory room access is further controlled by electronic badges; an access list only includes staff with specific need to enter the time scale rooms. B81 service corridor access is similarly restricted.

In B1, time scale assets are distributed among three adjacent rooms: maser rooms A and B, and the measurement room which also houses the cesium beam clocks. In B81, time scale assets are distributed among five rooms, all sharing a secured service corridor. The B81 measurement room houses one cesium beam device; the other rooms house between one and three masers each. Three of the B81 rooms containing masers also are home to the primary frequency standards projects (NIST cesium fountains) and optical frequency standards in development.

I.A.3.2. Cyber security

We consider the time scale cyber security in the information security framework known as the “CIA triad,” an initialism for Confidentiality, Integrity, and Availability. Of the three, Integrity and Availability have received the most attention in our design choices.

Confidentiality – The correct time is not considered to be classified information, nor are raw or processed data products resulting from our measurements of atomic clocks and time scale outputs. We publish our data and methods in publicly-available journals and conferences, and seek to support U.S. industry and commerce by freely discussing and promulgating best practices. So, in general, confidentiality is not a high-risk security component in NIST time scale operations, with the following exceptions.

- a) NIST performs calibration services for private entities whereby a customer’s atomic clock (typically a 5071A cesium-beam device) is shipped to NIST for calibration against UTC(NIST) for an interval of approximately one month. Out of an abundance of caution for the confidentiality needs of such customers, we anonymize the name of these clocks under calibration in the time scale data record. Reports transmitted to the customer containing calibration results are subject to NIST’s default customer-data security plan. Similarly, on occasion, NIST may measure signals from novel clocks or time-transfer hardware under development by private entities or publicly-funded research groups. Confidentiality expectations for such collaborations are typically made explicit by a signed Cooperative Research and Development Agreement (CRADA) or other memorandum of understanding. If ambiguous, NIST’s approach favors treating all private entity data and work products as consumer-confidential and subject to conservative data management plans.

- b) NIST staff are given many-factor login/security credentials to access various NIST laboratory virtual assets: e.g., email service, desktop workstation use, remote access to network resources via a virtual private network (VPN). Credentials include static PINs, timed one-time-use PIN-generating devices (“RSA tokens”), digital certificates stored on electronic badges, and passwords subject to length-/character-diversity/age requirements. These credentials are carefully secured by staff members, as their discovery could enable unauthorized privileged access to NIST resources. Computers employed in the NIST time scale operations (i.e. the MCMS machines, the TSPCs, the Hub PC) do not require and do not store any such NIST staff credentials, and are not usable as general-purpose computing tools. Instead, these computers are secured by physical means and lack all public network connectivity except for the Hub PC, which is subject to special NIST-managed security controls (see Fig. I.A.10). Instead, staff members perform time scale data analysis, report generation, and general-purpose computing on NIST-networked computers subject to NIST credential and security requirements.

Integrity – In our context, the integrity component of cyber security means that the atomic clock devices themselves are not exposed to disruption, and that measured data and time scale computation results are accurately acquired and preserved. Our approach begins with air-gapping all atomic clock devices: clock "health" telemetry is collected without use of general-purpose computer networks. No remote or arbitrary command/control of the atomic clocks is allowed. For example, the health status of the hydrogen masers is collected by a microcontroller with NIST-authored firmware that is triggered to execute a pre-programmed, rate-limited, poll command on the clock's serial (RS-232) communication port. Collected data is transported by an isolated serial (RS-485) bus for health checks, storage, and analysis.

The atomic clock measurement hardware (MSMC machines) and data acquisition computers (TSPC) are each deployed on non-public networks that reside behind the highest practical physical security. MCMS machines run the manufacturer's operating system (OS) and firmware from solid state storage with a file system that is mostly read-only. A small portion of the MSCS filesystem allows logfile generation and rotation. TSPC software requires no dependencies other than what is available from a default, offline, OS installation. All critical TSPC command-and-control functions require physical operator presence at the keyboard and are not remotely accessible.

TSPC file systems are subject to a weekly backup schedule. The TSPC backup media is a set of several USB solid-state storage devices regularly rotated such that many USB devices contain recent images of the TSPC systems. Following every 12-minute run of the AT1 algorithm, custom data transfer protocols shift recently-acquired data from TSPC machines to a Hub PC. The network protocols are built on top of the standard Transport Control Protocol (TCP) and wired Ethernet, which together has packet and payload data check-summing, acknowledged packet delivery, automatic retransmission capability, guaranteed correct packet ordering, or software error codes when these standards cannot be met. On the Hub PC, the TSPC data is stored in both its native, compact, binary message format, and is also converted into a human-readable line-oriented text format. Data integrity can be confirmed by comparing the native binary records stored on the TSPC machines with those on the Hub PC, and comparing text-formatted versions with original or archived binary records.

The Hub PC has one primary administrator; other time scale staff cannot generally access data stored on the Hub PC. Instead, a second layer of data transfer continuously shifts text-formatted TSPC data to two or more network file servers operated with “public read-only” permissions. One such network file server is NIST-managed and subject to frequent backups stored offsite. The Hub PC filesystem is subject to its own backup plan to attached storage.

As shown in Fig. I.A.10, the public-network interface of the Hub PC is located in NIST's so-called "network DMZ" ("demilitarized zone") and is therefore subject to special security controls managed by NIST's IT security office, OISM. For example, the Hub PC is required to run software that monitors the integrity of operating system files and binary executables, successful and unsuccessful login attempts, and maintains default-deny network firewall rules for both inbound and outbound connections.

Availability – For the NIST time scale, the availability component of cyber security translates to the goal that the Primary, Alternate, and Contingency systems operate as continuously as possible, minimizing avoidable downtime due to preventative maintenance, development, upgrades, faulty software, denial-of-service attempts, etc. Availability is increased by several of our design choices. For example, isolating critical functions from public networks removes a significant attack surface. The FreeBSD operating system has proven to be a good choice for the TSPC systems: in eight years of operation there have been no unplanned reboots or shutdowns, except for one case involving a failed hard disk.

For similar reasons, the Hub PC runs FreeBSD and is afforded additional availability provisions: mirroring RAID storage and redundant computer power supplies sourced by different circuits. An online, near-duplicate Hub PC is configured to run in "warm standby" mode as a backup system.

The failsafe-by-default design of TA(NIST) comes with marginal increased data availability risk. In the case of MCMS malfunction, the time scale control software on the associated TSPC halts and will not restart without manual intervention. While this leaves UTC(NIST) broadcast signals in a stable holdover configuration, data collection is temporarily halted, making some analyses of the interval unavailable. We implement three mitigations specific to this risk. First, any unexpected data gaps or delays at the Hub PC are detected by an alarm monitoring process; alerts are generated to "on call" staff within 24 minutes (two AT1 measurement periods), prompting rapid intervention. Second, the parallel Alternate time scale system generally continues to function given a malfunction causing the Primary system to pause, and vice-versa. Manual software adjustments can substitute clock data derived from the Alternate time scale system for gaps in the Primary system. Third, three newly developed "crate reader" (CR) computers are deployed on the private MCMS LANs with the sole purpose of automatically (re)connecting to and storing all data generated by each MCMS chassis with no design to halt after any anomalies. These data are currently archived as an independent data stream not subject to the live operational constraints of the TSPC systems.

Network links shift data from their initial acquisition point (TSPC machines) to initial processing stages (Hub PC) and then publicly-networked archival locations. All application protocols automatically recover from availability faults; for example, data gaps present on the Hub machine are automatically filled during subsequent connections with TSPC machines if possible. Data gaps on publicly-networked archives are automatically filled when availability is restored.

All MCMS, TSPC, and networking devices are powered by generator-backed UPS units. Due to their relative simplicity (uniform media, unmanaged switching, no routing, static IP addressing, no DNS services, no VLANs), the time scale LANs have demonstrated unavailability rates many times lower than the general-purpose NIST network.

I.A.4 - Validation of the Primary Time Scale as a Trusted Time Reference Source

In this section, we consider the use of NIST assets other than the Primary time scale that are useful in verifying correct operation of the primary time scale.

I.A.4.1. Coarse tests of UTC(NIST) pulse-per-second distribution

A coarse but important test of the validity of UTC(NIST) is that its physical signals correctly describe the time-of-day. Once synchronized such that a UTC(NIST) pulse-per-second (PPS) pulse N is identified with a named second in UTC (equivalent to a date and time-of-day to one second resolution), N' , counting M additional UTC(NIST) pulses should match the UTC count $N' + M$ exactly. Certain failure modes in the PPS generation and distribution hardware can result in extra or missing pulses, to which high-resolution tests (see Section I.A.4.2) are often insensitive. In particular, the MCMS will not detect errors that only affect the UTC(NIST) PPS.

NIST performs three such coarse verifications of the UTC(NIST) PPS signal. First, the PPS directly drives several so-called “impulse clocks,” which advance their time-of-day displays by one second upon reception of pulses (Fig. I.A.16). At the time of writing, NIST operates six such devices within the clock measurement room: three are adjusted to read UTC time codes, two are adjusted to read TAI time codes (equivalent to UTC plus a defined integer number of leap seconds; see Section I.A.5.2), and one is adjusted to read Mountain Standard Time (equivalent to UTC – 7:00). When additional leap seconds are announced and implemented, these clocks (excluding those set to display TAI) are manually adjusted by use of an interlock switch which commands them to “skip” a pulse. The impulse clocks are not networked and do not provide a time code reference for any user. Nonetheless, staff examine these display clocks at least weekly to detect any pathologic PPS signal issues.



Fig. I.A.16. Photograph of three impulse clocks driven by the primary UTC(NIST) PPS; each time code advances by one second whenever a PPS is received. The clocks are not networked and receive no other time- or frequency reference. Impulse clocks labeled "atomic time" are never adjusted to account for leap seconds; "UTC" time codes are manually corrected for new leap second insertions by manual use of the "hold" interlock switch shown here. Though this is a very coarse validity test of UTC(NIST), disagreements between these impulse clocks can help diagnose certain failure modes in the PPS generation/distribution network to which several high-resolution tests are insensitive.

Secondly, the PPS switch (see Fig. I.A.13, and Section I.A.2.3) maintains a count of pulses observed on the Primary and Alternate inputs. Should a relative offset develop and persist (e.g. due to either Primary or Alternate PPS generators “skipping” or “doubling” a pulse), the switch enters an alarm condition which is monitored by the B81 alarm system.

Thirdly, B81 hosts a WWVB network time appliance that receives the NIST radio signals from Fort Collins, CO (Section III.A) as its principal reference. It generates Inter-range Instrument Group (IRIG) time codes and a PPS which is compared against PPS signals from the Primary, Alternate, and Contingency (“KGA”) time scales at millisecond level accuracy.

In addition to these methods, downstream users of the UTC(NIST) PPS signal may also accumulate pulse counts and compare these to UTC-synchronized counts driven by independent atomic clocks, or GPS/GNSS time codes, WWVB receivers, or other time-of-day references. Time scale staff investigate reports of any such discrepancies as they may indicate signaling problems in the PPS distribution chain visible only to some users.

I.A.4.2. Comparisons of Primary, Alternate, and Contingency time scale outputs

Section I.A.2.3 discusses the redundant implementations of the Primary time scale, known here as the Alternate and Contingency systems. The Alternate system is composed of the same atomic clocks (and initial distribution segments) as the Primary, but with independent measurement hardware (MSMC), data acquisition systems (TSPC), and output signal generation (AOGs and counter/divider). All AOG outputs are distributed and measured as channels in the MCMS systems. This allows closed-loop control and verification of the Primary time scale in the Alternate MCMS and vice-versa. The Contingency system currently measures three clock signals in common with the Primary and Alternate systems, and signals derived from the 5 MHz UTC(NIST) output.

In general, computing the so-called “double difference” $(TA(NIST) - \text{clock } i) - (TA(NIST) - \text{clock } i)_{\text{alt}}$ signal for any clock i measured in both a Primary and Alternate (or Contingency) system produces a useful high-resolution (< 100 ps) validation of correct operation. These double-difference functions are routinely monitored for large changes, which are usually an indication of a signal distribution fault in only one of the MCMS routes.

Distribution amplifiers located before the primary PPS switch route Primary and Alternate PPS signals to TIC A in Fig. I.A.13. Periodically, the Hub PC samples TIC A to establish online agreement of the Primary and Alternate PPS; alarms are generated by the B1 alarm system if deviations exceed 3 ns.

The Contingency time scale system, known as KGA, consists of a different set of atomic clocks and is housed in an adjacent NIST building, B81. As shown in Fig. I.A.14, several diverse links transmit the Primary time scale output to the B81 measurement room; one link transmits the Alternate PPS output. All are compared against the KGA time scale output by TIC B and TIC C. Together, these measurements establish agreement between any pair of Primary, Alternate, and Contingency systems, and can identify the outlier should one of the three experience failure. The B81 alarm system makes measurements of each TIC C signal every minute and notifies staff when the relative deviations between any of the UTC(NIST) representations, the Contingency KGA output, or the Alternate time scale output, exceed 2 ns. Example data are shown in graphs (a) and (b) of Fig. I.A.17.

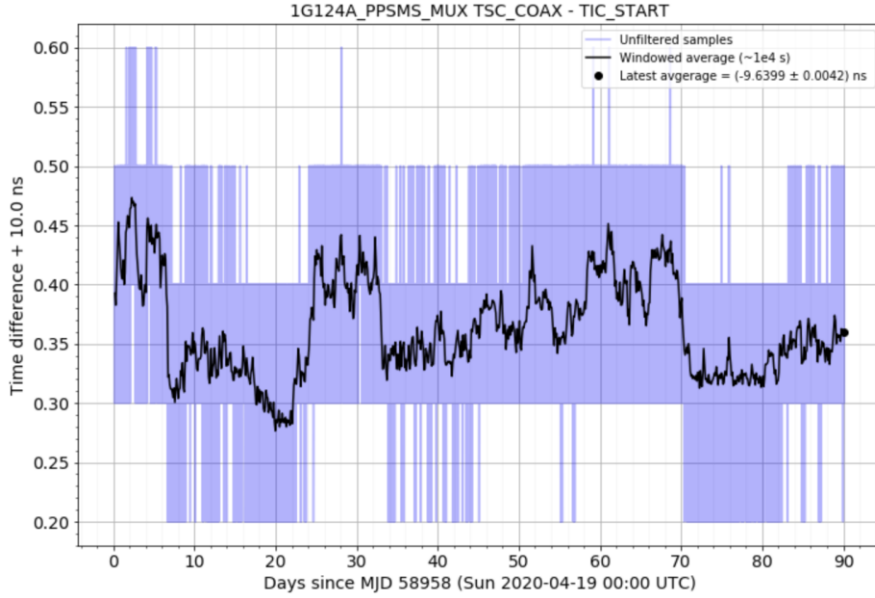


Fig. I.A.17(a).

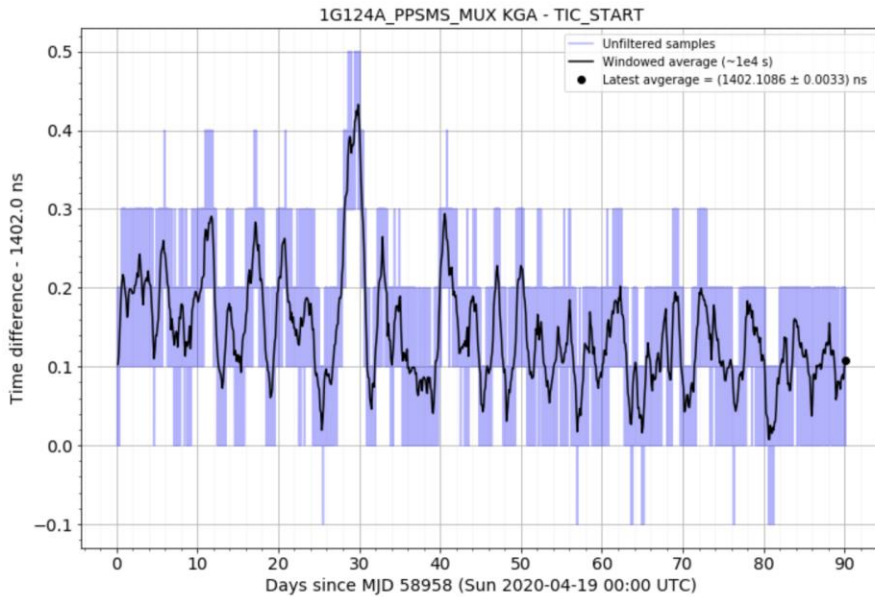


Fig. I.A.17(b).

Fig. I.A.17. Example validation measurements of the Primary time scale output against the (a) Alternate time scale PPS (labeled “TSC_COAX”) and (b) Contingency time scale PPS (labeled “KGA”) as measured by the B81 monitoring system, TIC C. Refer to Fig. I.A.14 for a signal network diagram. Both show long-term variation at the level of 100 ps or below. In this measurement system, the Primary signal “UTC_REGEN” is input to TIC C’s “start” input, while the other signals are switched through a multiplexer to appear on TIC C’s “stop” input. Samples are taken of all signals once per minute. The TIC measurement quantization limit is observed as discrete steps in the samples (blue data). For clarity, a moving average of each signal is also plotted (black line).

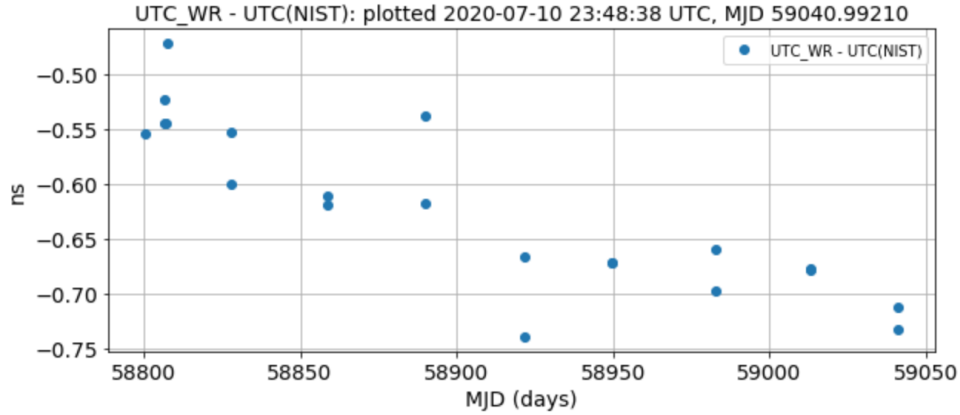


Fig. I.A.18. Example of regular calibration of delays between UTC(NIST) signals appearing in B81 to the reference point defined in B1. In this case, a “travelling” White Rabbit technique [15] is used to test the stability of the delay of a “stationary” White Rabbit link labeled “UTC_WR” in Fig. I.A.14. The stationary White Rabbit receiver utilizes two-way time transfer to obtain UTC(NIST) with only a small offset (approximately -0.50 ns near the beginning of this series, to be compared with the total signaling delay of about 1400 ns). Over the six months shown, the travelling White Rabbit technique has revealed a gradual change of approximately -0.2 ns in the stationary White Rabbit’s output. Each month, at least two “travelling” measurements are performed a few minutes apart; their scatter implies stability of the technique at the 50 ps level.

In addition to the online measurement TIC C, monthly calibration trips are performed with “travelling” White Rabbit receiver [15] using fiber optics links installed between B1 and B81. An example history of these monthly calibrations is shown in Fig. I.A.18, confirming that drifts in the signaling delay of building-to-building cable links and stationary White Rabbit links are small.

I.A.5 - Accuracy and Stability of Primary Time Scale with Respect to UTC

As discussed in Section I.A.1.1, all local time scales UTC(*k*) time scales, including UTC(NIST), that seek to approximate Coordinated Universal Time (UTC) must use UTC as an ultimate time reference. In this section, we first discuss the observed offsets between UTC and UTC(NIST) to characterize UTC(NIST)'s accuracy and stability. These offset data also inform NIST's coordination actions to steer its free-running atomic time scale TA(NIST) into UTC(NIST). Finally, for completeness, we discuss the origins of UTC and how its rate is disciplined to the definition of the SI second through a collection of the world's observation of independent clocks and frequency references.

I.A.5.1. UTC(NIST) accuracy and stability measures

The BIPM publication Circular-T [6] is the official record of UTC – UTC(NIST) time differences. Figure I.A.11 (see Section I.A.1.7) graphed four years of data from mid-year 2016 to mid-year 2020. Concerning stability, Fig. I.A.19 and Fig. I.A.20 show the time deviation (TDEV) and overlapping Allan deviation (a measure of frequency instability) of these data, respectively. Over almost all averaging intervals, the mean variation of the UTC – UTC(NIST) time offset is constrained to 2 ns or lower, and the mean magnitude of UTC – UTC(NIST) rate instability is $\leq 2 \times 10^{-15}$. Over averaging intervals longer than one year, the rate instability of UTC – UTC(NIST) $< 10^{-16}$, which is near the absolute uncertainty limits of the world's best primary frequency references.

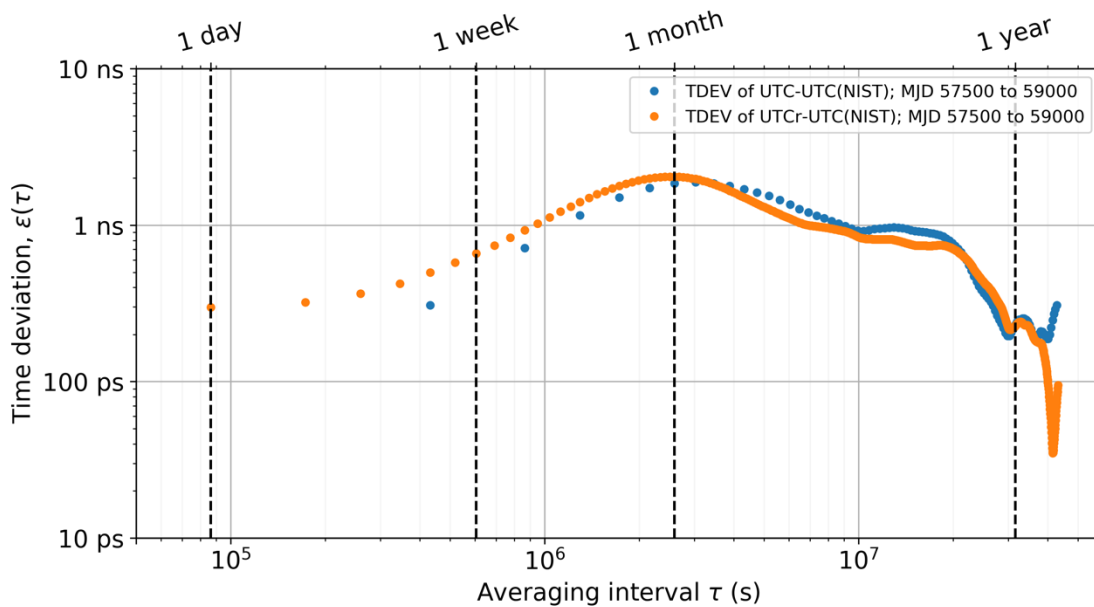


Fig. I.A.19. The time deviation [8] statistic is a measure of mean variation over a variable averaging interval τ . For intervals shorter than about one month the increasing trend is a consequence of random-walk phase noise characteristic of uncorrelated atomic clocks and time scales, as well as excess “servo variation” due to coordination actions. For intervals longer than about one month the time deviation tends to decrease as a result of successful coordination, meaning that UTC(NIST) tends to closely approximate UTC on average.

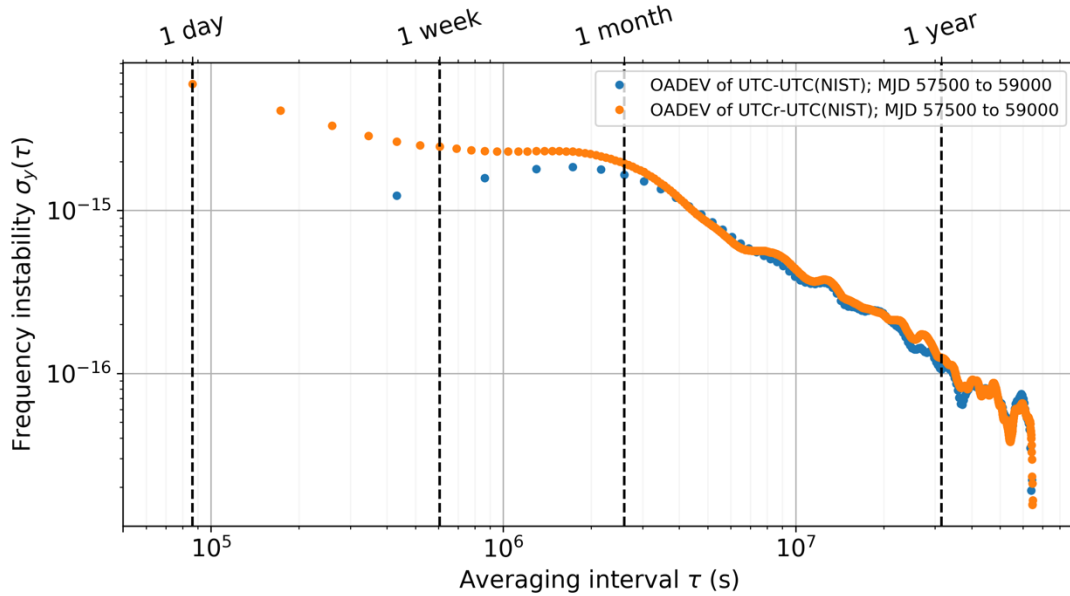


Fig. I.A.20. The overlapping Allan deviation [8] is a measure of frequency instability as a function of averaging interval τ . At intervals shorter than one month, the excess noise present in the series UTCr over UTC is observable. At intervals between one week and one month, both series show slight “servo peaks” due to coordination actions. At intervals longer than about one month, the trend is decreasing showing that UTC(NIST) realizes a stable rate with respect to UTC.

I.A.5.2. The synthesis of EAL, TAI, and UTC; fixing international time scales to the SI second

Numerous national metrology institutes (NMIs), including NIST, transfer their respective UTC(k) signals to the BIPM (see Section II.C). Briefly, this is accomplished either by physical transmission to a common “pivot” laboratory (currently assigned to be Physikalisch-Technische Bundesanstalt (PTB) in Braunschweig, Germany), accounting for all signaling delays, or by common-view GPS/GNSS observations and data transfer.

To create UTC, the BIPM does not simply “average” the many received UTC(k) signals. Instead, laboratories also transmit time difference data UTC(k) – clock i_k (where clock i_k means the i^{th} clock at laboratory k). By convention, all time differences are recorded at 0 h UTC; either daily for inclusion in the UTCr series or, at a minimum, every five days on a specified schedule (days with modified Julian dates ending in 4 or 9). As shown in Fig. I.A.21, the number of clocks reported has grown to more than 430.




Fig. A.I.21. A historical plot of the number of free-running atomic clocks reported to the BIPM for inclusion in the international time scales EAL and TAI. We show separately the populations of two common types of clocks: hydrogen masers and 5071A-type high-performance cesium beam clocks. The "other" category historically included rubidium and non-5071A cesium clock types, and more recently includes a few continuously operating primary and secondary frequency references (e.g. fountain apparatuses).

At BIPM, the received UTC(*k*) signal information and submitted clock data are used to compute clock time-differences with respect to a common, arbitrary pivot oscillator. A time scale algorithm called ALGOS [16] (similar to AT1; see Section I.A.1.6) computes a free-running atomic time scale called EAL according to the same principles underlying the AT1 algorithm: more predictable clocks are given higher statistical weight, and deterministic clock behaviors are discovered and discounted by filtering/averaging. Current administrative rules limit the contribution of any single clock to $\leq \frac{400\%}{N_{clocks}} \approx 1\%$. Notes in Table I.A.1 list the NIST clocks currently contributing to EAL. As an overview, Fig. I.A.22 illustrates the transformation of EAL into UTC and draws an analogy to the synthesis of UTC(NIST).

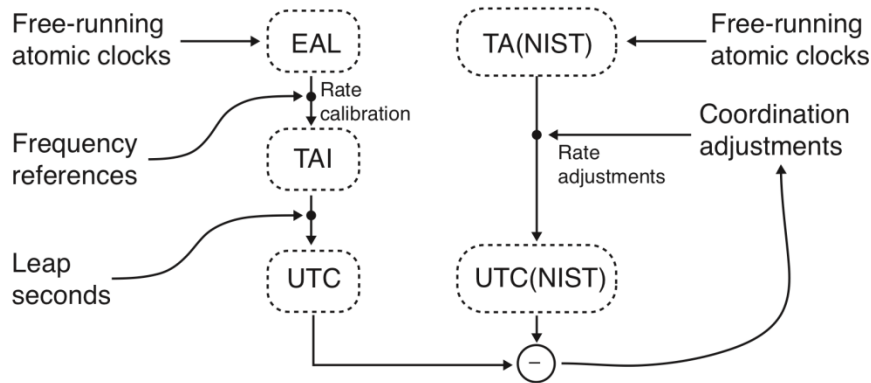


Fig. I.A.22. EAL is the BIPM’s time scale of the world’s free running atomic clocks. Observations of the rate of EAL compared to measurements of the SI second (realized by primary/secondary atomic frequency references) lead to rate calibrations which, when applied to EAL, create TAI. TAI is identical to UTC except for an integer number of seconds, known as leap seconds, meant to keep UTC roughly aligned with astronomical time scales. By analogy to EAL, TA(NIST) is synthesized from several free-running atomic clocks at NIST. Rate adjustments applied to TA(NIST) keep the observed difference $|UTC - UTC(NIST)|$ small. Leap seconds are incorporated into the time codes of UTC(NIST) but are irrelevant to physical signals (5 MHz and pulse-per-second). See text for further discussion.

Separately, several NMIs operate primary (or secondary) atomic frequency references such as cesium fountain apparatuses (for instance, see [7]) or advanced optical atomic “clocks” [17]. Ultimately, frequency references relate the rate of EAL to the interval of the SI second with an estimated uncertainty. In the present context, frequency references are technically distinguished from clocks in at least three important respects:

1. Frequency references are sometimes operated intermittently, such that they produce average rate offset data while operating, but no time offset information between operational intervals.
2. Frequency references have been studied thoroughly to characterize all known or suspected systematic biases, and to estimate the residual uncertainties due to imperfect bias compensation. Peer-reviewed, published analysis is required before a frequency reference is allowed to contribute to the rate determination of EAL.
3. While relative elevation above Earth’s rotating geoid (and local gravitational potential) is effectively arbitrary for clocks and NMIs’ realizations of $UTC(k)$, these effects must be characterized as any other systematic frequency bias due to the consequences of General Relativity. The magnitude of one such effect, the gravitational redshift, is significant: roughly 1 m of elevation change near Earth’s surface results in approximately a 10^{-16} change to the observed frequency of an atomic frequency reference. For instance, the effect in Boulder, CO amounts to about 1.8×10^{-13} .

For each Circular-T publication, the BIPM collects all valid submissions of frequency reference observations. By statistical weighting (where references with smaller systematic and statistical uncertainties are assigned increased weight), the BIPM specifies the apparent rate offset of EAL with respect to the realized SI second. By interpreting and applying this variable rate offset as a coordination action upon EAL, International Atomic Time (TAI) is created. Neither EAL and TAI are produced as physical signals. Instead, they are defined by their time differences, for example of the form $TAI - UTC(k)$ or $TAI - \text{clock } i_k$, for all participating laboratories k , and for all contributing clocks i_k .

UTC is identical to TAI except for an integer number (N_{LS}) of SI seconds, known as leap seconds [18]. The International Earth Rotation and Reference Systems Service (IERS) furnishes observations and forecasts about the behavior of $UT1 - UTC$. $UT1$ is a universal time scale referenced not to atomic clocks but to Earth’s orientation and relative position in space with respect to a model of the observable celestial background. By international agreement, and due to the current consensus desire to keep the UTC time scale roughly aligned with human/astronomical time scales, the IERS declares when leap seconds ought to be inserted into (or, theoretically, deleted from) UTC such that $|UT1 - UTC|$ is always < 0.9 s. Thus far, since this system was adopted in 1972, only leap second insertions have occurred as Earth’s average rotation rate gradually slows. Leap seconds presently occur roughly every 1.5 years, though with significant variability as Earth’s rotation rate exhibits unpredictable noise and variation: for example, due to large-scale changes in water distribution and earthquakes. By convention, leap second adjustments are scheduled as the final seconds in the months June and December. Regular announcements, even when no leap second adjustment is anticipated, are made periodically in the IERS Bulletin C [19]. At the time of writing, $N_{LS} = 37$. Consideration of leap seconds is important when calculating the time interval between any two UTC time stamps (written in terms of dates and time-of-day). TAI suffers no such ambiguities but sacrifices the traditional notion of the astronomical day in exchange.

Colloquially, the concepts of TAI and UTC may be interchanged only when the count of leap seconds is irrelevant, or when speakers implicitly share an understanding about how to discount the difference. For example, no coordination action whatsoever is required in the UTC(NIST) physical signals (5 MHz and PPS) to account for leap seconds. The physical UTC(NIST) signals proceed as an atomic time scale, much

like TAI. However, UTC(NIST) distributions, such as the NIST Internet Time Service, or the time transfer mechanisms to the BIPM, requiring a time code to correctly identify each physical signal feature do require special leap second procedures. Of relevance to this discussion, recall from Section I.A.1 that the time scale's MCMS measurements are normally sampled every $\Delta = 720$ s on a schedule aligned with 0 h UTC. In the rare cases of leap second insertions, the interval directly preceding 0 h UTC (beginning at 23:48:00) will contain 721 s instead. The additional second is given the time code 23:59:60 UTC and *is not* the same as 00:00:00 UTC the following day. Without special handling, each clock's apparent rate during these exceptional intervals would be biased by $1 / 720 \approx 0.1$ %.

I.B Realization of Secondary NIST Time Scales in Fort Collins, Colorado and Gaithersburg, Maryland

NIST's secondary time scales are located far enough from Boulder, Colorado campus that direct connection via cabling or fiber-optics is impractical. Nonetheless, NIST aims to produce UTC(NIST) representations to meet the accuracy and reliability needs of the Fort Collins WWV/WWVB radio station and the NIST Gaithersburg campus. As illustrated in Fig. I.B.1, NIST uses GPS common-view techniques to monitor the difference between the remote time scales and UTC(NIST). Should GPS become unavailable at either site, the secondary time scales will exhibit holdover performance no worse than a high-performance 5071A cesium beam clock that is optimally adjusted; Fig. I.A.2 depicts this level of dispersion as approximately 30 ns in the first month of holdover.

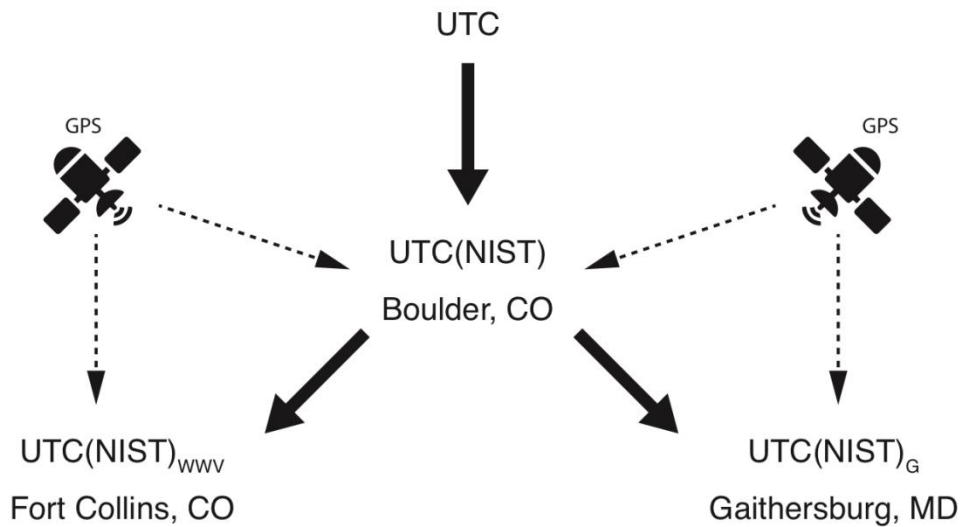


Fig. I.B.1. Illustration of the traceability of NIST's secondary time scales at Fort Collins, CO and Gaithersburg, MD to UTC(NIST) in Boulder, CO. GPS common-view observations (dashed lines) are used to monitor the time difference of each secondary time scale with respect to UTC(NIST); coordination actions (approximately biweekly or monthly) steer the secondary time scale frequencies to remove accumulated time offsets. UTC(NIST) itself is coordinated to approximate UTC, as discussed in Section I.A.5.

If all alternate, contingency, and emergency time scale systems (see Section I.A.2.3) in Boulder were to fail, the time scales at Fort Collins, Colorado and Gaithersburg, Maryland, would enter a holdover state exhibiting the time deviation of one, or a statistical ensemble of a few, cesium atomic clocks. Government and private-sector U.S. entities with a critical interest in accurate time may participate in common-view observation or two-way time transfer (including fiber-optics) with one, two, or all three NIST sites to take advantage of assurances the redundancy provides. In a hypothetical scenario of complete failure at the Boulder, Colorado site, accurate dissemination of time at or below 1 μ s could proceed indefinitely from these secondary sites assuming that NIST can (re)establish common-view or two-way time transfer with other metrological laboratories in cooperation with BIPM.

It is important to realize that the geographic diversity of the three sites eliminates the possibility of many simultaneous catastrophic failure modes. It is also important to realize that the time scales at all three sites share no failure modes with any GPS or GNSS systems. We emphasize that GPS is currently used to coordinate the secondary time scales, but other time-transfer methods could be substituted at higher cost.

I.B.1 - Technical Description of Secondary Time Scale Architectures

In this section, we describe the implementation of each remote, secondary NIST time scale. As many principles are shared with the primary time scale design, we refer to sections above as appropriate rather than repeat technical detail.

I.B.1.1. Time scale at Fort Collins, Colorado

NIST has operated a secondary time scale at Fort Collins, Colorado since 2007 [20]. The Fort Collins site is the home of radio stations WWV and WWVB (Section III.A), which distribute UTC(NIST) signals to many millions of users, and the site also houses a node of NIST’s popular Internet Time Service. Therefore, we implement online redundancy to help ensure continuous availability of the local time reference. Typically, the local time reference is coordinated to be within ± 25 ns of UTC(NIST). This accuracy is more than sufficient since the station’s transmission modes (WWV, WWVB, and NTP over the public Internet) typically permit no better than 100 μ s end-to-end accuracy.

As is the case with the Primary time scale in Boulder, two redundant time scale multi-channel measurement systems (MCMS) observe the phase evolution of several atomic clocks. Table I.B.1 lists the current roster of 5071A high performance cesium beam clocks. Output synthesizers (high resolution offset generators, or HROGs) are programmed to steer the output of an atomic clock to realize the UTC(NIST) secondary time scale.

The Fort Collins MCMS hardware is an older (Timing Solutions), but similar, version of the MCMS deployed in Boulder. Each Fort Collins MCMS can measure eight clock signals. HROGs are similar in function to the AOGs deployed in Boulder but with reduced resolution in offset programming.

Table I.B.1. Atomic clocks (HP: 5071A cesium beam device) and time scale output synthesizers (HROGs) participating in NIST’s Fort Collins primary/alternate secondary time scale. NIST clock identifiers are derived from serial numbers. Codes in the second and third columns indicate the clock’s roles in each time scale’s measurement and control system. M: measured in the MCMS. W: receives statistical weight in the AT1 calculation. S: serves as the time scale’s output synthesizer. P: serves as the synthesizer’s reference signal. 1: serves as the MCMS NCO-tuning reference frequency (“channel 1”). Information is accurate at time of writing but subject to change.

NIST identifier	Roles in primary Fort Collins time scale	Roles in alternate Fort Collins time scale
HP0282	M	M
HP1519	M, W	M, W
HP2031	M, W, P	M, W
HP2032	M, W	M, W, P
HP2034	M, W	M, W
HROG-A	M	M, S, 1
HROG-B	M, S, 1	M

As in Boulder, the two Fort Collins time scale systems are operated by separate, dedicated computers (TSPCs). The TSPCs calculate and implement HROG steering corrections every 12 minutes. Data are aggregated and transmitted to the Boulder Hub PC via a virtual private network (VPN) connection, where they are processed and analyzed along with the Boulder time scales' data.

I.B.1.2. Time scale at Gaithersburg, Maryland

The NIST campus in Gaithersburg, Maryland houses a node of NIST's Internet Time Service and the fiber-based remote time transfer calibration service (Section II.A). Currently, two high-performance 5701A cesium beam clocks (HP3366 and HP3367) serve as primary and alternate time references. No MCMS, AOGs, or HROGs are present. Instead, one TIC monitors the time-difference between the two clocks PPS output, and a second TIC measures the difference between one clock and a locked GPS common-view receiver. Data are relayed to the Boulder Hub PC via VPN. Occasional (approximately monthly) frequency corrections are programmed directly to the primary clock (as in Fig. I.A.3a) with the goal of coordinating its time output to within ± 20 ns of UTC(NIST). Similar, distinct frequency steers are applied to the alternate clock.

I.B.2 Methods Ensuring Continuous Operation of Secondary Time Scale Architectures

The subsections here mimic the structure of Section I.A.2. Because several design principles are shared with the primary time scale system, we refer to that material as appropriate. In general, the relative simplicity and looser performance goals of the secondary time scales also afford less complicated redundancy provisions.

I.B.2.1. Alarm system and on-call staff

The Boulder Hub PC receives regular data updates (via VPN) from the secondary time scales; these updates are analyzed to produce alarms in parallel to the B1 alarm system described in Section I.A.2.1. Table I.B.2 provides a list of the conditions that are monitored and recorded. Alarms related to the secondary time scales are typically routed to a subset of specialized staff. Data is analyzed weekly by specialized staff to predict needed coordination actions.

Table I.B.2. List of conditions which are monitored and recorded (code: M) and alarmed to specialized staff (code: S) by the B1 alarm system. See related alarm functionality in Table I.A.2.

Conditions	B1 monitoring/alarm
Fort Collins time scale or GPS common-view data missing	M, S
Fort Collins MCMS or TSPC halt; holdover condition	M, S
Fort Collins clock has successive “resets” in AT1 algorithm	M, S
Fort Collins: large AOG steering commands by AT1 algorithm	M, S
Fort Collins: large PPS difference between the two scales	M
Fort Collins: utility- and UPS-power voltage monitoring	M
Gaithersburg: clock or GPS common-view data missing	M, S
Gaithersburg: clock PPS large offset with respect to GPS	M, S

I.B.2.2. Failsafe by default

Without any intervention or coordination actions, the secondary time scales will exhibit the holdover characteristics of a 5071A high-performance cesium beam clock. A relatively long mean-time-to-repair is allowable considering the lower accuracy requirements of the secondary time scale users and is necessary due to the lack of on-call staff availability at the remote sites.

I.B.2.3. Redundancy and reliability

Fort Collins. Primary and alternate clock measurement and time scale control systems are continuously operated in parallel. As a contingency procedure, each TSPC can be configured to program either time scale output synthesizers. This provision acts like a “software switch” in lieu of a hardware signal switch for selecting between the primary and alternate implementations. Should both time scale systems fail, any of the five 5071A cesium clocks could be reconfigured to be directly steered to follow UTC(NIST).

Gaithersburg. Two 5071A cesium clocks are coordinated to approximate UTC(NIST). One is designated as the site’s primary reference; the other is an alternate. No physical signal switch is currently implemented, but site users can obtain access to both clocks and implement a “software switch” if appropriate.

I.B.2.4. Clock health telemetry

The atomic clocks comprising both secondary time scales are 5071A high-performance cesium beam clocks. The same data discussed in Section I.A.2.4 are collected and monitored by specialized staff, mainly with the aim of anticipating an approaching end of the clocks' service life.

I.B.2.5. Environmental control

Because no hydrogen masers are present at either secondary time scale sites, and because the stability requirements of the secondary time scales are relaxed compared to UTC(NIST), typical laboratory or office environmental control is sufficient. A severe loss of environmental control resulting in overheating can be addressed by non-specialized on-site staff.

I.B.2.6. Redundancy in power supply

Fort Collins. The secondary time scale assets are provisioned with utility power backed up by a diesel generator and automatic transfer switch mechanism (details about the generator are provided when discussing WWVB in Section III.A.2). This power is buffered by several dedicated UPSs such that the two time scale implementations depend on different UPS systems. Many components (clocks, MCMS, distribution amplifiers) accept DC power from a 24 V series of high-capacity marine batteries, which is kept on a trickle charger.

Gaithersburg. The secondary time scale assets exist in a central computer facility which provides generator- and facility-grade UPS redundant power sources. Due to facility policy, no dedicated UPS or battery backup sources are allowed within the equipment racks.

I.B.2.7. Preventative maintenance

Fort Collins. Typically, specialized NIST staff visit the site weekly to load test the DC battery bank and react to any negative clock health telemetry items. When necessary, coordination actions, and some contingency procedures, can be programmed remotely.

Gaithersburg. No preventative maintenance by specialized NIST staff is required. Extraordinary conditions can be diagnosed remotely with the assistance of non-specialized on-site staff.

I.B.3 - Physical and Cyber Security of Secondary Time Scales

A security plan similar to that of the primary time scale system is applied to the secondary time scales. Generally, the secondary time scales are simpler to secure as all assets exist in single rooms, with simpler network connectivity, more restrictive physical access procedures, and fewer staff interacting with the equipment.

I.B.3.1. Physical security

Both the Fort Collins and Gaithersburg sites are fenced; primary access-control is via electronic badge readers. In Fort Collins, the atomic clocks and time scale measurement hardware are located in a single-purpose locked room. In Gaithersburg, the atomic clocks and time scale measurement hardware are collocated in a central computing facility that is subject to strict room- and rack-access controls.

I.B.3.2. Cyber security

Like the Boulder time scale, integrity and availability are the most important aspects of the information security plan. The same concerns and methods apply as described in Section I.A.3.2. The links between the secondary time scale data sources and the NIST time scale Hub PC are dedicated VPN configurations.

I.B.4 - Validation of Secondary Time Scales as Trusted Time Reference Sources

As in Section I.A.4, we wish to be able to verify correct performance of the secondary time scales at any time without making direct use of the secondary time scale's assets, including the GPS common-view data linking them to UTC(NIST).

Both secondary time scales provide time references to nodes in NIST's Internet Time Service. Though using NTP as a validation check offers relatively low (millisecond level) precision, it readily confirms correct date/time-code generation and signal persistence.

Fort Collins. The alternate time scale system at Fort Collins validates the performance of the primary, to within a common absolute offset uncertainty of the GPS common-view technique. NIST maintains a redundant common-view calibration link at Fort Collins, which confirms the low offset of the secondary time scale with respect to UTC(NIST). This link is independent from the main common-view GPS link that is used for the time scale's coordination. Likewise, local staff at the Fort Collins site operate other GPS receivers and can verify that the time error of the secondary time scale is always $< 1 \mu\text{s}$, even without the common-view links. If GPS becomes unavailable for an extended interval, a transportable 5071A high-performance cesium beam clock is available at NIST's Boulder Colorado laboratory, and a day trip between the sites can likely accomplish time transfer of UTC(NIST) at the level of $\pm 10 \text{ ns}$.

Gaithersburg. Primarily, validation is accomplished by the inter-comparison of the two 5071A cesium beam clocks (independent oscillators between coordination actions) and GPS/GNSS signals. Validation of the GPS common-view link is performed by first stabilizing a transportable rubidium clock (or similar) to received GPS signals obtained in clear sky outside the facility. The rubidium clock can then be brought into the facility to confirm various time offsets at the measurement reference points with an uncertainty at the level of $\pm 10 \text{ ns}$. Given the inconvenience, this procedure is only attempted as necessary. Should GPS become unavailable for an extended interval, special arrangements must be made for a higher quality transportable clock referenced to UTC(NIST).

I.B.5 - Accuracy and Stability of Secondary Time Scales with Respect to UTC

Both secondary time scales are composed of fewer and less-stable clocks than the primary realization of UTC(NIST). Both secondary time scales are coordinated to approximate UTC(NIST) but accumulate time errors in between coordination actions. Therefore, neither secondary realization of UTC(NIST) can exceed the accuracy or stability of the primary realization of UTC(NIST) which itself is coordinated to approximate UTC. The coordination “stiffness” is a balancing act: coordination steers are frequent enough to keep the secondary time scale maximum time errors (typically ± 25 ns) far below the accuracy requirements of the sites’ users (typically ± 1 μ s), but also as infrequent as is practical to maximize the short-term independence of the secondary time scales as reference sources. Likewise, to maximize the secondary time scales’ frequency stability, time offsets are never abruptly corrected. Instead, NIST attempts to make smooth frequency offset adjustments calculated to remove time offset errors gradually.

Regular GPS common-view observations are recorded and transmitted to the Boulder Hub PC as part of the coordination plan for the secondary time scales. The Hub PC regularly computes the common-view GPS offset between UTC(NIST) and the Fort Collins time scales. Frequency offset corrections to the Fort Collins time scales are applied manually every few weeks, consistent with the goal of keeping the time offsets with respect to the primary realization of UTC(NIST) to within about ± 25 ns.

Figure I.B.2 shows the time difference of the Fort Collins time scale with respect to UTC(NIST) in Boulder as measured via common-view GPS for the six-month period from September 2019 to February 2020. The graph shows one-hour averages recorded continuously during the 182 day period. The average time difference is 12 ns, and the range of values is 34 ns. The maximum time difference was -29.5 ns (note that the uncertainty of the common-view measurement is ~ 10 ns).

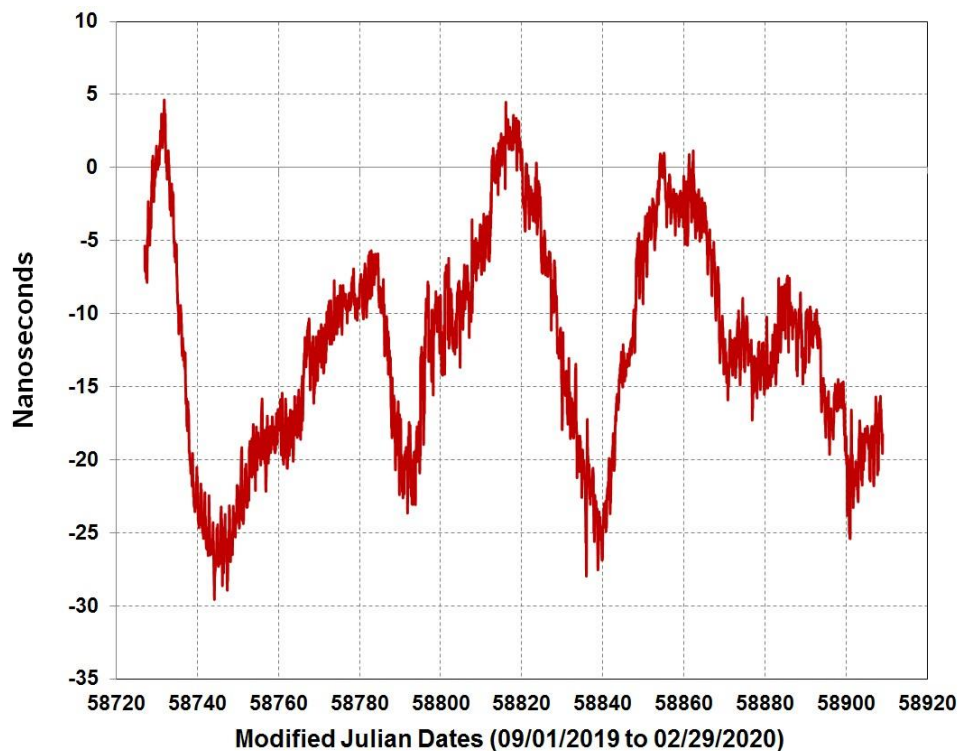


Fig. I.B.2. Six-month comparison of secondary time scale in Fort Collins to primary time scale in Boulder. For nearly all critical infrastructure applications, the time delivered by this simple, robust secondary time scale will be indistinguishable from the time delivered by the primary time scale.

II. Methods of Distributing UTC(NIST) at Sub-Microsecond Accuracies

Several existing methods of distributing UTC(NIST) can currently achieve sub-microsecond accuracy. These methods have been implemented either as part of NIST's internal operations, or as a calibration service or special test offered by NIST on a subscription basis to the general public. Each method, however, would benefit from future investment and refinement. The goals of this investment would be to reach more users, to improve reliability, and to improve ease-of-use. In some instances, the transfer of these distribution technologies to U. S. industry would likely be necessary to accomplish all of these goals.

The importance of making multiple methods of time distribution available cannot be overemphasized, as diversity is a fundamental component of any resilient architecture. Thus, three distinct UTC(NIST) distribution methods are described in detail in the following sections. Section II.A describes time distribution via optical fiber, Section II.B describes time distribution via two-way satellite time and frequency transfer (TWSTFT), and Section II.C describes time distribution via multi-source common-view disciplined clocks (MSCVDC).

II.A Distribution of UTC(NIST) via Optical Fiber

In 2019, NIST introduced a “Special Calibration Test” [21] that is designed to provide signals traceable to UTC(NIST) from the secondary time scale (Section I.B) at Gaithersburg, Maryland. The signals travel through third-party optical fibers to a customer’s system at a remote location. The description of the test was that it would provide a time signal to end users with an initial accuracy of $\pm 1 \mu\text{s}$; the accuracy was expected to improve to $\pm 100 \text{ ns}$ after six months of testing and evaluation. The preliminary results that are presented in this section suggest that the ultimate goal of $\pm 100 \text{ ns}$ has already been realized. Although the initial announcement of the special test described the test based on signals derived from the secondary time scale in Gaithersburg, the same service can also be based on signals derived from the primary time scale located in Boulder, Colorado [22]. The following discussion describes the offering from either location.

Our descriptions of the various ways that NIST can distribute time via optical fiber do not define or constrain the details of the transport channel or the end-point interface hardware. This omission is intentional as it allows us to support solutions that propose hardware or firmware that is unique or proprietary, and also allows us to evolve and improve the service as new techniques and hardware become available. Thus, the description is focused on the goal of providing accurate time to remote systems, while at the same time enabling maximum flexibility in the methods that can be used to realize the goal.

Section II.A.1 provides a technical description of the architecture for a time over fiber distribution system. Section II.A.2 covers the reliability of the distribution system, its levels of redundancy, and the safeguards in place to ensure continuous operation. Section II.A.3 describes the physical and cyber security of the system. Section II.A.4 describes the frequency and time stability of the signals distributed by fiber with respect to UTC(NIST). Finally, Section III.A.5 discusses the feasibility of distributing UTC(NIST) via fiber to a large number of critical infrastructure systems.

II.A.1 – Technical Description of Architecture for Time Over Fiber Distribution System

As of September 2020, NIST has its first test project from a commercial firm named OPNT (Optical, Positioning, Network and Timing, www.opnt.nl), and we have several other test projects that are in early discussions. The initial OPNT project is to implement a Special Calibration Test [21] between the NIST campus in Gaithersburg, Maryland, and a colocation network facility in McLean, Virginia, approximately 30 km directly south of Gaithersburg. The hardware at the colocation facility would redistribute the timing signals to end users. The physical length of the fiber circuit between the end points is not known now, but is likely to be greater than the straight-line distance. We have installed, configured, and tested the end-point hardware at the Gaithersburg end of the link and the fiber circuit between Gaithersburg and McLean. We have preliminary data on the statistical characteristics of the link between Gaithersburg and McLean. We also have preliminary measurements of the characteristics of a longer link between McLean and Atlanta, Georgia, a one-way distance of about 700 km.

The initial design for the OPNT test was based on a pair of fibers linking the two end-points, where each fiber transmitted signals in only one direction. The current design, which has now been commissioned, is based on using a single fiber strand to transmit data in both directions on different wavelengths. In addition to the hardware that supports the conversion between the optical and electrical components of the link, the single-fiber realization also requires optical filters at both ends of the path to separate the signals at the different wavelengths that propagate on the single fiber. As we will discuss below in more detail, the current implementation uses two pairs of wavelengths to transmit two timing signals in both directions. These two signals are derived from two cesium clocks at Gaithersburg.

The remainder of this section provides a technical description of the system architecture, beginning with a discussion of the types of optical fibers, followed by a discussion of transport methods and message formats, and information specific to the Special Calibration Test with OPNT.

II.A.1.1 - Types of optical fibers

Although both single-mode and multi-mode optical fibers are commercially available, only single mode fibers are used for long-distance communication because they exhibit much lower signal attenuation. Multi-mode fibers typically attenuate 850 nm signals by about 3 dB/km whereas single-mode fibers have much smaller attenuation factors; typically about 0.25 dB/km for signals at 1550 nm, and 0.35 dB/km for signals at 1310 nm [23, 24]. Therefore, the long-distance circuits that realize the Special Calibration Test will always be single-mode fibers. If the optical fiber length is about 100 km, the fiber path will attenuate the signal from 25 dB to 35 dB. The attenuation of 25 dB reduces the signal by a factor of about 316, and 35 dB attenuation reduces the signal by a factor of about 3160. Note that these values are examples; the fiber linking Gaithersburg and McLean is almost certainly shorter than 100 km.

II.A.1.2 - Transport method and message format

All of the currently known solutions depend on the two-way protocol, in which the one-way channel delay is estimated as one-half of the measured round-trip value. The details of how this estimate is computed varies from one method to another, but all methods assume that the inbound/outbound path delays are equal, so that the time offset that would arise from a path asymmetry is not a problem. When the signals sent in the inbound and outbound directions are transmitted over an optical fiber by means of different wavelengths, there will be a systematic asymmetry in the delay because of the dispersion of the optical fiber (the magnitude of this effect can be estimated from the physical characteristics of the fiber or measured as will be described). All of the protocols use the exchange of multiple messages to estimate the channel delay, and implicit in this exchange is that the underlying delay does not change during the time of the message

exchanges. Finally, as will be discussed in Section II.A.4, the optimum strategy is to use the message exchange to synchronize a remote clock rather than to use messages to directly apply time-stamps to events.

The initial implementation of the link between Gaitherburg and McLean is based on the White Rabbit protocol (described later in this section), an extension to the IEEE-1588/Precision Time Protocol (PTP) message exchange [25, 26]. The end-point hardware is interfaced to the reference time system by using signals at 1 PPS and 10 MHz. This method has two problems. First, the signals contain no time-of-day information, so that they must be augmented by a system that assigns time-stamps to the 1 PPS pulses. This assignment is needed during a cold start and to indicate the occurrence of a leap second. Second, any method based solely on signals that convey no time-of-day information is sensitive to the problem of cycle slips – or errors that result in a time step of the interval that corresponds to exactly one pulse. This type of error cannot be detected without ancillary information.

To support these two requirements, the White Rabbit message exchange that supports the link between Gaithersburg and McLean uses Network Time Protocol (NTP) messages [27, 28] to acquire time-of-day information. The NTP messages are exchanged with the other systems at Gaithersburg that support the NIST Internet Time Service. It would be useful to use the fiber circuit to transmit time-of-day information in NTP format to provide an independent method of detecting transmission errors, but this is not planned for the initial configuration.

The Precision Time Protocol (PTP) is typically realized in hardware, which is its primary advantage relative to NTP. The channel that links the local and remote systems operates in half-duplex mode, meaning that the messages exchanged between the two end points are transmitted consecutively rather than simultaneously to estimate the one-way delay, as is also true with NTP, although the details are somewhat different.

A potential advantage of PTP over NTP is in the methods it includes for measuring and correcting for the delays in the elements that are present in a public network. The delay through each network element is measured by a special “boundary clock,” which must be a part of every network element in the path between the two systems. However, most of the routers on the public Internet do not include this hardware to support this, so PTP is not suitable for a wide-area network such as the Internet except in special situations.

Generally, PTP is realized as a hierarchical configuration of reference and secondary clock systems, and it includes a method for establishing the pyramid structure by automatically configuring a group of clock systems on a network. However, these capabilities are not important in the context described here, where messages are exchanged between a single reference clock and a single secondary clock over a dedicated fiber with no other traffic. In this case, PTP can easily support $\pm 1 \mu\text{s}$ timing accuracy. Unlike NTP, which includes extensive tests to detect a remote system that is transmitting the wrong time (a “false-ticker” in NTP jargon), PTP trusts the correctness of the remote time, and must therefore be augmented by control and supervision methods that are outside of the protocol. These methods are already part of the reference time systems at all of the NIST sites.

The White Rabbit protocol [29, 30] is an extension of PTP and it therefore inherits the general features described in the previous paragraph. The White Rabbit protocol is realized as PTP over synchronous Ethernet. The physical-layer clocks in the network nodes of a White Rabbit network are phase-locked together. This increases the resolution of the message exchange by eliminating the phase jitter that is present in standard Ethernet, where the clocks in the various network elements are not synchronized. A White Rabbit system can realize sub-nanosecond time accuracy, but the price is that all of the network elements must be designed to preserve the synchronous Ethernet clock. As with the PTP protocol discussed in the previous paragraph, this is not an important limitation in the current discussion, where time signals will be transported over a dedicated fiber link with no other traffic. The accuracy of White Rabbit depends on the

fact that the physical-layer clocks are phase-locked to each other, and cycle slips of a physical clock are a potential problem that must be addressed and monitored.

Both PTP by itself and PTP with the White Rabbit extension can support a $\pm 1 \mu\text{s}$ timing requirement when the link between the UTC(NIST) time scale and the remote system is realized over a dedicated fiber circuit. The White Rabbit protocol can support greater resolution, and can support sub-nanosecond stability in principle, making sub-nanosecond accuracy possible. However, it is not clear that this level of accuracy can be realized in the overall system. For example, there is currently no way to validate the time offset between the secondary time scale in Gaithersburg and the primary time scale in Boulder at the sub-nanosecond level. However, the uncertainty of the secondary time scale in Gaithersburg with respect to the primary time scale is significantly less than 100 ns, and will not limit the ultimate accuracy specification of the Special Calibration Test.

II.A.1.3 - Technical description of the OPNT architecture

The root of the proposed distribution system is a pair of devices from OPNT that can each function as either a grand master or boundary clock in White Rabbit format. They can also serve as a PTP reference clock for a time distribution system because they can receive timing information from a primary reference system. The time transmitted by these devices is normally accepted unconditionally by down-stream devices, so that the system will require ancillary monitoring and supervision to detect problems.

The White Rabbit hardware that is currently installed at NIST in Gaithersburg can support two different configurations. The initial, currently operational, configuration is shown in Fig. II.A.1. The two cesium clocks at Gaithersburg are configured as the reference clocks for two White Rabbit grand masters. Both of the grand master devices receive time-of-day information in NTP format from time servers that are collocated with this hardware. The time-of-day information only needs to be accurate enough to uniquely identify the time stamp associated with a 1 PPS pulse and to detect possible cycle slips. This requirement is easily satisfied – the accuracy of the NTP messages exchanged over the private network is measured and monitored routinely and is typically better than $\pm 100 \mu\text{s}$ peak-to-peak.

The optical outputs of the two grand masters are configured to transmit and receive data on two different pairs of optical wavelengths (typically 1550 nm/1490 nm for one pair and 1530 nm/1470 nm for the other). The passive optical filters multiplex the signals at the two output wavelengths onto a single optical fiber, and route the wavelength signals received from the remote end back to the appropriate grand master device. The hardware at McLean has a similar optical configuration and recovers the two signals and measures the time difference between them. The time difference between the two data streams measured at McLean is compared with the time difference between the two cesium clocks as measured at Gaithersburg, as shown in Fig. II.A.1.

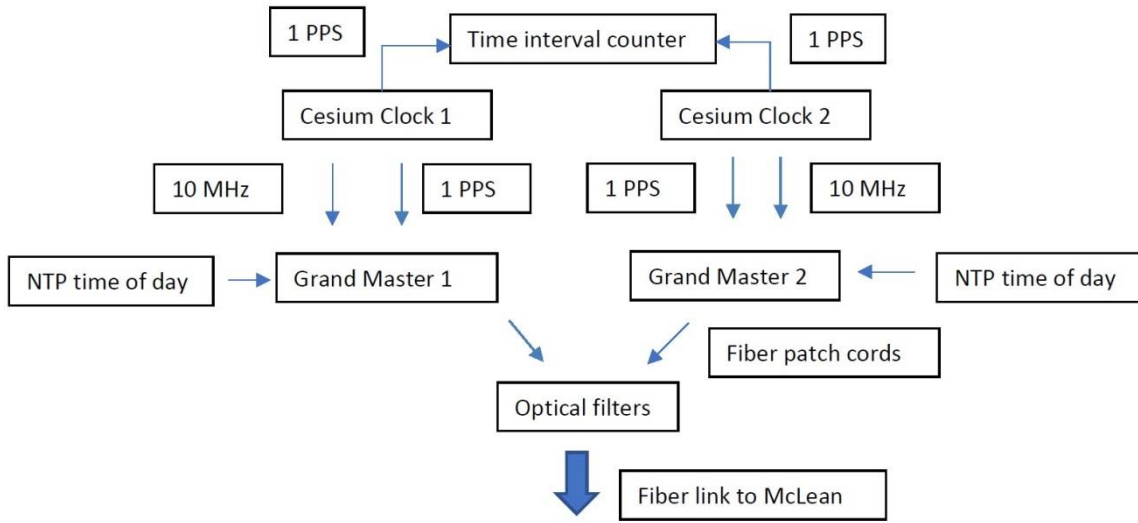


Fig. II.A.1. The initial configuration of the White Rabbit hardware at NIST in Gaithersburg as described in the text. This configuration compares the time difference measured locally by a time interval counter with the time difference recovered by the receiving hardware at McLean. The results are shown in Fig. II.A.2.

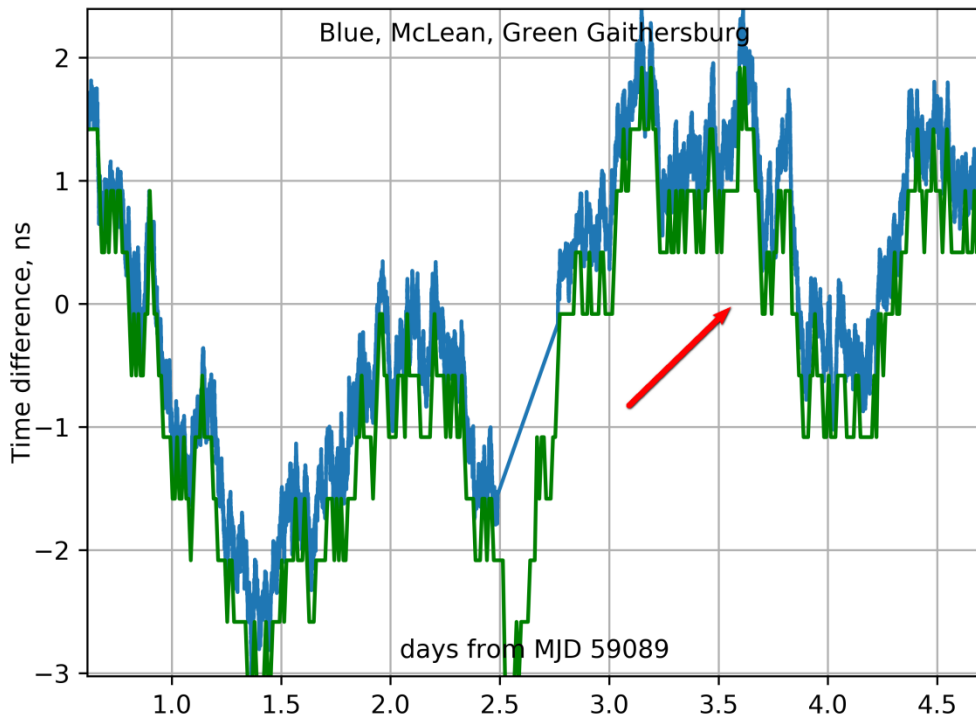


Fig. II.A.2. A comparison between the time differences measured at Gaithersburg by the time interval counter (Fig. II.A.1) and the time differences recovered by the hardware at McLean. The mean values of both series have been subtracted as described in the text. An intentional steering adjustment was applied to Clock 1 to validate the measurements. The time of this adjustment is shown by the red arrow.

Figure II.A.2 compares the local and remote time difference data. The two data sets have a relative time offset that is due to the differential delays in the distribution amplifiers and cables at the Gaithersburg site. The mean value of each series has been subtracted; the two mean values differ by this time offset, which is ~17 ns. Although it is not shown in the figure, the times of both cesium clocks are calibrated and monitored with respect to the primary UTC(NIST) time scale in Boulder, and both clocks have an offset with respect to UTC(NIST). The goal of these initial tests is to validate the performance of the link between Gaithersburg and McLean, and the offset with respect to Boulder is not important for this evaluation.

The comparison data shown in Fig. II.A.2 are very encouraging. The time-difference measurements at Gaithersburg are limited at short times by the 500 ps (0.5 ns) resolution of the TIC, but there is no indication of any long-term variation in the initial time offset over the five-day measurement period. We plan to upgrade the distribution hardware at the Gaithersburg site to address delay differences between channels of the 1 PPS distribution hardware.

The design in Fig. II.A.1 has the useful advantage of minimizing the single points of failure by exploiting the two independent cesium clocks and the two independent grand masters. Although the results obtained with this configuration validate many of the design parameters of the White Rabbit configuration, it is not sensitive to an overall calibration offset, and it cannot confirm the accuracy of the time signals at McLean with respect to UTC(NIST) without an ancillary calibration of the recovered time at that end.

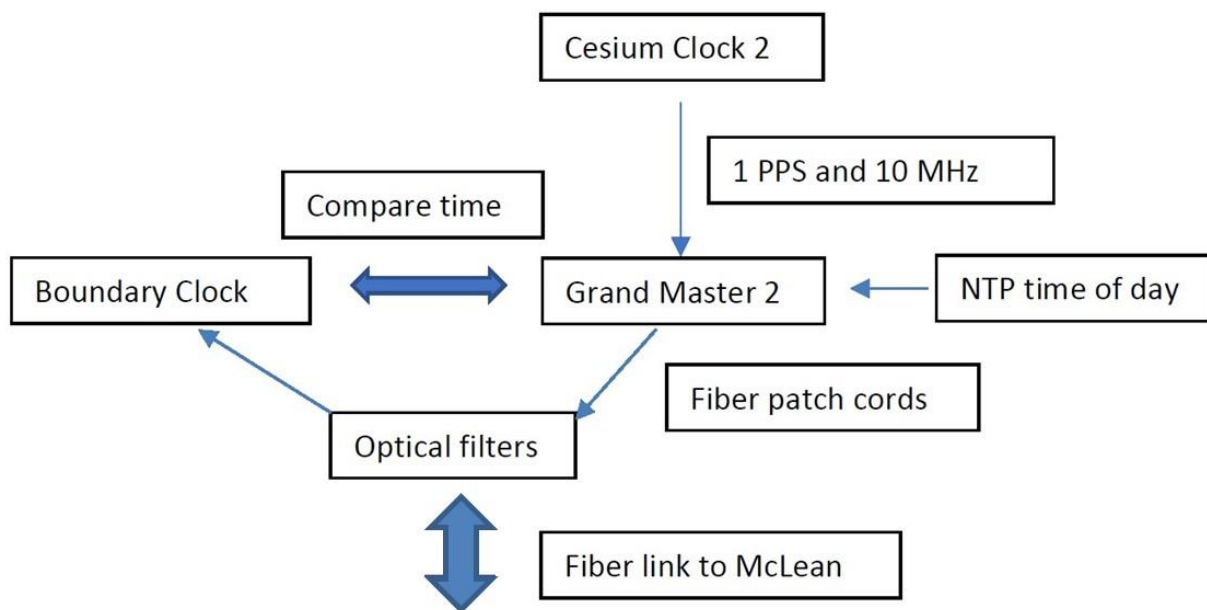


Fig. II.A.3. An alternate configuration at Gaithersburg in which the fiber link is configured in a loop-back arrangement. The grand master transmits time information to McLean; the hardware at McLean echoes the data back to Gaithersburg, where it is used to synchronize the boundary clock. The time recovered by the boundary clock is compared with the time of the cesium clock that acts as the reference for the grand master. The configuration validates the calibration of the network delay but has a single point of failure.

Figure II.A.3 shows a different configuration that has not yet been tested that uses the same basic hardware. It trades the redundancy of the configuration shown in Fig. II.A.1 for the ability to monitor the overall accuracy of the link between Gaithersburg and McLean. The time signals from cesium clock 2 provide the reference for grand master 2. The time data are transmitted to McLean in White Rabbit format and are echoed back to Gaithersburg by the remote hardware. The boundary clock receives the data from the remote

end and recovers the timing information. These data are compared with the time of the reference clock. This configuration would validate the estimate of the network delay and would be useful for monitoring the performance of the entire configuration. It has the disadvantage relative to the first configuration that it has more single points of failure and does not use the data from the second cesium clock at Gaithersburg. The time difference between cesium clocks 1 and 2, and the ancillary operating parameters of both clocks, are monitored to detect a failure of either clock. These data are acquired in both configurations.

The initial design of the fiber link to the remote system at McLean, Virginia, used a pair of optical fibers, with each fiber transmitting signals in only one direction. The current design, which is illustrated in all of the previous figures, uses a single fiber transmitting signals in both directions by means of two different wavelengths. For example, the signals in one direction could use a wavelength of 1550 nm, while the signals in the other direction would use 1490 nm. These wavelength pairs (and others) are industry-standard values, so that commercial equipment to handle these signals is readily available.

The choice between two one-way fibers and one two-way fiber involves a complex analysis of various factors, and there is no obvious best choice. Some important points that must be considered when choosing between these two configurations are:

1. After the system has been installed and configured, a single-fiber configuration may have lower recurring costs because the principal recurring cost is the cost of the circuit. However, the one-time non-recurring costs are greater because of the filters required at the end points to separate the signals at the two wavelengths (or two pairs of wavelengths, when the redundant configuration is used).
2. The design must consider potential interference between the two signals. A single-fiber configuration requires additional optical components at each end of the path to separate the signals at the two wavelengths. Commercial optical filters are readily available, but are not inexpensive. A typical filter designed to pass the signal at 1550 nm will attenuate the signal at the adjacent wavelength of 1490 nm by at least 35 dB. The filter that passes 1490 nm will attenuate the signal at 1550 nm by at least 32 dB. As we discussed earlier, a single-mode optical fiber that is 100 km long will attenuate the transmitted signal by about 25 dB to 35 dB at these wavelengths [31].

The problem of interference between the two wavelengths can be illustrated by a numerical example. Suppose there is an air-gapped splice in the fiber located 1 km downstream from the NIST hardware. The outbound signals from the NIST system will be attenuated by about 3.2 dB when it reaches the splice. The Fresnel reflection at the splice will have an amplitude of about -11 dB (8%) relative to the incident signals. The reflected signal will return to NIST with an overall attenuation of $3.2 \text{ dB} + 11 \text{ dB} + 3.2 \text{ dB} = 17.4 \text{ dB}$. The filter at the NIST end will attenuate this signal by about 32 dB, so that the overall attenuation of the initial output signal is 49.4 dB. The signal from the remote end has been attenuated by about 32 dB in travelling over the 100 km fiber, so that the signal-to-noise ratio is 17.4 dB. When expressed as an amplitude signal-to-noise ratio, 17.4 dB is equivalent to an amplitude signal-to-noise ratio of about 7.4. The documentation for the end-point hardware does not provide information on the relationship between amplitude signal-to-noise ratio and the probability of a transmission error, so it is difficult to decide how serious this issue would be in practice.

The message of the preceding discussion is that reflections in general; and air-gapped splices in particular, are potentially serious, and should be avoided, especially when the

reflection is generated close to the source of the signal. Fusion splices, and other splices and connectors with no air-gap, produce much smaller reflections. However, the maintenance of the fiber is not under the control or supervision of NIST or OPNT, so it is difficult to know if this is a significant concern.

3. The configuration will also include real-time monitoring and control functions. Some of these will be implemented by performance data acquired over the private reference clock network that is already in place, and additional monitoring will be provided over the fiber circuit from the McLean end of the connection. Some of these capabilities – especially those that are based on the private reference clock network, have already been implemented. If the supervision and control of the circuit is implemented by messages transmitted over the fiber link, then additional wavelengths could be required to implement these functions. This could increase the problems of interference between signals of the different wavelengths discussed in the previous paragraph.

II.A.2 – Reliability and Redundancy of Time over Fiber Distribution System

The initial Special Calibration Test has been realized with only a single fiber circuit between the reference clocks at NIST/Gaithersburg and the remote site at McLean, Virginia. Both ends of this circuit have redundant clock hardware in the current operational configuration, which is shown in Fig. II.A.1, so that only the fiber and the interface hardware represent single points of failure. We have no data on either the mean time between failures or the mean time to repair either of these parts of the service. It is likely that the mean time to repair a significant failure will be at least 24 to 48 hours.

It is difficult to reduce the mean time to repair a serious fiber cut, so that designing adequate holdover capabilities at the remote end is probably the best use of limited resources. The mean time to repair a failure of the end-point hardware can be reduced by having spare parts readily available either at the end-point facilities or by overnight shipping from another location.

The alternate configuration shown in Fig. II.A.3 has an additional point of failure since there is only a single, non-redundant, signal sent from Gaithersburg to McLean. On the other hand, this configuration has the advantage that it provides a direct comparison between the local reference and the time recovered from the data returned to Gaithersburg from the remote hardware at McLean. This comparison would be useful in identifying problems that cannot be detected by the primary configuration without ancillary measurements, and measurements in this configuration may be performed in the future.

II.A.3 – Physical and Cyber Security of Time over Fiber Distribution System

At both the NIST/Gaithersburg and NIST/Boulder sites, the physical security requirements are satisfied by means and methods that are outside of the scope of the Special Calibration Test. This service has no special physical security requirements beyond those that are provided at both sites for other aspects of the time services.

The secondary time scale NIST/Gaithersburg and the primary time scale at NIST/Boulder each operate on air-gapped networks, which have no physical or electrical connection to any other network, whether inside or outside of NIST. The fiber connection between Gaithersburg and McLean has no connection to any other network.

The physical and cyber security at the remote end in general, and at the McLean, Virginia end in particular, is the responsibility of the remote user.

II.A.4 – Accuracy and Stability of Time Over Fiber Distribution with respect to UTC(NIST)

The usefulness of this service will depend on the accuracy with which a remote user can receive a signal that is traceable to UTC(NIST), and then use that signal to apply a time stamp to an event of interest. Although this section does not explicitly define the timing accuracy of the service, we set as a goal that the service would initially provide signals within at least $\pm 1 \mu\text{s}$ of UTC(NIST), and that the accuracy would be improved by at least a factor of 10 after an initial testing period.

OPNT has conducted a preliminary test of a wide-area time distribution system based on the timing signal received at McLean, Virginia from the Special Calibration Test. The test was based on a loopback configuration between McLean and Atlanta, Georgia, a one-way distance of about 1300 km. A time signal was sent from McLean to Atlanta and retransmitted back to McLean via a PTP link. The receiving system in Atlanta has an internal oscillator that is synchronized by means of PTP data from McLean and it drives the return path back to McLean. The signal returned to McLean from Atlanta is compared to the reference oscillator that was the source for the original signal from McLean to Atlanta.

The time difference between the time of the local reference at McLean and the time received over the round-trip path to Atlanta is shown in Fig. II.A.4.a and b. The data in II.A.4.a were averaged for 100 s; the II.A.4.b data were averaged for 400 s. Data like these are typical for measurements of network-based time differences. The standard deviation of the data is not significantly improved by the additional averaging, but Fig. II.A.4.b suggests that the standard deviation would be significantly smaller if a number of relatively large outliers were rejected. An algorithm for detecting and rejecting outliers is a matter for further study. In addition, the difference in the mean values of the two data sets is significant ($0.75 \mu\text{s}$) and is also a matter for further study. In an operational network, such changes in the mean value upon resetting a link would need to be addressed – possibly by comparing multiple sources for fault detection and recovery actions.

These data should be considered preliminary. In spite of the issues discussed in the previous paragraph they are consistent with realizing an overall timing accuracy of 100 ns at a remote location. However, the uncertainties of these results are only one contribution to the overall accuracy of the time distribution.

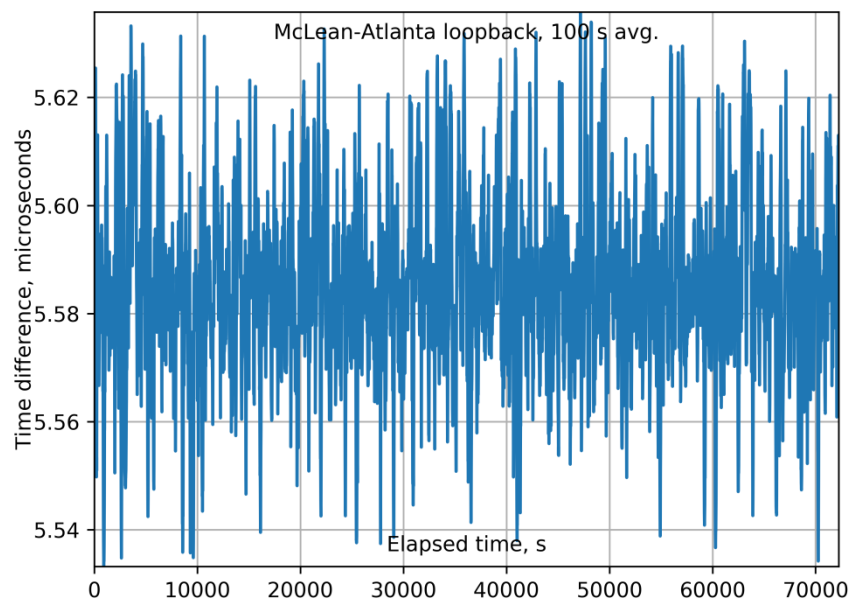


Fig. II.A.4.a

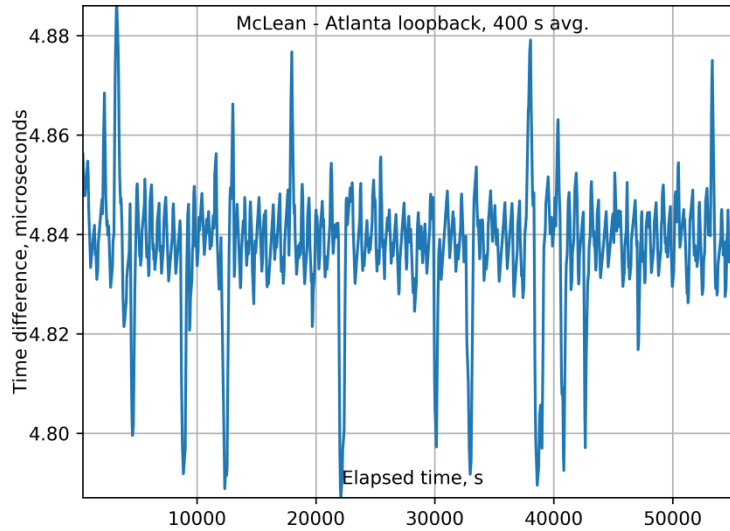


Fig. II.A.4.b

Fig. II.A.4. Results of a loop-back test between McLean, Virginia and Atlanta, Georgia. The time signal is sent from McLean and the hardware in Atlanta echoes the signal back to McLean. The difference between the transmitted time and the time of the reply is shown in Fig. II.A.4.a after averaging for 100 s, and in Fig. II.A.4.b after averaging for 400 s. The longer averaging period reduces the standard deviation, but does not significantly reduce the peak-to-peak variation, as discussed in the text. The difference in the mean values of the two data sets is significant and requires further study.

The overall accuracy with respect to UTC(NIST) will be determined by the uncertainties of several additional contributions:

1. **The time offset of the reference clock from UTC(NIST), and any variation in this offset.** This is much less of a problem when the fiber connection terminates at the NIST Boulder laboratories since this offset can be determined relatively easily as the physical delay in the cable connecting the fiber interface hardware to the time scale reference plane. The situation is more complicated at the NIST/Gaithersburg site, because the time scale there is a secondary representation of UTC(NIST), and the static time offset, and any variation in this offset, with respect to the primary clock ensemble in Boulder must be measured and continuously monitored. This measurement will include some uncertainty, which is small relative to the goal of $\pm 1 \mu\text{s}$ accuracy, but not completely negligible. A detailed discussion of this point is provided later in this section.
2. **The uncertainty in the calibration and characterization of the end-point hardware and the channel connecting the reference time scale to the remote end point.** This is probably the most important aspect of the service, and it will be discussed in the following text in more detail. A central aspect of this discussion is the boundary between the portion of the connection where the calibration and the certification of traceability are the responsibility of NIST, and where that responsibility must be borne by the user.

We have considered three possible locations for the boundary of NIST monitoring, control, and responsibility for the timing accuracy. The simplest solution from the NIST perspective would be to define the boundary as located at the reference plane of the time scale at either NIST/Gaithersburg or NIST/Boulder. Everything after that point would be the responsibility of the user.

A second possibility would be to extend the NIST responsibility to the end-point of the fiber circuit from NIST. There are two general methods that could be used to support this extended responsibility:

- The hardware at the end points of the circuit could provide real-time diagnostic information that would be used to detect relatively rapid changes in the characteristics of the fiber link. The most serious problem would be a change in the equality of the inbound and outbound delays, since the equality of the inbound and outbound delays is a fundamental assumption of the two-way method of time transfer, and the validity of this assumption determines the accuracy of the estimate of the one-way delay (we assume that the asymmetry in the delay due to the variation of the index of refraction of the fiber as a function of wavelength is known and its effect has been removed). A change in the symmetry of the delay would appear to be a change in the time difference between the local and remote clocks, and the sensitivity of this detection algorithm would depend on comparing this apparent change in time offset with the stability of the clocks at the end points. We do not consider this possibility in detail because this method depends on the characteristics of the hardware at the remote sites, and these characteristics are excluded from this discussion. Even if this technique were available, it would be very difficult to detect slow changes in the characteristics of the fiber link because these changes would be masked by the long-period random-walk and flicker processes that dominate the stability of cesium clocks at longer periods. Also, NIST will not have access to the remote sites, and we have not considered installing NIST-supplied monitoring hardware at the remote sites for that reason.
- This additional responsibility could also be supported by periodic calibration trips to one remote location specified by the operator of the fiber link. This technique has the advantages that it provides an estimate of the time difference between the clocks at the end points that is completely independent of the details of the fiber circuit itself, or of the assumptions that underlie the two-way method. This is a very important advantage, since it can provide an independent estimate of any systematic offsets in the two-way method that would be very difficult to detect otherwise. It would also be sensitive to slow changes in the characteristics of the circuit, which are also difficult to detect by the methods discussed previously.

The third and final possibility would be for NIST to support occasional calibration trips to multiple end nodes of the service. This would be possible for only a relatively small number of end nodes, as periodic calibration trips are time-consuming and expensive. Of course, the success of the time distribution service will depend upon the long-term stability of the transmission delays through the fiber links. If the link delay symmetry is stable, then only infrequent calibration trips would be needed to validate the service accuracy. If the link delay symmetry is not stable, and requires frequent calibration trips, then successful operation of the service will be impractical.

3. **The uncertainty of the hardware that receives the time data at the remote end of the link.** Although it is possible in principle to use the signal that is received via the fiber link to directly apply time tags to events, this strategy is strongly discouraged for the following reasons:
 - a. When the data received over the fiber link is used to discipline a local clock, instead of being directly applied as a time tag, the local clock can provide holdover if the link to NIST fails for any reason. The time to repair a break in the physical fiber is likely to

be at least 24 hours, and a robust design should be prepared for this possibility. For example, a commercial rubidium clock can support a timing accuracy of better than 200 ns for 24 hours once it has been calibrated and its static frequency offset characterized. A commercial cesium clock can provide a timing accuracy of about a factor of 100 better under the same conditions. Therefore, the use of a disciplined local clock allows the $\pm 1 \mu\text{s}$ accuracy requirement to be satisfied even if it takes several days to repair a fiber cut.

- b. No communications channel can ever be perfect and noise free, and a robust design recognizes these limitations and designs the end-point hardware to compensate for them to the greatest extent possible. In a well-designed system, the combination of the statistical characteristics of the channel and the end-point receiving hardware complement each other and the resulting system has better performance than either component by itself. The weaknesses of each portion of the system are compensated by strengths in the other portion. For example, it is common for the clock at the remote end of the circuit to have better short-term stability than the reference time scale, which is seen at the remote site only through the fiber circuit. In general, the end-point hardware that receives and processes the data is not under the control of NIST and may not be designed based on considerations consistent with optimal performance. For this reason, the following discussion considers only the performance of the basic fiber link without considering the potential advantages of an implementation that uses receiving hardware and software whose characteristics are matched to the characteristics of the fiber as discussed in this section. A well-designed system would be expected to have better performance than what we will describe in this section, especially for very short averaging periods.
4. **The offset and jitter in the delay between the instant when a user application receives a time signal, and when the application applies that time stamp to an event.** This delay is internal to the hardware that applies a time stamp and is in addition to all the other delays discussed in the previous sections. The magnitude of this delay is a crucial issue because it can be difficult to maintain a timing accuracy of $\pm 1 \mu\text{s}$ if this time stamping process is realized on general-purpose hardware [32]. Although this aspect is important, it will not be considered here, and is mentioned here because it is often both more important and more difficult to address when the reference time signals do not originate from a locally-connected GNSS receiver, which is the configuration described in this text.

II.A.4.1 - Metrics utilized to evaluate the stability of a time over fiber system

The various forms of the Allan deviation [8, 33, 34] are often used to describe the stability of clocks, oscillators, and channels that transport time and frequency information [32]. These metrics are generally useful, but are not adequate by themselves for characterizing the performance of the Special Calibration Test for these reasons:

1. All stability estimating metrics are averages, computed from multiple data points. However, an application that applies a single time stamp to a single event cannot benefit from an average specification, because such a specification does not provide any guidance on the maximum offset that might be obtained in a single measurement. To add to this ambiguity, many stability estimators are computed only after “outliers” have been detected and removed from the data, and may significantly underestimate the offset of a single measurement.

An outlier is a measurement that is not consistent with the deterministic and stochastic parameters that characterize the remainder of the data. However, accurately identifying outliers in a real-time data stream can be difficult, because what is thought to be an outlier may actually represent a change in the deterministic or stochastic characteristics of the data. Outliers are much easier to detect when post-processing a data set.

2. Stability estimating metrics are not sensitive to static offsets in time or in frequency, even though the magnitude of these offsets contributes to the timing accuracy.

We typically use three calibration techniques to address the limitations of the traditional descriptions in terms of Allan deviation. We have tested and validated the performance of these three methods by performing calibration trips between various sites in Boulder and Fort Collins, Colorado, as described in the following paragraphs:

1. **Portable rubidium clock trips** - these trips use a portable rubidium clock, a time interval counter, a portable power supply, and a set of calibrated cables (Fig. II.A.5).




Fig. II.A.5. The portable rubidium clock method of calibration.

These clock trips typically start at the NIST/Boulder laboratory, where the output 1 PPS of the rubidium clock is measured with respect to UTC(NIST). The clock is then carried to a remote site – either the JILA building on the University of Colorado campus, about 1 km from NIST, or to the NIST radio station WWV, about 78 km from NIST/Boulder. The time difference between the rubidium clock and the local reference clock is measured at the remote location and the rubidium clock is then returned to NIST/Boulder where a second measurement is made with respect to UTC(NIST). The magnitudes of the time differences are not important, since the time transfer will use the differences of the time differences recorded at each site.

The uncertainty in the calibration process is dominated by the frequency noise of the rubidium clock. A round-trip from NIST/Boulder to JILA and back typically takes about 1500 s, so the frequency instability of the clock, which is about 5×10^{-12} for that averaging period; would result in a time dispersion of 7.5 ns or less, depending on the statistical model assumed for the frequency fluctuations. The closure in a round trip between NIST/Boulder to JILA and back is typically from 10 ns to 15 ns; the round trip between NIST/Boulder and WWV typically takes 15 000 s, and the closure on return to NIST/Boulder is typically about 75 ns.

If we assume that the uncertainty in the calibration is the full measured value of the lack of closure, then the time offset between JILA and NIST/Boulder can be determined with a worst-case uncertainty of 15 ns, and the time offset between WWV and NIST/Boulder can be determined with a worst-case uncertainty of 75 ns. Figure II.A.8 compares the measurements of the time difference JILA – UTC(NIST) made with this method to the other methods discussed below. In all portable clock

measurements, the closure on return to NIST is larger than the value that would be expected based on the frequency stability of the rubidium clock when it was operated in a laboratory environment, and we attribute this excess to the vibration and temperature changes associated with moving the clock from one place to another. Nevertheless, the agreement between the calibration estimated by the portable clock trip and the value known from previous measurements is smaller than the worst-case uncertainty.

2. **Portable GPS receiver calibrations** - the simplest version of this system consists of a GPS receiver, a TIC, a computer to record the measurements, and a long (200 m) cable that has been calibrated by measuring the delay of a 1 PPS signal through the cable (Fig. II.A.6).




Fig. II.A.6. The portable GPS receiver calibration method. The time differences between the local clock and GPS system time are recorded every second and stored on a laptop computer, which also computes 5-minute average time differences for each GPS satellite in view.

To perform the calibration, the GPS receiver is located at a convenient place outside of the NIST/Boulder laboratories with a clear view of the sky, and the time interval counter is set to record the difference between a direct connection to UTC(NIST) at its reference plane (using the long, calibrated cable) and GPS system time as reported by the GPS receiver. The cable delays of the receiver and antenna are adjusted so that the measured time difference $\text{UTC(NIST)} - \text{GPS time}$ agrees with the differences measured by the various standard receivers at NIST at the same time. This process is typically repeated five times over a period of about two weeks; each observation lasts about one hour, and the data are analyzed by using standard common-view techniques. The calibration constants estimated by the five measurement periods typically agree with each other to within ± 4 ns peak-to-peak (Fig. II.A.7).

When this process is completed, the system is carried to JILA and then to radio station WWV, and the measurements are repeated multiple times. The time offset between JILA and UTC(NIST) estimated with this method, typically agrees with the previously discussed portable rubidium clock trip method to within ± 8 ns (Fig. II.A.8).

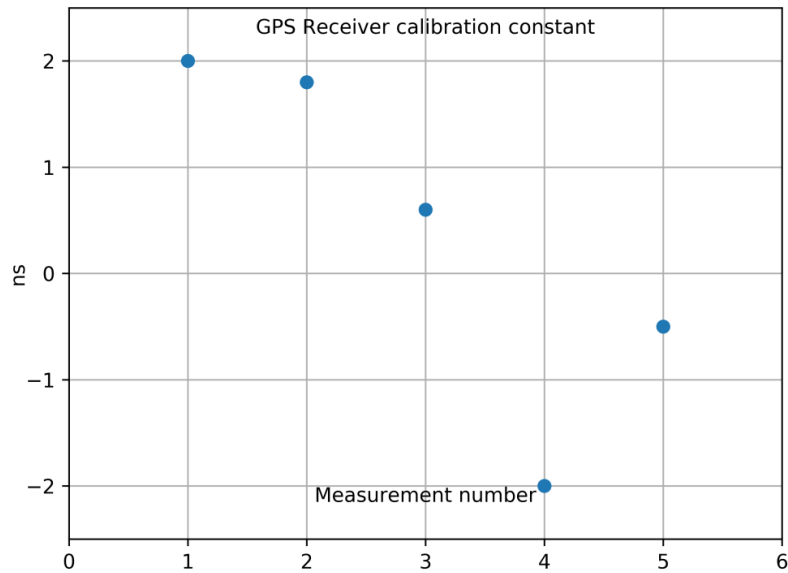


Fig. II.A.7. Variations in the calibration constant of a portable GPS receiver obtained from five measurements. This is the difference between UTC(NIST) – GPS time, as measured by the official NIST receiver and by the portable GPS receiver used for calibrations. The average time difference (353 ns) has been subtracted from all of the values.

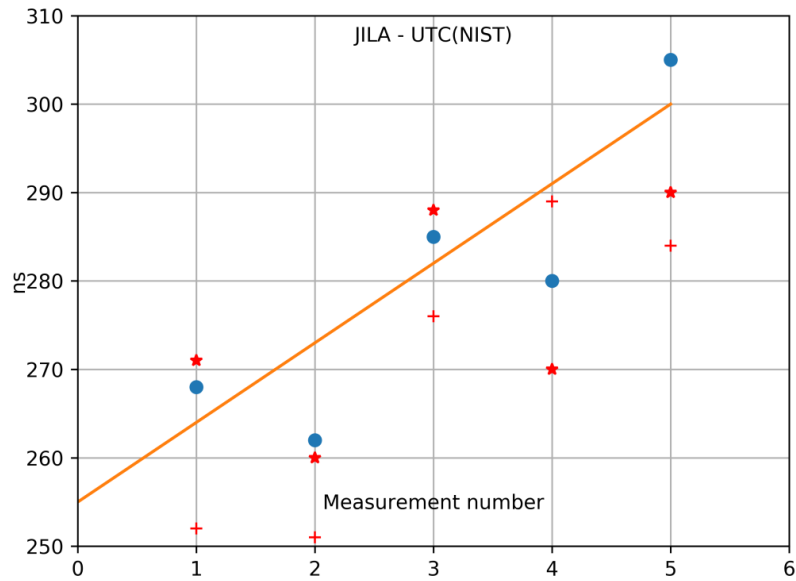


Fig. II.A.8. The time difference between the reference clock in JILA and UTC(NIST) estimated by the initial calibration (blue dots), the portable clock method (red +) and the portable GPS receiver (red *). The line represents a least squares fit to the initial calibration data.

The corresponding uncertainty in the time offset of the clocks at WWV is typically ± 45 ns, about a factor of two better than the uncertainty of the calibration based on the portable clock trip. Figure II.A.9 compares these results with measurements acquired by other methods to be described later. Although the mean of the calibration measurements at WWV is consistent with the time offset known from other measurements, the relatively large scatter in the calibrations was troubling. The portable GPS receiver was located in the field outside of the WWV building for these measurements, and there was some indication that the receiver, and especially the long cable from the receiver to the reference plane of the time scale in the building; were affected by the strong radio frequency fields outside of the building. To address this concern, a third measurement strategy was developed as will be discussed in the next paragraph.

3. **The third calibration strategy is a combination of the previous two methods.** The calibrated GPS receiver is used to measure the time difference between a portable rubidium standard and GPS system time, and the portable rubidium standard is then carried to the clock ensemble to be calibrated and the time-difference measurement is repeated. In the context of the calibration campaign at WWV, this method made it possible to operate the GPS receiver far away from the sources of radio frequency interference and then use the rubidium standard to transfer GPS system time to the location of the time scale in the WWV building. This method has been tested at all three sites (NIST/Boulder, JILA, and WWV), and the calibration constants in each case agreed with values known from other methods to within ± 30 ns RMS, and generally considerably better than this value (Fig. II.A.9).

The largest contributions to the variation from one measurement campaign to another are the stochastic variations in the frequency of the portable rubidium standard when it is carried from place to place. Unfortunately, access restrictions at remote sites may preclude the use of the calibrated cable method (Fig. II.A.5) so that the portable clock method may be the only possibility. The differences in independent calibrations of the same time difference are always negligible relative to ± 1 μ s accuracy, and are small, although not negligible, with respect to the ultimate accuracy goal of ± 100 ns. This confirms our assertion above that the uncertainty in the secondary time scale at NIST/Gaithersburg will not limit the overall accuracy of the service.

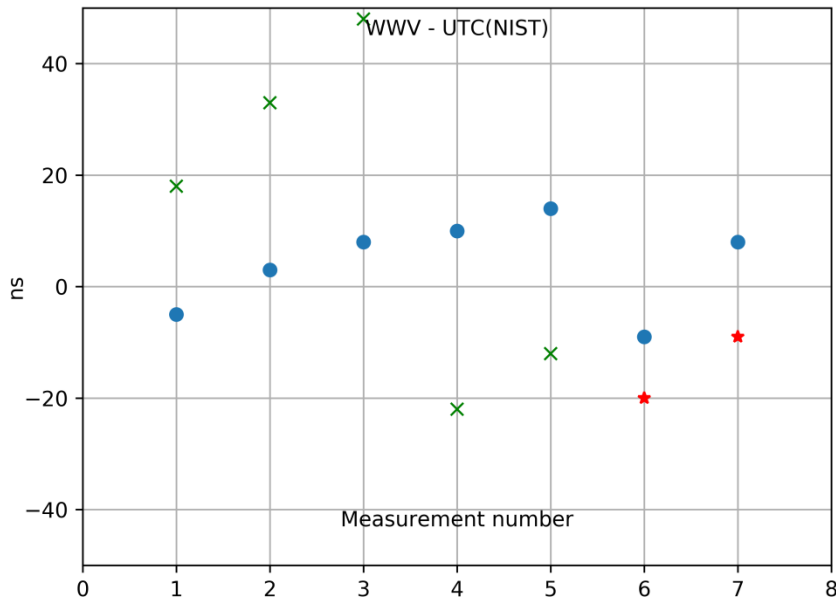


Fig. II.A.9. The time difference, $WWV - UTC(NIST)$ measured by the GPS common-view hardware, calibrated when it was first installed, (blue dots), by direct connection to the portable GPS receiver as described in method 2 (green x) and by the portable clock and GPS receiver as described in method 3 (red *).

Method 3 has been used to verify the time offset of the clock ensemble at NIST/Gaithersburg with respect to UTC(NIST) in Boulder. The portable GPS receiver was set up in front of the Technology Building (Building 225), and the rubidium standard was measured at this point and again at the reference plane for the secondary time scale, which is in the basement of the same building. The time offset measured in this way agrees with the time offset based on the calibration when the system at Gaithersburg was installed with an uncertainty of ± 45 ns RMS, which is consistent with the combined uncertainty of this calibration method, ± 30 ns RMS, and the previous calibration, ± 35 ns RMS. This calibration campaign will be repeated periodically to verify that the calibration constants do not have a significant time variation. Either Method 2 or Method 3 could also be used to calibrate the time offset of a signal with respect to UTC(NIST) at any arbitrary location, and one of these methods may be used to validate the transmission delay in the circuit between the NIST reference time scale and a remote location.

Because the position of the antenna of the calibrated receiver is known only approximately at all of the sites, the measurements at each site were preceded by operating the portable receiver in “survey mode,” in which the receiver uses the data from all of the satellites in view to estimate the position of the antenna. The vertical position determined in this way tends to be anti-correlated with the measured time offset of the local clock with respect to GPS time, and this effect introduces a bias in the measured time offset of the local clock. The magnitude of this bias can be estimated if multiple satellites are tracked from horizon to horizon, but this would have required observation periods on the order of a full day, which was not practical. The bias in the determination of the vertical position can be up to several meters, and the corresponding bias in the determination of the time offsets would be of order 10 ns to 15 ns, which makes a significant contribution to the uncertainties quoted above.

This problem will be more serious if the portable GPS receiver is used to calibrate the time signal at a remote site. The positions of the antenna at the NIST/Boulder, JILA, and WWV sites are known, at least

approximately, because there are other GPS receivers at the same site. For example, the location of the antenna at NIST/Gaithersburg is almost directly underneath the GPS antenna that is already installed at that site on the roof of the building, so that the horizontal components of the position were known and the vertical position could be estimated based on the height of the building. This advantage probably will not be possible at a remote site, where there is unlikely to be any first-order estimate of the antenna position. Therefore, a more extensive campaign to determine the position of the antenna will be required. Our initial plan is that the survey campaign will probably require 24 hours of data with post-processing to determine the clock offset and the position of the antenna. A geodetic receiver will be useful for this analysis. It will also be useful if we can establish some kind of benchmarks at the remote sites, so that we do not have to repeat the position survey on subsequent calibration trips.

II.A.4.2 - Calibration of the delay in the channel from NIST to the remote location

The end-to-end delay in the channel between the UTC(NIST) time scale and the receiving equipment at the remote site is estimated by means of the usual two-way protocol: the round-trip delay can be measured at either end of the circuit, and the one-way delay is estimated as one-half of the measured value. Both the single fiber and the fiber-pair circuits support true full-duplex transmissions – signals can travel in both directions simultaneously. Methods that are based on packet-switched networks, such as the Internet, must operate in half-duplex – the inbound and outbound signals are transmitted sequentially. In principle, the full-duplex transmissions should have a smaller difference between the inbound and outbound delays, but this theoretical advantage is unlikely to be important on a dedicated circuit and may not be fully realized in practice because different end-point hardware must be used to process the signals in the two directions.

This estimate of the one-way delay is correct if the inbound and outbound contributions to the round-trip delay are equal – the magnitude of the round-trip delay is not important. The situation is more complicated when the channel is implemented with an optical fiber, because the signals in the two directions are transmitted by means of different wavelengths. The index of refraction of the fiber is different for the two wavelengths, so that the transmission delays in the two directions are different. This static asymmetry will be discussed in the next section in more detail. This discussion assumes that the asymmetry due to the optical dispersion is either small enough to be ignored or is known from the characteristics of the fiber and that its effect has been removed. The following discussion is concerned with any residual asymmetry that is not due to optical dispersion.

If the measured round-trip delay is Δ , and if the outbound contribution the round-trip delay is $k\Delta$, where k is a fraction between 0 and 1, then the channel delay estimated by the two-way method, which assumes $k=0.5$, will be in error by

$$\varepsilon = (k - 0.5)\Delta . \quad (\text{II.A.1})$$

This equation can provide an estimate of the maximum static asymmetry that can be present to realize an overall timing accuracy. For example, suppose that the fiber link between NIST and the remote location is 100 km long, and that the speed of light in the fiber is 66% of the value in a vacuum. The round-trip distance is then 200 km, the velocity of light is 2×10^5 km/s, and the round-trip transit time is 1 ms. If the goal is to provide a time accuracy of better than 1 μ s, and if the path asymmetry is to contribute no more than 10% to the overall error budget, then the asymmetry parameter ($k - 0.5$) must be not greater than $\pm 10^{-4}$. The coefficient of thermal expansion of optical fibers varies somewhat from one configuration to another but is near $10^{-6}/^\circ\text{C}$ [35], so that changes in the length of a fiber circuit composed of two one-way fibers, caused by changes in the ambient temperature, would not be a problem in the configuration. The thermal expansion of a single fiber, with signals transmitted in both directions, does not make any first-order contribution to the error budget. There might be some higher-order effects due to the dispersion of the fiber, since the

wavelengths used to send the signals in both directions will see slightly different indices of refraction. For example, if the wavelengths in the two directions are 1550 nm and 1490 nm, and if the length of the fiber is 100 km, the optical delays in the two directions differ by about 420 ns [36]. However, the dispersion of the fiber is assumed to be known, so that the uncertainty in this asymmetric delay is a small fraction of the effect.

The transmission delay estimated by the two-way method can be validated by comparing the estimate with the value obtained from a portable clock trip or by using the portable calibrated GPS receiver as discussed in the previous section. The second method, based on the calibrated GPS receiver and the long cable, is potentially more accurate than the third method that also includes the portable rubidium standard, but it requires a physical connection between the reference plane of the clocks at the remote location and a location outside with a clear view of the sky. This configuration is not possible at NIST/Gaithersburg, and may not be possible in general.

The two-way method will almost certainly provide a measurement of the path delay with greater resolution than the absolute calibration methods described above, so that validating the accuracy of the two-way estimate will depend primarily on an accuracy determination of the position of the antenna of the portable GPS receiver and on decreasing the impact of transporting the portable rubidium standard.

The asymmetry in the two-way method due to dispersion - When the two-way method is implemented by using an optical fiber between the two end-points, the signals in the two directions are carried by signals at different wavelengths. The optical fiber is dispersive, so that the indices of refraction, and therefore the transit times of the signals, are different for the signals propagating in the two directions. It is helpful to describe the two-way method in detail to understand the magnitude of this effect.

The two-way message protocol starts when the local system sends a time packet to the remote system at time t_1 as measured by the clock on the local system. The remote system receives the message at time t_2 , as measured by its clock, and responds to the request at time t_3 , also measured by its clock. The local system receives the reply at time t_4 as measured by the clock on the local system. The round-trip transit time, Δ , is given by

$$\Delta = (t_4 - t_1) - (t_3 - t_2), \quad (\text{II.A.2})$$

where the first term on the right side of Eq. (II.A.2) is the total elapsed time as measured by the clock of the local system and the second term is the delay between when the remote system received the request and when it replied. All of the quantities on the right side are measured, so that the round-trip time is known.

If the signal from the local to the remote systems is transmitted at wavelength λ_1 with index of refraction n_1 , and if the signal in the opposite direction is transmitted at wavelength λ_2 with index n_2 , then the outbound travel time from the local system to the remote system is $n_1(L/c)$ and the return travel time is $n_2(L/c)$, where c is the speed of light and L is the length of the path. The round-trip transit time is then

$$\Delta = (n_1 + n_2) \left(\frac{L}{c} \right). \quad (\text{II.A.3})$$

The time difference between the local and remote systems is given by

$$\Delta T_{tr} = (t_1 + d) - t_2 = (t_1 - t_2) + d, \quad (\text{II.A.4})$$

where d is the outbound (local to remote) path delay. The outbound delay can be calculated from the round-trip delay by using Eq. (II.A.3),

$$d = n_1 \left(\frac{L}{c} \right) = \left(\frac{n_1}{n_1+n_2} \right) \Delta , \quad (\text{II.A.5})$$

and the second term in the brackets reduces to the usual value of $\frac{1}{2}$ when there is no dispersion. Note that Eq. (II.A.5) is independent of the length of the path, L , and depends only on the measured round-trip time, which is measured, and the indices of refraction, which are assumed to be known from the properties of the fiber. In other words, the one-way delay in the dispersive medium can be estimated from the round-trip delay without requiring an explicit knowledge of the physical length of the fiber (the physical length of the fiber implicitly determines the round-trip delay).

If the inbound and outbound wavelengths are exchanged, then the asymmetry changes sign, but the time difference between the local and remote systems and the round-trip delay are unchanged. If we use Eq. (II.A.4) and (II.A.5) for each choice of the outbound wavelength, then

$$(t_1 - t_2)_{\lambda_1} + \frac{n_1}{n_1+n_2} \Delta = (t_1 - t_2)_{\lambda_2} + \frac{n_2}{n_1+n_2} \Delta \quad (\text{II.A.6})$$

and

$$(t_1 - t_2)_{\lambda_1} - (t_1 - t_2)_{\lambda_2} = \frac{n_2 - n_1}{n_2 + n_1} \Delta \approx \frac{\delta n}{2\bar{n}} \Delta , \quad (\text{II.A.7})$$

where δn is the dispersion and \bar{n} is the average index of refraction. This is a useful approximation because the dispersion is generally much smaller than the average index of refraction. For example, if the outbound and inbound wavelengths are 1550 nm and 1490 nm, respectively, then $\bar{n} = 1.44$ and $\delta n = 1.2 \times 10^{-6}$ [36]. If the length of the path is 30 km, $\Delta = 288 \mu\text{s}$ and the right-hand side of Eq. (II.A.7) is about 120 ps. Therefore, the asymmetry that results from the optical dispersion in the fiber circuit is not important for a path length of only 30 km.

II.A.4.3 - Results obtained by similar time distribution systems at other locations

Previous tests conducted at other locations indicate that sub-microsecond accuracy can be achieved by the NIST Special Calibration Test by use of commercial components – in other words, the Special Calibration Test is not a research proposal. For example, The National Physical Laboratory (NPL), which is the national metrology institute of the United Kingdom, has previously implemented a service [37] that is similar to the Special Calibration Test discussed in this text. The NPL implementation is based on dedicated fiber circuits between the NPL primary time scale in Teddington and various locations in London and elsewhere in the United Kingdom. The data are transmitted over the fiber circuits by using PTP, and there are cesium clocks at the remote locations. The general characteristics of the service are more or less identical to the configuration that is envisioned for the Special Calibration Test, and the UK results demonstrate that distributing time signals to commercial and financial customers with an accuracy of better than $\pm 1 \mu\text{s}$ with respect to the national reference time scale, UTC(NPL), can be achieved. As was discussed in the previous sections, the use of cesium clocks at the remote locations means that the service accuracy can be maintained for significant holdover periods, so that the service will not be critically compromised by failures of the fiber links. The NPL accuracy specification is likely very conservative, and the actual accuracy at the remote reference stations probably exceeds the specification by an order of magnitude or more. Other demonstrations of time transmission over long distances that did not use the White Rabbit protocol include other tests previously conducted by OPNT between McLean, Virginia and Atlanta, Georgia [38, 39].

II.A.5 – Feasibility of Distributing UTC(NIST) via Fiber to a Large Number of Users

Each PTP reference clock and its associated boundary clocks can support up to 18 downstream nodes, so that the hardware at the end points will not limit the scaling of the service. The recurring costs of the fiber circuits are more likely to be a concern, and these costs could be minimized by a distribution method that was based on a pyramid configuration with only a relatively small number of connections at the NIST facilities. The details of the distribution network will be the subject of further study as more data become available.

II.B Distribution of UTC(NIST) via Two-Way Satellite Time and Frequency Transfer (TWSTFT)

Two-way satellite time and frequency transfer (TWSTFT) is one of the primary space-based time and frequency transfer techniques for comparisons and synchronization of remote clocks [40, 41]. When TWSTFT systems have been calibrated to compensate for all delays, time transfer with an uncertainty near 1 ns can be achieved [42], as well as frequency uncertainty near 1×10^{-15} at an averaging period of one day. Currently, TWSTFT is widely used in the international network of time-keeping metrology institutes for the realization of International Atomic Time (TAI) and Coordinated Universal Time (UTC) by the International Bureau for Weights and Measures (BIPM, Bureau International des Poids et Mesures), and serves as a primary time transfer method for ground-based reference clock ensembles for GNSS, including GPS (USA), GLONASS (Russia), Galileo (Europe) and BeiDou (China). TWSTFT is also used by the United States Naval Observatory (USNO) for worldwide distribution of their time scale, UTC(USNO), for military applications.

The TWSTFT method involves the simultaneous exchange of time signals between pairs of satellite earth stations (point-to-point comparisons). The time signals originate from each station's reference clock. Both earth stations simultaneously transmit time signals to the other station through a transponder (or transponders) on the same geostationary satellite. The received signal contains time information from the remote station's reference clock with a propagation delay of about $\frac{1}{4}$ of a second (~ 250 ms). This delay represents the time required for a radio frequency (RF) signal to travel from Earth, through a geostationary satellite, and back to Earth. The arrival time of the signals received at both earth stations is measured with respect to the time of the local reference clock. The propagation delays through the atmosphere are highly reciprocal in both directions, and therefore will be canceled to the first order when the measurements recorded at both earth stations are subtracted from each other.

Some TWSTFT systems can provide remote clock differences in real-time with an uncertainty of less than 5 ns for averaging periods of two to five minutes. The clock difference data obtained via TWSTFT can be used to discipline or steer a remote clock so that it is synchronized to a reference clock. By utilizing UTC(NIST) as the reference clock, we can use TWSTFT in a star architecture to distribute UTC(NIST) in real-time to any location within the coverage area of the satellite.

This section discusses the potential use of TWSTFT to distribute UTC(NIST) for critical infrastructure timing applications. The use of TWSTFT has several potential advantages – it can easily support the sub-microsecond accuracy requirement, it can continue to operate in hostile environments (such as during some types of electromagnetic interference, including jamming and spoofing of the GPS signals, and also during situations where wired connections, such as optical fiber links, are damaged or unusable). In addition, the method is flexible, it can be implemented via the deployment of portable earth stations.

Section II.B.1 describes the principles of TWSTFT and presents the architecture for a TWSTFT distribution system. Section II.B.2 discusses the reliability of a TWSTFT distribution system, its levels of redundancy, and the safeguards in place to ensure continuous operation. Section II.B.3 describes the physical and cyber security of the TWSTFT station at NIST. Section II.B.4 describes the accuracy and stability of TWSTFT with respect to UTC(NIST). Section II.B.5 discusses the feasibility of NIST establishing a TWSTFT time distribution service for critical infrastructure users, as well as the feasibility of scaling such a service so that it can meet the needs of a large number of clients.

II.B.1 – Technical Description of Architecture for TWSTFT Distribution System

A TWSTFT time comparison link (a pair of remote timing facilities) requires two satellite earth stations that each transmit and receive signals that pass through a geostationary satellite. The time information distributed by a TWSTFT signal is carried by pseudorandom noise (PRN) codes generated by a modulator/demodulator (modem) device. The use of unique PRN codes at each station allows multiple pairs of remote stations to compare their reference clocks simultaneously based on the code-division multiple-access (CDMA) principle.

The unique PRN code sequences sent by each station are modulated using binary phase-shift keying (BPSK) on an intermediate frequency (IF) that is frequency locked to the local reference clock frequency. The phase modulation is also synchronized with the local reference clock time in the form of 1 PPS signals. The modem derives a 1 PPS signal (called the transmit or TX 1PPS) from the BPSK sequence. In general, the timing of the BPSK sequence and TX 1PPS contains a delay with respect to the 1 PPS of the reference clock, or REF 1PPS.

The modem generates a replica of the BPSK sequence being transmitted by a remote station using a known PRN code. By correlating this replicated code sequence to the incoming BPSK signal received from the satellite, the modem can extract a 1 PPS signal, called the receive or RX 1PPS. A time of arrival measurement, or $TX\ 1PPS - RX\ 1PPS$, is made by a time interval counter (TIC). The TIC is also used to measure $REF\ 1PPS - TX\ 1PPS$ to determine the delay of the TX 1PPS with respect to the REF 1PPS.

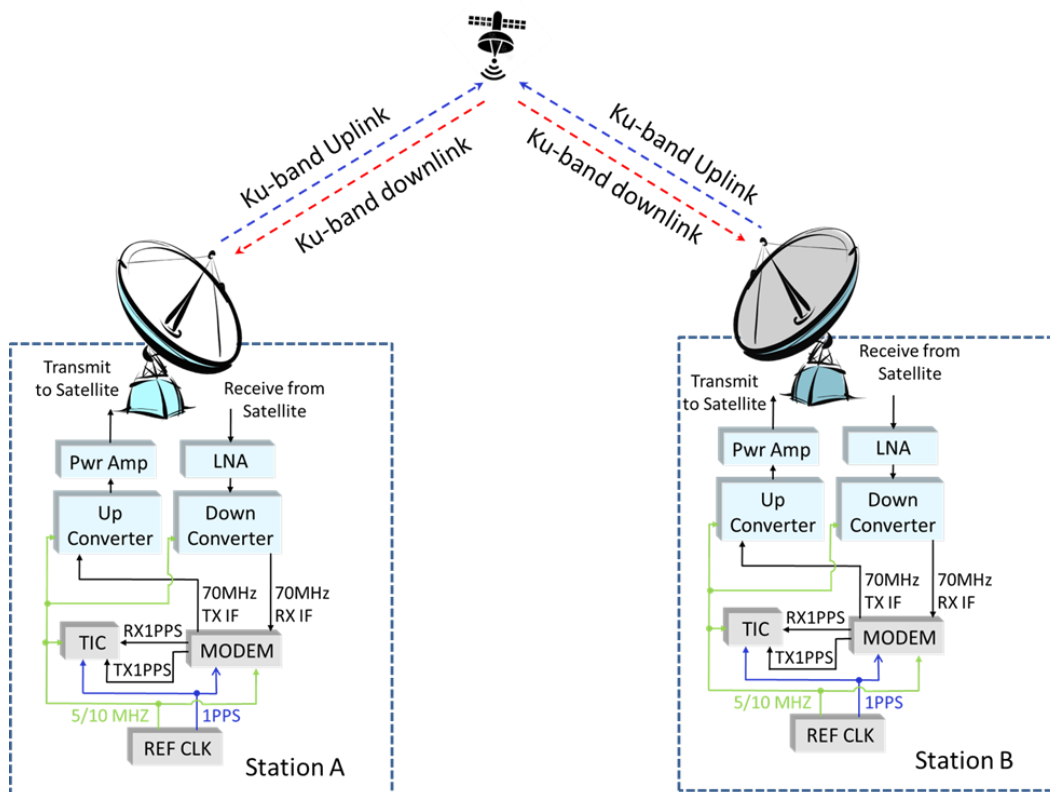


Fig. II.B.1. Block diagram of components for TWSTFT operation between Stations A and B.

Figure II.B.1 shows a block diagram of the TWSTFT operation and the major components. Each earth station consists of a reference clock, a TWSTFT modem, a TIC, a up converter, a down converter, a power

amplifier, a low noise amplifier, and a satellite dish with a feed horn antenna that is used to both transmit and receive signals. The reference clock produces a 1 PPS time signal and a 5 or 10 MHz frequency signal. A TWSTFT modem is essential for both the transmission and reception of signals. The transmit part of the modem generates the PRN codes, utilizes BPSK to modulate the codes onto the IF frequency to produce the TX signal, and synchronizes the TX signal to REF 1PPS. The receive part of the modem decodes the PRN code from a remote station and generates the RX 1PPS. Figure II.B.2 shows the spread spectrum TWSTFT signal, centered on a 70 MHz IF, as it appears on a spectrum analyzer.




Fig. II.B.2. Spectrum of pseudorandom noise codes for TWSTFT at 70 MHz intermediate frequency. The yellow trace is for a 2.5 MChip/s (Mega-chip per second) code, the blue trace is for a 1.0 MChip/s code, and the purple trace is for a 2.5 MChip/s code after passing through a special bandpass filter.

Because different modem manufacturers use different PRN codes, each of the participants in a TWSTFT network must use modems from the same manufacturer. Some modern modems utilize the software-defined-radio (SDR) scheme [43, 44]. Some modems include a built-in TIC, so a standalone instrument is not needed.

The up converter translates the IF frequency to the proper signal band for transmission, and the down converter translates the received RF signal back to the IF. The power amplifier (Pwr Amp) amplifies the transmitted RF signal after it is up converted, and the low noise amplifier (LNA) amplifies the received RF signal before it is down converted. The RF signals are transmitted and received from the same satellite dish with a feed horn antenna.

When the signal reaches the satellite the transponder amplifies the RF uplink signal, converts it to the downlink RF frequency, amplifies it again, and then retransmits the signal back to earth. Radio frequencies in several bands, including C-band (3.7 to 6.4 GHz), Ku-band (10.7 to 14.5 GHz), and X-band (8 to 12 GHz), have been used in TWSTFT operations. Figure II.B.3 shows the components used in a NIST TWSTFT earth station.




Fig. II.B.3. Components of a NIST TWSTFT earth station.

Because TWSTFT involves both the transmission and reception of signals, it is inherently more complex and expensive than GNSS time transfer, which only requires the reception of signals. After the initial purchase of receiving equipment, the use of GNSS signals for time and frequency transfer is free. In contrast, TWSTFT participants not only need to invest in the earth station and TWSTFT equipment but must also pay for the use of a satellite transponder or multiple transponders. The cost of satellite usage is proportional to the bandwidth, which is in turn determined by the chip-rate of the PRN codes used.

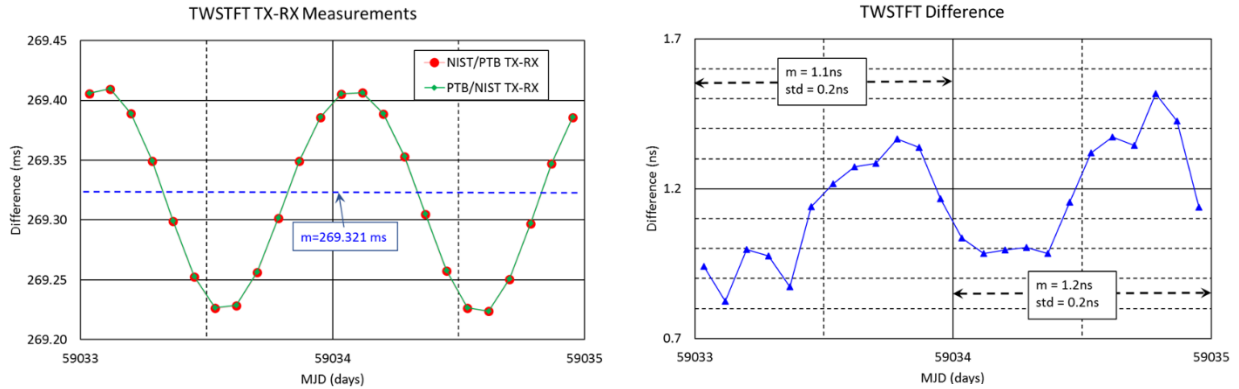
As illustrated in Fig. II.B.2, the higher the chip-rate, the more bandwidth is required. For a given code chip-rate, a bandwidth of at least $1.4\times$ the chip-rate is needed to prevent the signal from distorting and from interfering with neighboring signals on the satellite transponder. With a wider bandwidth, the satellite cost is higher, but there are advantages. Each earth station can transmit a higher power signal to improve the carrier-to-noise ratio (C/N_0) of the received signal and can obtain higher measurement resolution to improve the accuracy of TWSTFT. Cost estimates for bandwidth are described in Section II.B.5.

NIST is currently a participant in the Europe-to-Europe and transatlantic TWSTFT network. The participants in this network follow a pre-arranged schedule for comparison of their clocks. At a time assigned by the schedule, a pair of stations begins transmitting on their respective PRN codes, while simultaneously locking to the PRN code transmitted by the other station. Both stations continue to transmit and receive signals for a predetermined interval, called a session, and recording the measurements, $TX\ 1PPS - RX\ 1PPS$ and $REF\ 1PPS - TX\ 1PPS$, described earlier. After exchanging the time of arrival measurements and differencing them with the link calibration corrections, the time difference between the two clocks is obtained by

$$\begin{aligned}
 CLK_A - CLK_B = & 0.5 \times (TX\ 1PPS - RX\ 1PPS)_A + (REF\ 1PPS - TX\ 1PPS)_A \\
 & - 0.5 \times (TX\ 1PPS - RX\ 1PPS)_B - (REF\ 1PPS - TX\ 1PPS)_B \\
 & + CAL_{(A,B)}
 \end{aligned}
 \tag{II.B.1}$$

where $CAL_{(A,B)}$ is the calibration of delay correction for the TWSTFT between the clock A (CLK_A) and clock B (CLK_B).

Some of the TWSTFT modems can exchange the time of arrival measurements in real-time by modulating the measurements in a data stream transmitted together with the TWSTFT signal to the satellite transponder. By transmitting and receiving both TWSTFT signals and measurement data, the modem can produce the difference between the remote clocks in real-time.



(a) The time of arrival ($TX\ 1PPS - RX\ 1PPS$) measurements made by NIST (red) and PTB (green) in even hours in UTC time. Each measurement session is 2 minutes.

(b) The TWSTFT difference between UTC(NIST) and UTC(PTB) from the two time of arrival measurements shown in (a). Each difference is the average of a 2-minute session.

Fig. II.B.4. Example of a 2-day TWSTFT comparison between NIST and PTB.

Figure II.B.4 shows the results of a two-day TWSTFT comparison between NIST and the Physikalisch-Technische Bundesanstalt (PTB), the national metrology institute of Germany. During the two-day period, one two-minute session was conducted every two hours. Part a) of the figure overlays the data recorded at PTB with the data recorded at NIST. Note that the magnitude of the data is about 269 ms, reflecting the one-way propagation delay between the two earth stations, when the signal passes through the transponder of a geostationary satellite. The range of the data is about 200 μs due to satellite motion, as a geostationary satellite is not completely stationary but moves in a figure-eight pattern. Part b) shows the results of the two-way comparison after being processed as shown in Eq. (II.B.1). These data indicate that the average time difference between UTC(NIST) and UTC(PTB) is about 1 ns, with a peak-to-peak variation of less than 1 ns over the two-day period.

II.B.1.1 - Star configuration network

With the RX 1PPS signal and the real-time remote clock difference information, we can use TWSTFT to distribute UTC(NIST) to remote timing facilities in a star configuration network, where the earth station at NIST continuously compares UTC(NIST) to a network of remote earth stations that each have their own local clock. The comparisons are performed sequentially, according to a prearranged schedule.

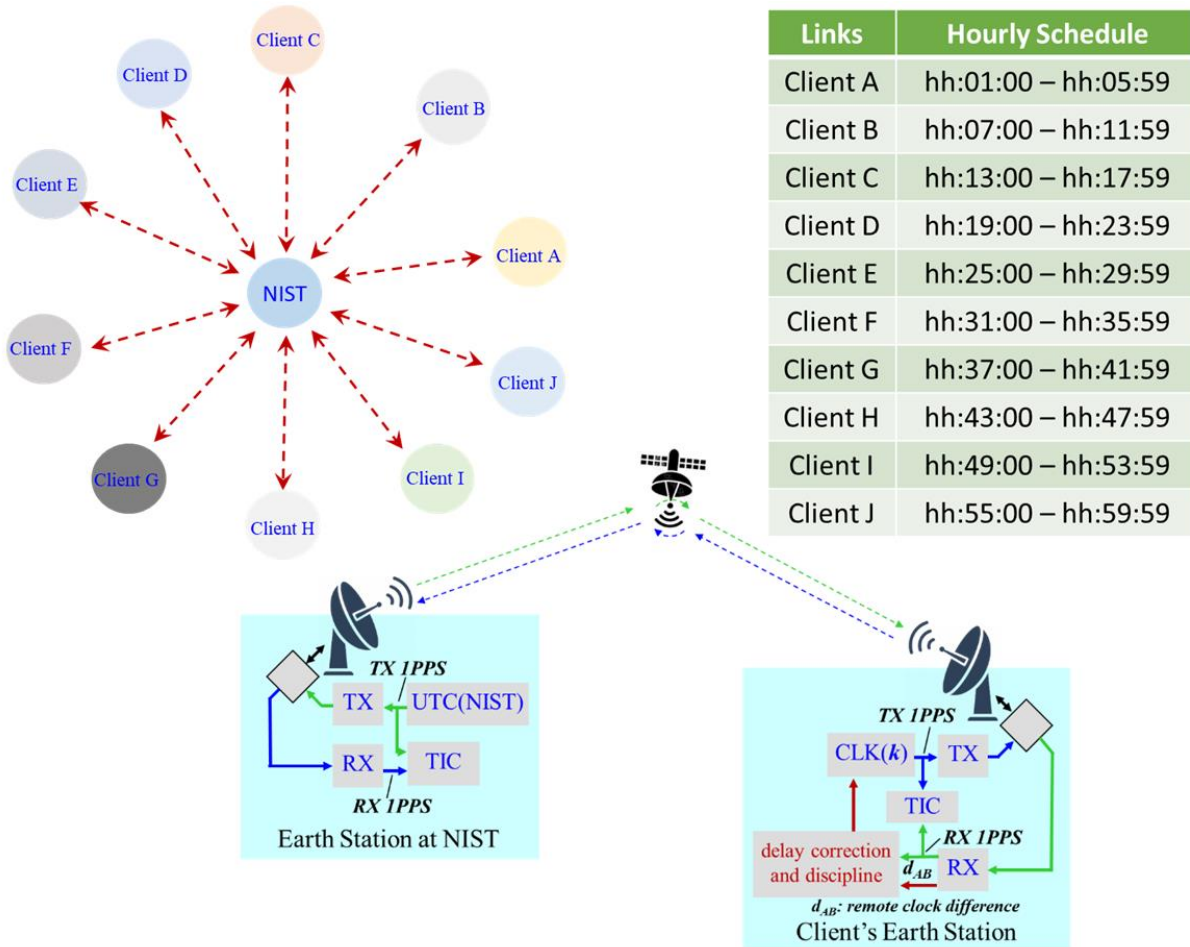


Fig. II.B.5. Star configuration network for distributing UTC(NIST) via TWSTFT.

Figure II.B.5 shows an example of star configuration for a 10-station network. The figure also contains a block diagram for the distribution of UTC(NIST) to one of the remote timing facilities using TWSTFT and a prearranged hourly schedule. As described earlier and illustrated in the block diagram, the RX 1PPS at a client's earth station contains UTC(NIST) plus the propagation delays introduced by the signal path. If the TWSTFT link between NIST and the remote earth station is calibrated, the real-time remote clock difference (d_{AB}) reveals the time difference between UTC(NIST) and the reference clock of the remote earth station. The remote clock difference can be used to advance the RX 1PPS to compensate for the propagation delay. The corrected RX 1PPS then becomes a replica of UTC(NIST) that can be applied to synchronize the reference clock at the remote earth station. The delay correction can be applied with a micro-phase stepper or similar device that can shift the phase of the timing signal. The real-time remote clock difference is also utilized to discipline the frequency of the reference clock of the earth station to keep its time

difference with respect to UTC(NIST) minimized while the NIST earth station is engaged in TWSTFT comparisons with other earth stations.

Each TWSTFT session can be from two to five minutes long, depending on the number of remote earth stations in the distribution network and remote clock performance requirements. For example, if there are only 10 remote timing facilities in the distribution network, the hourly sessions can be five minutes long. However, to support a $\pm 1 \mu\text{s}$ accuracy requirement, fewer sessions will be necessary. A session every few hours or even one session per day may be sufficient to correct the remote clock's frequency offset and drift and maintain sub-microsecond synchronization. The use of short and infrequent time sessions could allow using TWSTFT to distribute UTC(NIST) to a large number of remote timing facilities.

II.B.2 – Reliability and Redundancy of TWSTFT Distribution System

NIST has participated in TWSTFT remote clock comparisons for more than 30 years. Although we have not yet operated a TWSTFT distribution network, NIST has demonstrated that we can build and operate TWSTFT systems reliably. In the past 10+ years, NIST has missed very few TWSTFT measurements that contribute UTC(NIST) data to the BIPM, and these comparisons are conducted 24/7 during each day of the year. When the failure of NIST TWSTFT does occur, it is usually related to equipment malfunctions. NIST personnel are alerted via monitoring and alarm systems when irregular conditions or failures occur, allowing us to take preventive maintenance actions and to troubleshoot and repair or replace equipment. In nearly all cases, measurements are resumed within 10 to 12 hours.

The monitoring and alarm systems are an integral part of the NIST TWSTFT operation. The systems continuously measure the temperature and other environmental parameters of the TWSTFT equipment, check all the parameters of reference signals, TX and RX signals, and inspect the remote clock difference data. When an abnormal condition is detected, alarm e-mail and text messages are immediately sent to staff members, who investigate the alarm and go onsite if the problem cannot be fixed remotely.

If a TWSTFT time distribution network is constructed, its reliability will likely be worse than that of the NIST TWSTFT comparisons with other national metrology institutes. This is simply because there may be no technical expert at the remote TWSTFT station to troubleshoot and replace defective components, and time would be required for NIST to deliver replacement parts to remote stations. However, some of this could be solved by having spare equipment available at each client station. Also, it would be possible to design earth stations that are simpler and more robust, perhaps at the expense of some small loss of performance capability.

Because TWSTFT involves both transmitting and receiving signals to and from a satellite, interruption in any of the signal paths (such as the equipment malfunction or rare geomagnetic events that could potentially affect signal propagation) will stop the collection of TWSTFT data. However, if the reference clock at a remote earth station is still functioning, it can provide time referenced to UTC(NIST) at reduced accuracy, essentially going into “holdover” mode in the same fashion as GNSS disciplined clock.

It is critical to provide uninterruptable electrical power and Internet access to the TWSTFT system for continuous operation. Continuous Internet connectivity is essential for the remote monitoring and alarm systems, and for troubleshooting.

II.B.3 – Physical and Cyber Security of TWSTFT Distribution System

The satellite dish and the RF components for a TWSTFT station are installed on the roof of a NIST building and are only accessible to trained and authorized NIST personnel. The other TWSTFT equipment is located in a laboratory that only trained and authorized NIST personnel are allowed to access. Physical security measures should also be implemented on each of the client TWSTFT stations.

A computer that controls the TWSTFT operation and disciplines the clock needs restricted Internet access so that the system status and alarms can be sent to NIST (via e-mails and text messages) and NIST can do remote troubleshooting. This restricted access includes remote desktop access from a few specific computers for monitoring and troubleshooting, and for transferring data to a specific server. The computer is located on a secure network with a restricted firewall administered by NIST IT personnel.

II.B.4 – Accuracy and Stability of TWSTFT Distribution System with respect to UTC(NIST)

The accuracy of TWSTFT is determined by the calibration of the link between NIST and a remote timing facility. Following the TWSTFT calibration guidelines [45], the best accuracy is 1 ns if the link is calibrated by a TWSTFT mobile station, and the accuracy is 1.5 ns if the link is calibrated by a dual-frequency GPS receiver. Each remote TWSTFT station has to be calibrated on the client site after all of the reference signals and components are connected and installed. We estimate the uncertainty of delay correction and steering to be less than 5 ns. So, the combined uncertainty, U_c , for distributing UTC(NIST) to a remote timing facility via TWSTFT is better than 5.2 ns, $U_c = \sqrt{(1.5^2 + 5^2)}$.

The stability of TWSTFT (measured by the time deviation) is about 0.2 ns at an averaging period of two minutes (the minimum time period we plan to compute a new steering parameter). If we assume the stability of the delay correction and steering device is better than 1 ns, then the combined stability for distributing UTC(NIST) to a remote timing facility via TWSTFT will be better than 1 ns at an averaging period of two minutes.

II.B.5 – Feasibility of Establishing a TWSTFT Distribution Service and Estimated Costs per Client

It is feasible to develop a service at NIST that utilizes the TWSTFT method to distribute UTC(NIST) to clients on a subscription basis, in much the same fashion as the time over optical fiber service (Section II.A) discussed previously, and the MSCVDC services (Section II.C) to be discussed next. However, such a service would likely be expensive. The cost would include both the investment in earth station equipment and operational expenses. The cost of an earth station with a 1.8 m Ku-band satellite dish and all other equipment needed for TWSTFT would likely be near \$150,000 USD. It could be amortized over a period of months or years as part of a monthly or annual fee paid by the client, but the cost would still be high.

Satellite time is one of the major parts of the operational expenses. To obtain 1.6 MHz bandwidth on a Ku-band transponder, the annual subscription for the non-preemptive use of the satellite is about \$160,000 USD, but this cost can be shared amongst clients. For example, if we compare and discipline a remote clock every hour, we can have 10 clients and the satellite cost for each client would be about \$16,000/year. If we only compare and discipline remote clocks every two hours, we can double the number of clients (to 20) and cut the cost per client in half (to about \$8,000/year). These costs would not include the cost of NIST labor, or NIST administrative surcharges. Even so, clients who require the best possible level of accuracy might find that a TWSTFT service is comparable to the recurring costs of other high accuracy methods, such as the dedicated fiber connection previously discussed in Section II.A.

Scaling a TWSTFT service so that it can distribute UTC(NIST) to larger numbers of clients, for example to hundreds of clients, would likely require a different approach. One possibility would be to combine the TWSTFT distribution network to regional timing distribution systems that utilize fiber optic links, as described in Section II.A. In this case, UTC(NIST) could be distributed to perhaps 10 to 20 tier-1 critical infrastructure clients using TWSTFT over long distances, and the tier-1 clients could in turn distribute UTC(NIST) to tier-2 clients via fiber optic links.

II.C Distribution of UTC(NIST) via Multi-Source Common-View Disciplined Clocks (MSCVDC)

A multi-source common-view disciplined clock (MSCVDC) is a clock, based on a concept developed at NIST [46], that replicates the performance of a reference clock such as UTC(NIST), at a remote location. The name of the device explains what it does – a MSCVDC is continuously compared to the reference clock using the common-view measurement technique, hence the “CV” part of the acronym. The common-view technique requires the MSCVDC and a receiver connected to the reference clock to each receive, nearly simultaneously, a signal broadcast from the same transmitter. The MSCVDC is a disciplined clock, meaning that its time and frequency are continuously adjusted to be synchronous with the time and frequency of the reference clock, hence the “DC” part of the acronym. Finally, the MSCVDC is a multi-source clock; hence the “MS” part of the acronym. We mention this part last because multi-source capability is both valuable and essential when developing resilient time distribution architectures. It allows the clock to be disciplined by multiple time scales, with a primary time scale selected by the user and with secondary time scales used as backups to protect against primary time scale failures. It also allows the clock to use multiple sources as common-view signals, protecting the user from failures if a given common-view signal (CVS) becomes unavailable.

An MSCVDC is well suited for critical infrastructure systems that need a very accurate and very reliable clock. To be more specific, “very accurate” means that an MSCVDC can easily keep time within $\pm 1 \mu\text{s}$ of UTC, and “very reliable” means that it has enough redundancy to continue to be accurate, even when one or more parts of the clock have failed. The multiple layers of built-in resilience to failures provide the MSCVDC, perhaps to a greater extent than any other time distribution method, with true fail-safe redundancy.

NIST currently distributes UTC(NIST) via its NIST disciplined clock (NISTDC) service, which is a partial implementation of the MSCVDC technique. Over 20 NISTDC units are currently in operation, with several located outside of the United States [47]. A primary application of the NISTDC has been the synchronization of stock exchanges and financial markets [48]. The following sections describe the MSCVDC technique in detail, using the partial implementation of the current NISTDC in some examples, but focusing on what a full implementation of the method can accomplish.

Section II.C.1 provides a technical description of the architecture for a MSCVDC distribution system. Section II.C.2 discusses the reliability and redundancy of a MSCVDC distribution system, and Section II.C.3 discusses its physical and cyber security. The accuracy and stability of a MSCVDC with respect to UTC(NIST) is discussed in Section II.C.4. Finally, Section II.C.5 discusses the feasibility of a commercially-available MSCVDC product.

II.C.1 – Technical Description of Architecture for MSCVDC Distribution System

A MSCVDC time distribution system is based on common-view comparisons of clocks. Before a clock can be synchronized to agree with another clock, the time difference between the clocks must be measured and known. Ideally, the time difference would be measured by direct comparison, after bringing both clocks to the same location. However, if the clocks are geographically separated, for example, if we want to synchronize a clock in Chicago to agree with the UTC(NIST) time scale in Boulder, then direct comparison is not possible. Fortunately, synchronization can still be accomplished via a common-view comparison.

In our example, a common-view comparison can be arranged if there is a signal that can be simultaneously observed both in Chicago and Boulder. If such a common-view signal (CVS) exists, a name chosen because the signal can be “viewed” from both locations, then the clocks in Chicago and Boulder can be nearly simultaneously compared to the CVS. The difference between the two “indirect” comparisons effectively substitutes for a direct comparison and reveals the time difference between the Chicago clock and UTC(NIST). Even though the CVS signal originates from its own clock, the time signal it delivers does not have to be accurate, because it is cancelled out when the two indirect comparisons are subtracted from each other. This works exceptionally well if the propagation times are equal or if the differences in propagation time can be measured and corrected. In short, the CVS is not a reference clock used for synchronization, but instead is just a vehicle that relays time information from one site to another.

Figure II.C.1 shows a common-view time transfer system where an RF transmitter produces the CVS, and where the CVS is received at sites *A* and *B*. Both sites have a local clock and a receiver that each produce a 1 PPS signal, and the 1 PPS signals are compared with a time interval counter (TIC). The site *A* measurement compares the CVS received over the path d_{SA} to Clock A, producing the time difference *Clock A – CVS*. The site *B* measurement compares the CVS received over the path d_{SB} to Clock B, producing the time difference *Clock B – CVS*.




Fig. II.C.1. A common-view time transfer system. The system allows accurate time to be easily transferred from clock A to clock B, without requiring the time signal broadcast by the CVS to be accurate.

After the time difference values are obtained; data transfer needs to take place before the measurement results can be processed. Figure II.C.1 shows bidirectional data transfer, but the data transfer could be unidirectional, if only one site needs to know their time difference with respect to the other site. Or, both sites could send their data to a “neutral” site, such as a data repository residing on a network server, from which each site can retrieve the other site’s measurements. Regardless of how the data transfer is configured, once the two measurements are available, simple subtraction reveals the time difference between the two clocks, because the contribution from the CVS is cancelled out.

To look at this in more detail, delays that are common to paths d_{SA} and d_{SB} cancel even if they are unknown, but delays that are not common to both paths contribute measurement uncertainty, and result in an uncertainty term of $d_{SA} - d_{SB}$. This term represents the relative or differential delay between the two common-view measurement systems. Thus, the basic equation for common-view measurements is

$$Clock_A - Clock_B = (Clock_A - CVS) - (Clock_B - CVS) + (d_{SA} - d_{SB}). \quad (\text{II.C.1})$$

The delays included in the $d_{SA} - d_{SB}$ term can be measured or estimated and applied as a correction to the measurement to reduce the uncertainty. The largest part of both d_{SA} and d_{SB} is the propagation delay, or the time interval required for the CVS to travel from the transmitter to the receiving antenna. For radio signals travelling in free space, the propagation delay is nominally equal to the distance divided by the speed of light. If the CVS originates from a GNSS satellite, the receiving antenna’s position is determined by GNSS, and ranging measurements are automatically performed by the receiver to compensate for propagation delay. However, if the CVS originates from a non-GNSS source, such as a geostationary satellite or a terrestrial based transmitter, and if we can assume that the transmitter position is known, then the position of the receiving antenna must still be independently determined before d_{SA} and d_{SB} can be compensated for propagation delay.

Once the propagation delay has been compensated for, other smaller delays, typically measured in nanoseconds, might still need to be measured or estimated. For example, if the CVS is a GNSS satellite, delays are added as the signal passes through the ionosphere and troposphere on its way to Earth. Other delays are added by multipath signal reflections, and by antenna coordinate errors. After the signal reaches the antenna, delays are introduced by the antenna, antenna cable, and receiver; and these delays must also be measured or estimated (details about how a MSCVDC is calibrated are given in Section II.C.4). The goal is to reduce measurement uncertainty as much as possible by making $d_{SA} - d_{SB}$ as small as possible, and some common-view systems routinely compare clocks with uncertainties of < 10 ns.

Unlike the remote synchronization methods described in Sections II.A (Time over Fiber) and II.B (TWSTFT), common-view is a passive, receive-only method of transferring time. The CVS travels from the transmitter to the receiver, but the receiver does not exchange messages or interact in any way with the transmitter, thus the transmitter has no knowledge of how many receivers exist. As previously noted, data transfer is required before measurement results can be obtained, but because the data stream is not a timing signal its latency is not important, at least not at the level of microseconds or milliseconds. Depending on the system, measurement data delayed by seconds, minutes, hours, or even days, might be acceptable.

To better illustrate this point, for some common-view comparisons, such as comparisons with PTB for data submitted to the BIPM (Section I.A.5.2), it may be days or weeks before the results are processed and published. This type of latency is still usually acceptable for the purpose of providing time difference data to UTC(k) time scales that, because they are based on very stable oscillators, require infrequent steering adjustments. However, if the latency is low enough to generate corrections and compensate for frequency drift in a quartz oscillator, for example, if it is reduced to seconds or minutes; then common-view becomes

a very powerful technique. It can then be used not only just to compare high accuracy atomic clocks, but to automatically control any clock, in much the same way that GPS signals control a GPSDC. The next section describes a common-view disciplined clock (CVDC).

II.C.1.1. Common-view disciplined clock (CVDC)

A common-view disciplined clock (CVDC), an instrument first described in [49], is a clock whose frequency and time are disciplined to agree with a reference source of UTC through the use of common-view comparisons. A block diagram of a CVDC is shown in Fig. II.C.2.

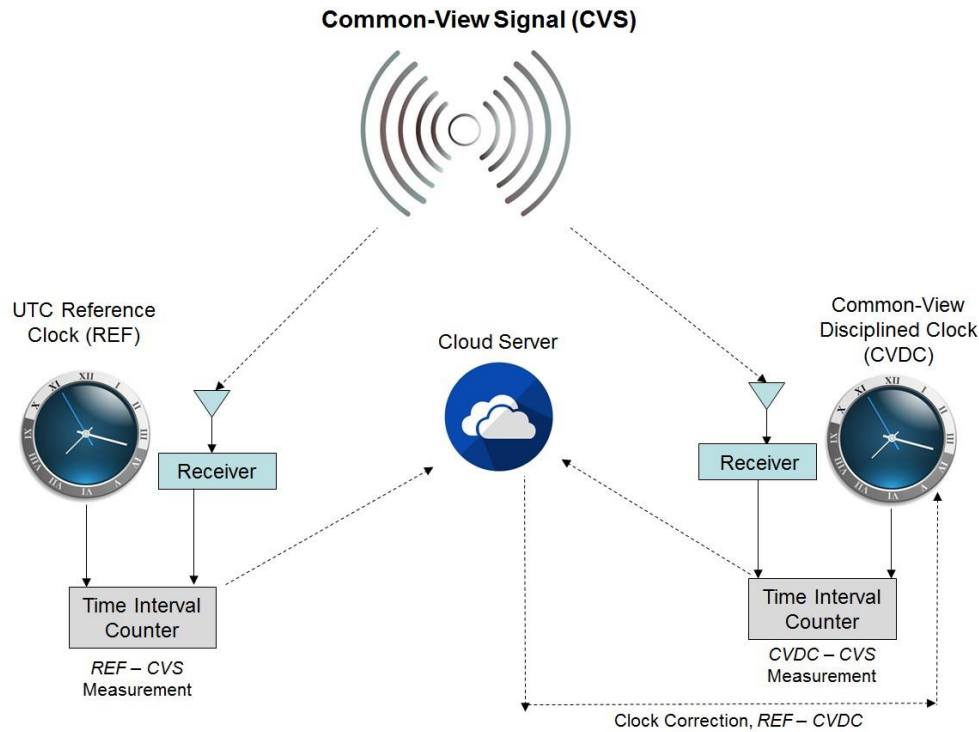


Fig. II.C.2. A common-view disciplined clock. Here, the common-view comparisons are used not only to measure a remote clock, but also to control it. This is done by processing the common-view comparisons in real-time, converting them to clock corrections, and applying the corrections to the remote clock so that it agrees with the reference clock.

Time from a reference UTC time scale is transferred to the CVDC by use of the common-view disciplining method [46, 48, 49]. As illustrated in Fig. II.C.2, a UTC reference clock, such as the UTC(NIST) time scale in Boulder is compared to a CVS by measuring the time difference $REF - CVS$. A simultaneous comparison at a remote site measures the time difference $CVDC - CVS$.

Periodically, the UTC reference clock and the CVDC each send their measurement result to an Internet cloud server where the difference between them, $(REF - CVS) - (CVDC - CVS)$, is obtained to remove the contribution of the CVS to produce the time difference, $REF - CVDC$. The CVDC then retrieves $REF - CVDC$ from the server and converts this time difference to a frequency correction, by use of an adaptive proportional-integral-derivative (PID) controller or a similar control loop. The correction is then applied to the CVDC's local oscillator. The process is continuously repeated to keep the CVDC locked to the UTC reference.

A major advantage of a CVDC is that any UTC reference time scale can serve as the CVDC's reference clock if *REF* – *CVS* measurements from that time scale are made accessible to CVDCs in real-time. This means that a CVDC, unlike a GPSDC, is not controlled by GNSS time. This makes CVDC distribution attractive to national metrology institutes such as NIST that want to replicate their time scales at remote facilities. NIST has deployed CVDCs as part of its calibration programs since 2010 [50], and similar systems have been subsequently introduced by timing laboratories in Canada [51], China [52, 53], the European Union [54], Japan [55], Mexico, and perhaps elsewhere.


Despite its versatility and overall usefulness, a CVDC has vulnerabilities that must be addressed when developing resilient time distribution architectures. Possible failure modes include a reference UTC time scale that is either unavailable or inaccurate, the inability to receive the CVS at either the reference site or the CVDC site, or a network outage at either site. The loss of CVS reception or a network outage for extended periods will cause a CVDC to go into holdover mode and become a free-running clock, just as a GPSDC will do if it is unable to receive GPS. These vulnerabilities are each addressed by the MSCVDC, described in the next section.

II.C.1.2. Multi-source common-view disciplined clock (MSCVDC)

A MSCVDC [46] is an enhanced version of a CVDC that adds multiple layers of redundancy to the basic design. Like a CVDC, it includes a receiver, a time measurement device such as a TIC, a local oscillator, and a network interface. However, unlike a CVDC its receiver is capable of receiving more than one type of CVS, and its control software is slightly more complex. It delivers on the CVDCs promise of allowing the user to select between different sources of UTC. It also provides fail-safe layers of redundancy that a CVDC (or a GPSDC) cannot provide. For example:

- If a UTC source is unavailable, the MSCVDC can switch to another UTC source (Section II.C.1.3).
- If a CVS is unavailable, the MSCVDC can switch to another CVS. For example, it can potentially switch from GNSS to a geostationary satellite or even a terrestrial radio station. This mitigates concerns about GNSS failures (Section II.C.1.4).
- If a network connection is unavailable, the MSCVDC can reference itself to the CVS source, or perhaps obtain clock corrections from a wireless source (Section II.C.1.5).

Figure II.C.3 is a simplified diagram of a MSCVDC distribution system. The right side of the diagram shows that common-view signals are being received at three different locations that maintain UTC – labelled as UTC-A, B, and C. In practice, this could be three time scales; for example, the primary NIST time scale in Boulder (Section I.A) and the secondary NIST time scales in Fort Collins and Gaithersburg (Section I.B). Other possibilities are discussed in Section II.C.1.3.



Lombardi, NIST, U. S. Patent 10,466,363 B2

Fig. II.C.3. Block diagram of a MSCVDC distribution system. Each MSCVDC has multiple layers of redundancy, including access to multiple CVS sources and multiple reference time scales.

The MSCVDC measures the time difference between the CVS and its local oscillator. At each UTC location, a device exists that measures the time difference between the CVS and UTC. The measurements taken at each UTC location are uploaded to a cloud server, shown in the middle of the figure, and then downloaded from the cloud by the clock. By subtracting the UTC measurements from its own measurements, the clock now knows its time difference with respect to each UTC source, so it can potentially use any of them for synchronization. The MSCVDC is configured to select one UTC source as its primary source. If the primary source fails or becomes inaccurate, the clock automatically switches to a secondary UTC source. If N represents the number of available UTC sources, then $N - 1$ is the number of UTC source failures that a MSCVDC can withstand. A detailed discussion of UTC reference options follows in Section II.C.1.3

The middle part of the diagram shows two CVS possibilities, one originating from a satellite and the other from a ground-based transmitter. If the CVS signal is unavailable, the clock can, if the necessary receivers are included in the design, automatically switch to another CVS signal. For example, if a multi-constellation GNSS receiver is included, and GPS signals are unavailable, the clock can switch to signals broadcast by another satellite navigation system such as Galileo. If the appropriate receivers are included, the clock can also switch to a CVS that originates from a geostationary satellite, or from a ground-based radio station. If N represents the number of available CVS sources, then $N - 1$ is the number of CVS source failures that a MSCVDC can withstand. A detailed discussion of CVS options follows in II.C.1.4.

If the network connection fails, the clock can begin using the CVS signal as a reference. For example, if no common-view corrections with respect to a UTC reference are available via the network, the MSCVDC can use any of its CVS sources to discipline its frequency and maintain synchronization until the network comes back. If the CVS originates from GNSS, the time steps that are introduced when the MSCVDC switches from common-view comparisons with a UTC reference to direct comparison with a GNSS source will be small, usually much smaller than $0.1 \mu\text{s}$, so it is normally not a problem for a $\pm 1 \mu\text{s}$ requirement.

Even so, if the time steps introduced by this type of a switch are a problem, they can be calibrated out, if the MSCVDC was keeping track of its time difference with respect to both the CVS and the remote UTC reference prior to the network outage.

Because of its ubiquitous nature, the Internet is typically utilized as the network source. However, because so little information (in terms of bytes) needs to be transferred to the MSCVDC, it would be possible to utilize existing direct broadcast wireless sources for the delivery of the common-view corrections. Again, if N represents the number of available network sources, then $N - 1$ is the number of network source failures that a MSCVDC can withstand. Section II.C.1.5 discusses data transfer options and configurations in more detail.

Table II.C.1 lists the approximate steps that an MSCVDC distribution system performs. Note that the table omits details that are specific to a particular MSCVDC design, such as the averaging interval used to obtain a data point, the frequency of data transfer, the criterion used to establish the validity of a reading, as well as the period of a CVS, reference clock, or network outage that must elapse before the MSCVDC switches to another source. Some details about the considerations taken into account when making these design decisions are provided in Sections II.C.1.3 to II.C.1.5. Before those sections, some details are provided about the current NIST implementation of an MSCVDC.

Table II.C.1. Steps performed by an MSCVDC distribution system.

Step	Action
1	MSCVDC compares its local oscillator to CVS_1 to CVS_N to obtain series of local time differences, LTD_1 to LTD_N
2	Perform sequential search of LTD series until valid reading is found.
3	Identify valid LTD as $MSCVDC - CVS_i$, or LTD_i .
4	Simultaneous with Step 1, each UTC site compares its reference clocks to CVS_1 to CVS_N to obtain a series of reference time differences, RTD_1 to RTD_N
5	Each UTC site, 1 to N , uploads their RTD series to the network.
6	The MSCVDC requests the RTD series from the network for the primary and secondary time scales it has selected. If available, skip to step 7. If no network sources are available, convert LTD_i to a frequency correction and skip to step 9.
7	The MSCVDC finds the value in the RTD series that corresponds to $REF_1 - CVS_i$. If valid, this value is identified as RDT_i . If invalid, the MSCVDC perform a sequential search of the RTD series to find a REF value that matches the CVS used by LTD_i .
8	Convert $RTD_i - LTD_i$ to frequency correction.
9	MSCVDC applies frequency correction to local oscillator using PID controller.
10	MSCVDC applies $RTD_i - LTD_i$ time correction when necessary to improve "lock" time.
11	Keep MSCVDC locked to UTC reference by repeating steps 1 to 10.

Current NIST Implementation of MSCVDC - NIST currently offers a partial implementation of the MSCVDC to subscribers through the NIST disciplined clock (NISTDC) service. The service typically utilizes a rubidium clock (Service ID 76102), but a cesium clock version of the NISTDC (Service ID 76103C) is also offered [47]. Currently deployed at more than 20 locations, NISTDC units synchronize clocks at several stock exchanges, including the New York Stock Exchange and the Nasdaq. A NISTDC also serves as the reference clock for NIST radio station WWVH in Hawaii. The NISTDC maintains synchronization to within about 10 nanoseconds, or 0.01 microseconds, of UTC(NIST), and is packaged as a rack-mount instrument as pictured in Fig. II.C.4.



Fig. II.C.4. NIST’s implementation of the MSCVDC, the NIST disciplined clock (NISTDC).

The NISTDC houses an internal rubidium oscillator (unless an external cesium clock is used), a GNSS receiver, a TIC, a time server synchronized to the NISTDC that supplies time in the network time protocol (NTP) and precision time protocol (PTP) format, an event timing board that helps monitor the accuracy of local time servers by comparing them to the NISTDC, and a distribution amplifier that produces 10 MHz and 1 PPS signals locked to UTC(NIST). Figure II.C.5 provides a block diagram and Fig. II.C.6 is a photograph of the interior of a NISTDC unit.

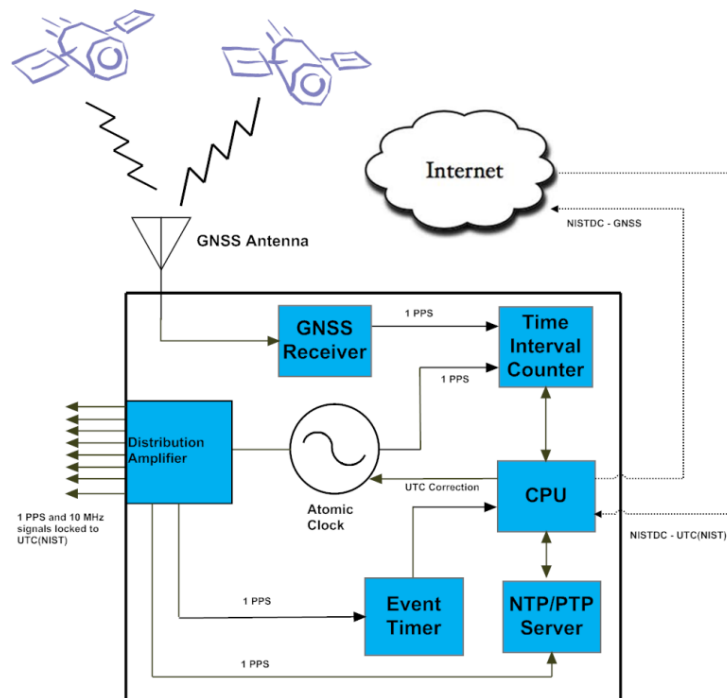


Fig. II.C.5. Block diagram of NIST Disciplined Clock (NISTDC).




Fig. II.C.6. Interior of NISTDC chassis. The visible components include a rubidium oscillator, a time interval counter, a GNSS receiver, an amplifier for distributing clock signals, and other hardware that is utilized to synchronize and measure computer time servers.

II.C.1.3. Reference UTC time scales for MSCVDC

In theory, any UTC time scale, including all of the UTC(*k*) time scales listed on the BIPM Circular-T [6] that provide traceability to the SI second, could serve as the reference clock for an MSCVDC. The only requirement is that the UTC reference makes *REF – CVS* data publicly accessible in a format that the MSCVDC can read, using a CVS that the MSCVDC also supports. Therefore, the MSCVDC can potentially replicate the performance of any time scale.

In practice, however, when our goal is developing resilient time distribution architectures for the United States, the reference clock for a MSCVDC should be a U. S. time scale, such as the primary UTC(NIST) time scale in Boulder (Section I.A) or the secondary time scales located in Fort Collins or Gaithersburg (Section I.B). The other UTC(*k*) time scales located in the U. S. could also be utilized if permissions were obtained; including UTC(USNO); UTC(NRL), located at the Naval Research Laboratory in Washington, DC; and UTC(APL), located at the Applied Physics Laboratory of John Hopkins University in Laurel, Maryland. Three UTC reference time scales that are geographically separated should be enough to provide fail-safe redundancy, but if sufficient data transfer links are in place (Section II.C.1.5), there is no technical limit to how many UTC references can be included in a MSCVDC design.

The MSCVDC software used by the NISTDC currently has switchover capability between the primary time scale in Boulder and the secondary time scale in Fort Collins. If data from either time scale are unavailable, the GNSS signal used as the CVS becomes the reference source. Figure II.C.7 shows the software in operation at a stock market location.

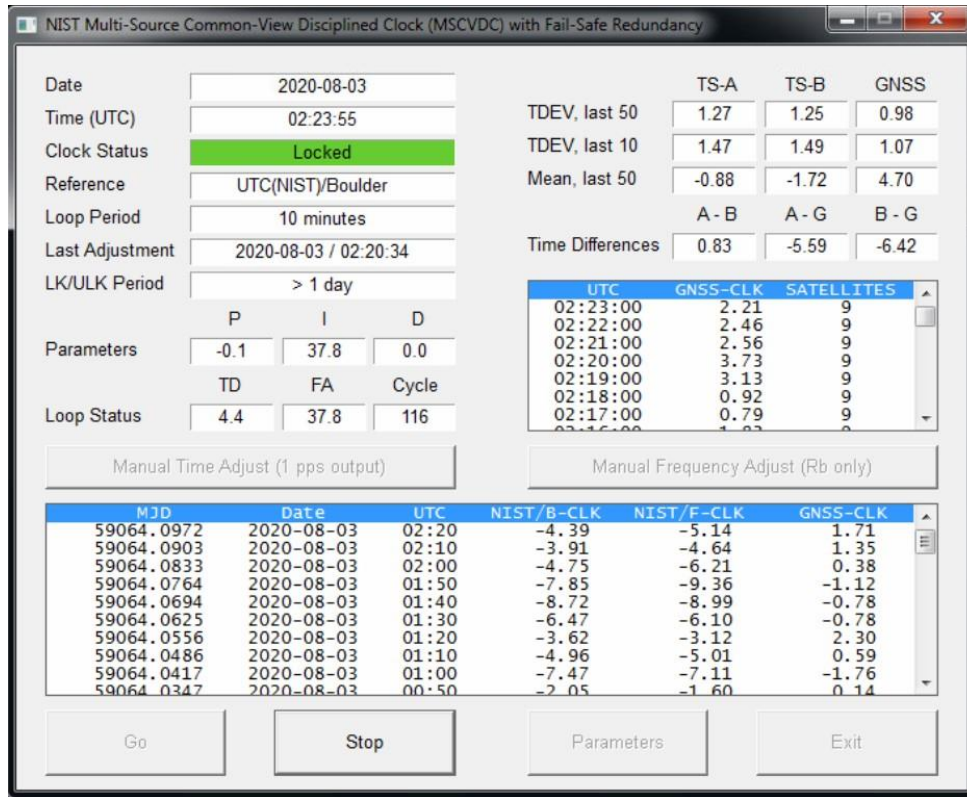


Fig. II.C.7. MSCVDC software that controls the NISTDC.

The large window near the bottom of the display compares the NISTDC, labelled here as CLK, compared to the primary time scale in Boulder (NIST/B), the secondary time scale in Fort Collins (NIST/F), and the GNSS signal used as the CVS. A new data point is obtained every 10 minutes. The boxes in the upper right corner of the display show the average time differences between the three sources (in nanoseconds) during the past eight hours. In this example, the differences are quite small, $A - B$ is the difference between the primary and secondary time scales, or 0.83 ns. The differences between the primary time scale and the GNSS source, or $A - G$, and the secondary time scale and GNSS, or $B - G$, or both near -6 ns. The time steps introduced when the MSCVDC switches between UTC references will not always be this small, but are usually much smaller than $0.05 \mu\text{s}$ (50 ns), and thus are typically not a problem for a $\pm 1 \mu\text{s}$ accuracy requirement. If the time steps introduced by switching reference clocks are considered large enough to be a problem, they can be calibrated out, simply by keeping track of the differences between all of the UTC reference sources during periods when all of them are healthy, and then applying corrections when necessary.

The criterion used for switching from the primary to backup time scale can vary, as they depend upon the hardware used in the MSCVDC design as well as other considerations. When using a rubidium clock, the NISTDC switches to the secondary time scale when the primary time scale has not uploaded data for one hour, as shorter outages will not cause any significant time deviations. Neither the primary or backup time scales will upload data if their time difference from UTC, estimated from cross comparisons at the NIST

sites, exceeds 0.05 μs (50 ns). Thus, if reference clock data are available to the MSCVDC, it can be assumed to be accurate to within that specification.

II.C.1.4. Common-view signal sources for MSCVDC

The common-view technique for comparing geographically separated atomic clocks predates GNSS, as signals originating from very low frequency (VLF) and low frequency (LF) time signal stations [56] and television transmitters [57] were utilized as a CVS source long before satellites. However, the technique became far more commonplace after the launch of the first GPS satellites in the late 1970s [58] and GNSS signals remain the predominant CVS source, capable of time transfer with uncertainties of < 10 ns.

The use of a multi-constellation GNSS receiver is an efficient way to provide a MSCVDC with multiple CVS signal sources. Numerous manufacturers now offer board-level GNSS timing products that receive BeiDou, Galileo, and GLONASS, in addition to GPS, all with a single antenna [59]. The characteristics of these GNSS constellations are summarized in Table II.C.2.

Table II.C.2. Signals received by multi-constellation GNSS timing receivers. The frequencies highlighted in blue are available on single-frequency products, all frequencies are available on multi-frequency products.

	BeiDou	Galileo	GLONASS	GPS
Host Nation	China	European Union	Russia	United States
Altitude (km)	21 150	23 222	19 130	20 180
Orbit Period (hh:mm)	12:53	14:07	11:15	11:58
Orbit Inclination Angle	55°	56°	64.8°	55°
Frequencies (GHz)	1.561098 (B1) 1.589742 (B1-2) 1.20714 (B2) 1.26852 (B3)	1.57542 (E1) 1.191795 (E5) 1.27875 (E6)	~1.602 (L1)* ~1.246 (L2)*	1.57542 (L1) 1.2276 (L2) 1.17645 (L5)
Number of Satellites	35	30	24	31

*GLONASS is a frequency division multiple access (FDMA) system, thus the satellites use multiple frequencies.

Multi-frequency GNSS timing products, for example, GPS receivers with multi-band capabilities (L1, L2, and L5) are still relatively expensive, as they are manufactured in relatively small numbers and marketed primarily to professional users (for geodetic applications, for example), rather than for use in consumer-grade navigation and timing products. However, multi-constellation timing products with just single-frequency capability for each constellation, for example, those that receive GPS L1, BeiDou B1, Galileo E1, and GLONASS L1 (the frequencies highlighted in blue in the table), are now available for less than \$100 in single quantity, making them well-suited for integration into a low-cost MSCVDC design. A single antenna, essentially the same size as a GPS only antenna, usually suffices for multi-constellation receivers (both multi- and single-frequency).

Of course, when GNSS signals are used as a CVS, the ever-present GNSS vulnerabilities that can impact GPSDCs can also impact MSCVDCs. For example, the two most common GNSS vulnerabilities, RF interference and spoofing, can cause a common-view link to either be intermittent or fail if they occur at either the reference clock or the MSCVDC location.

An interfering RF signal can completely block satellite reception, forcing a GNSS receiver to come unlocked. When this happens, the receiver inside an MSCVDC will either turn off its 1 PPS signal, meaning that no common-view data will be available (the preferred practice), or begin reporting incorrect clock

comparisons. The RF signals that interfere with GNSS time signals can have many spectral characteristics and can originate from many sources, making them difficult to categorize [60]. For example, unintentional RF interference can originate from sources including in-band signals (such as signals re-radiated from another GNSS receiver or antenna) and out-of-band signals (radar systems, television stations, microwave data links, and so forth). Intentional interference, known as jamming, occurs when an adversary deliberately blocks reception by broadcasting a signal at or near the GNSS frequency. While jamming constitutes a serious threat, unintentional interference is probably more common. In other words, GNSS signals are often blocked accidentally.

The use of multi-constellation GNSS receivers provides more protection against RF interference than single-constellation receivers, such as a GPS only device. The more costly multi-constellation, multi-frequency receivers, capable of receiving all of the signals listed in Table II.C.2, actually provide a high level of protection against RF interference, as the frequency diversity of the various GNSS signals makes it unlikely that a single interference or jamming source could block all satellite reception [61].

As noted, multi-constellation, single-frequency receivers (those that receive the frequencies highlighted in blue in the Table II.C.2), are better suited for widespread deployment in MSCVDCs because they cost less than multi-frequency receivers. Unfortunately, they offer less frequency diversity – for example, the L1 frequencies for GPS and Galileo are identical, and the BeiDou B1 and GLONASS L1 frequencies are nearby in the spectrum. Even so, they still offer far more protection against RF interference than a single-constellation receiver. A recent study demonstrated this by using a GNSS simulator to jam a receiver’s L1-band reception of both GPS and GLONASS. Narrow and wide-band jamming signals, at various power levels, were alternatively directed at both GPS and GLONASS. By raising power levels high enough, to about -74.8 dBm, it was possible to completely block either GPS or GLONASS reception, but not both. In all scenarios, the combination of GPS/GLONASS allowed at least 40% of the combined number of satellites in-view to remain usable [62].

Spoofing is the practice of tricking a GNSS receiver into reporting the wrong time or position. With few exceptions [63], spoofing is not accidental and originates from an adversary who deliberately broadcasts a false signal that the receiver misinterprets as authentic. If it occurred on either end of a common-view link, it would cause a false clock comparison to be reported. Fortunately, researchers have designed many techniques to detect, and defend against, spoofing attacks [64], and recently anti-spoofing algorithms have begun to take more advantage of the redundancy provided by multi-constellation receivers. For example, a receiver tracking the four constellations listed in Table II.C.2, should normally track about 30 satellites at once. This makes it easier to detect false signals in one or more of the constellations, even if the spoofing attack affects an entire constellation [65].

In addition to the methods cited above to combat jamming and spoofing, an MSCVDC that looks at multiple common-view links using different GNSS constellations can potentially do interference and spoofing detection of its own, and switch to another CVS whenever necessary.

Inclusion of a GNSS CVS in a resilient timing architecture - It can be argued, and often is, that the use of GNSS as a CVS in a MSCVDC distribution system represents only a partial, rather than a full dependency on GNSS, and that a partial dependency is acceptable. The justification for labelling GNSS common-view as only a partial dependency is easy to express; a clock being synchronized via common-view is not being synchronized to GNSS time, but rather to the time of the clock located at the other end of the common-view link. As noted in Section II.C.1, GNSS time does not have to be correct when used as a CVS, because it is cancelled out in the common-view calculation.

Purists will argue that any use of GNSS, even for common-view comparisons, is unacceptable in resilient time distribution architectures. Their viewpoint is also easy to express because the common-view clock

comparison will stop if all GNSS signals are unavailable. However, being a purist is not always easy. For example, it seems obvious that many if not most of the UTC(k) time scales, in the event that they completely lost synchronization, could not be easily synchronized again without relying either on some form of GNSS common-view, or on direct synchronization to GNSS signals. Therefore a third group will argue, perhaps more reasonably than either the purists or those who choose not to worry about a “partial” dependency, that the use of GNSS for common-view is acceptable if a non-GNSS CVS source is also available.

The MSCVDC method can potentially satisfy all three groups. It can work as a system that relies exclusively on GNSS for common-view, either single or multi-constellation, by use of very low cost receiving equipment and a single antenna. Or, if the appropriate receiving equipment and antenna is included, it could fall over to a non-GNSS CVS source whenever necessary. Finally, it could even be designed to rely exclusively on a non-GNSS CVS, to achieve complete GNSS independence. Two non-GNSS CVS sources that could be integrated into a MSCVDC design are geostationary satellites and low frequency (LF) ground-based transmitters, described in the remainder of this section.

Geostationary satellites – As discussed in Section II.B, geostationary satellites are used for TWSTFT clock comparisons, with exceptional nanosecond-level accuracy. As noted, the TWSTFT method requires both stations involved in a comparison to both transmit and receive signals through a transponder on the satellite. This method of time transfer is costly, not only in terms of equipment, but also in terms of leasing satellite time. However, geostationary satellites can also be used in receive-only mode to implement passive common-view comparisons with inexpensive equipment, albeit at lower accuracy levels.

Geostationary satellites are in geosynchronous orbit directly above the Earth’s equator, at an altitude (semi-major axis) of 42 164 km when measured from the center of the Earth. The term geosynchronous simply means that the satellite’s orbital period is the same as the Earth’s rotation period. To ensure that a satellite stays above the same point on Earth, a geostationary orbit must be circular, meaning that its eccentricity and inclination angle must both be near zero. Thus, when viewed from a fixed location on Earth, a geostationary satellite is permanently fixed in the same position in the sky, so that an antenna pointed at the satellite can continuously receive the signal. This differs from GNSS satellites that are only in view for parts of the day, generally during two “fly over” periods. The GNSS satellites circle the Earth approximately twice a day at an altitude that is approximately half that of a geostationary satellite, with inclination angles between 55° and 65° (Table II.C.2).

Because common-view time transfer does not require the CVS to be synchronized to a UTC reference, signals from communication satellites, such as the direct broadcast satellites (DBS) used for television, can potentially serve as a CVS. With the method shown in Fig. II.C.1, two clocks, A and B , are each simultaneously compared to a signal from the DBS satellite. The DBS signal can be, for example, an agreed upon pulse in the stream of sync pulses embedded in the video portion of the television signal [66].

The coverage area, or footprint, of a geostationary satellite is nearly 1/3 of the Earth’s surface. For example, Fig. II.C.8 maps the coverage area of Echostar 105, also known as SES-11, a DBS located on the equator at 105° West longitude, essentially on the same longitudinal line as Boulder, Colorado. As the map shows, this position allows Echostar 105 to provide television coverage to all 50 U. S. states.

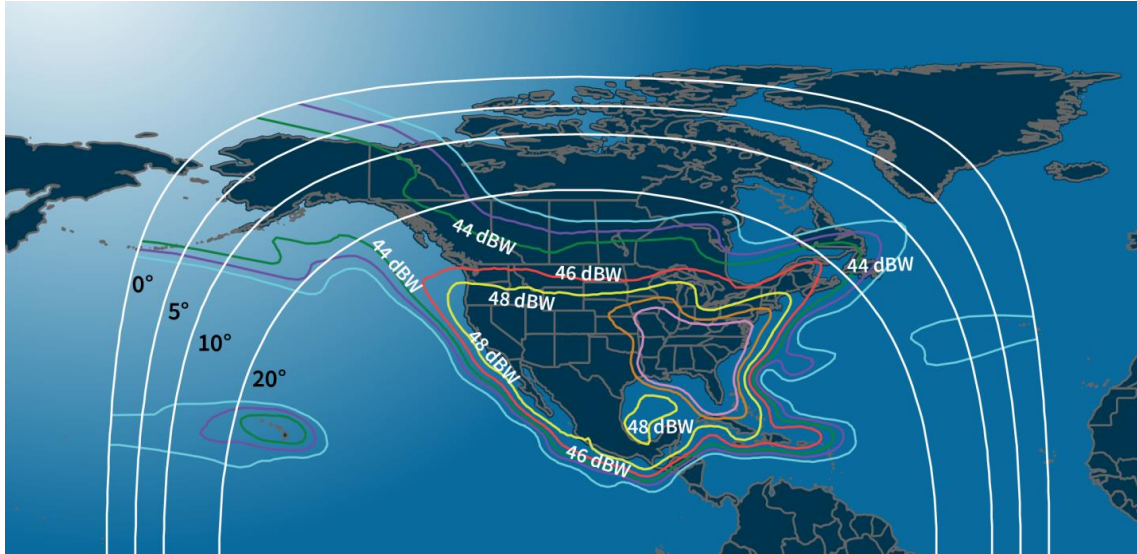


Fig. II.C.8. Coverage area of the direct broadcast satellite Echostar 105, located at 105° West longitude, which provides television signals to all 50 U. S. states (courtesy of Echostar).

When GPS is used as the CVS, the *CLOCK* – *CVS* time differences recorded at each site have already been compensated for the propagation delay between the CVS and the receiver; thus when the two clock measurements are subtracted, only small, nanosecond-level differences between the two signal paths remain. In contrast, when a DBS is used as the CVS the time differences recorded at both sites will be arbitrarily large, with an uncompensated propagation delay difference between the CVS and each receiver. This typically results in millisecond-level differences between the two signal paths.

To explain why the propagation delays can differ so much between two receiving sites, consider that the height of a geostationary orbit above the Earth’s surface is $\sim 35\,786$ km above the equator, so the minimum delay occurs if the receiver is located on the equator and the satellite is directly overhead. Applying the speed of light constant (299 792 458 m/s) produces a minimum propagation delay of 119.4 ms. The maximum delay occurs when the receiver is located at the extreme edge of the coverage area where the satellite appears to be just above the horizon, at an elevation angle near 0° . Then, the distance to the satellite is $\sim 41\,756$ km, and the propagation delay is 139.3 ms, or about 20 ms more than the minimum delay. Because a MSCVDC could be located anywhere in the coverage area, the propagation delay difference is likely to be multiple milliseconds, with 20 ms representing the worst case scenario.

Compensating for propagation delay requires knowing the position of both the MSCVDC and the satellite. If the MSCVDC was designed to use GNSS as its primary CVS and to “fall over” to a geostationary satellite when necessary, then its position should already be known, previously obtained from a GNSS position fix. Thus, we can assume that the MSCVDC position is known. The satellite position is more difficult to know accurately because, as discussed in Section II.B.1, geostationary satellites are not actually stationary but trace a repeatable figure-eight pattern in the sky, known as an analemma. The satellite returns to approximately the same position after a period equal to the length of the sidereal day, but during the day the propagation delay of the signal varies, with a range of $\sim 200\ \mu\text{s}$ (see Fig. II.B.4a). Because the two stations involved in a common-view comparison simultaneously measure the delay from the CVS, essentially taking snapshots of the satellite from different angles as it traces its figure-eight path, the propagation delay differences for a given measurement will be small compared to the $200\ \mu\text{s}$ range, and should be at the sub-microsecond level. However, it will still limit the uncertainty of the comparison as a function of the distance between the two stations. Depending on the length of the baseline between the two stations, the propagation delay difference might typically be tens or hundreds of nanoseconds, approaching

1 μ s for some baselines, when a constant is used for the satellite position, even if the position of both receivers is known [67]. This type of performance can satisfy a sub-microsecond requirement but will be perhaps an order of magnitude noisier than GNSS common-view.

Better compensation for propagation delay and hence better time accuracy can only occur if the actual position of the satellite is known in real-time and made available to the MSCVDC. This would be possible if the satellite broadcast included updated position data that the MSCVDC could download, which would allow calculating the propagation delay on the fly, and removing it from the time difference. If no position data are broadcast, the MSCVDC could estimate the satellite's position itself by using orbital elements downloaded from another source, and then use the calculation results to estimate propagation delay. Satellite providers typically make orbital elements publicly available; however they may only be updated every one or two weeks. The uncertainty of the position calculation will depend upon the age and quality of the orbital elements, with uncertainties increasing as the orbital elements age.

For best results, a precise ephemeris for the geostationary satellite should be generated in real-time and made available to the MSCVDC. A method for doing this was proposed by NIST in the 1980s [68] but not implemented. Since then, other similar methods that utilize ground stations and ranging measurements have been demonstrated [69, 70] that can potentially reduce uncertainties to near the GNSS common-view level. Regardless of what method is utilized to determine the satellite's position, once the propagation delay compensation is complete, subtracting the two *CLOCK* – *CVS* values produces the estimated time difference between the two clocks, Eq. (II.C.1), which the MSCVDC can convert to a frequency correction and apply to the local oscillator.

Common-view measurement programs via DBS geostationary satellites have been successfully implemented with sub-microsecond accuracy by timing laboratories in India [66, 67], the United Kingdom [71], France [72], China [73], and perhaps elsewhere, demonstrating GNSS independence.

NIST formerly operated a one-way timing service from geostationary weather satellites for many years [74], but has not operated a common-view DBS service. However, the use of a DBS as a CVS holds considerable promise for adding additional redundancy to a MSCVDC, as it makes use of already existing resources, and provides frequency diversity with respect to GNSS as the downlink frequencies that broadcast television signals to users are typically in the K_u band (12 to 18 GHz), with some broadcasts at higher frequencies (for example, in the “B Band”, 18.3 to 18.8 GHz, utilized by DIRECTV). An additional directional antenna would be required, but the receiver could likely be made small enough to integrate into the MSCVDC chassis. Table II.C.3 provides just a partial list of DBS satellites that could potentially be used to provide CVS coverage throughout the United States.

Table II.C.3. Partial list of direct broadcast satellites that deliver digital television signals to all 50 states and that could potentially provide a CVS source (sorted by longitude).

Satellite	Year Launched	Longitude	Operator
T11 (DIRECTV 11)	2008	99° W	AT & T
T14 (DIRECTV 14)	2014	99° W	AT & T
T16 (DIRECTV 16)	2019	101° W	AT & T
T10 (DIRECTV 10)	2007	103° W	AT & T
T12 (DIRECTV 12)	2009	103° W	AT & T
T15 (DIRECTV 15)	2015	103° W	AT & T
EchoStar 105 (SES-11)	2017	105° W	EchoStar
EchoStar X	2006	110° W	Dish Network
EchoStar XIV	2010	119° W	Dish Network
EchoStar IX	2003	121° W	EchoStar

LF ground-based transmitters – As noted at the beginning of this section, common-view observations using VLF (below 30 kHz), and LF (30 to 300 kHz) radio signals were utilized long before the launch of the first GPS satellite [56]. In fact, VLF common-view observations were an integral part of the 1950s transatlantic comparisons between clocks at the USNO and the National Physical Laboratory (NPL) in the United Kingdom that led to the current definition of the SI second [75]. Later, 100 kHz signals from Loran-C were used in common-view measurements that collected clock data for the UTC calculations. The first usage of a GPS common-view link to report to UTC was in 1981 [76], consisting of comparisons between NIST (then called the National Bureau of Standards, or NBS) and USNO. Prior to that, NBS contributed to UTC by directly comparing its time scale to Loran-C stations located in Fallon, Nevada and Dana, Indiana, enabling a common-view link to the USNO [77].

A 2003 experiment compared Loran-C common-view to GPS common-view. This was done by using two Loran-C links and a GPS link to simultaneously compare two atomic clocks for a period of two weeks. The clocks were separated by about 445 km, with one located in Gillette, Wyoming, and the other in Boulder, Colorado. The peak-to-peak variation for the GPS link was ~50 ns, with some variation due to the frequency difference between the two clocks. The Loran-C links were noisier by a factor of about two or three. The first link utilized the Loran-C transmitter in Boise City, Oklahoma, located about 860 km from Gillette and 440 km from Boulder, as the CVS. The peak-to-peak variation was ~100 ns. The second link utilized the Loran-C transmitter in Havre, Montana as the CVS, a station further away from Boulder than Boise City (~1050 km) but nearer to Gillette (~600 km). In this case, the peak-to-peak variation was ~150 ns. Even though the Loran-C links were noisier, the average time offset recorded with all three links were within about 10 ns of each other; and the use of Loran-C as a CVS source easily demonstrated its ability to transfer time with sub-microsecond accuracy [78].

In these examples, and in any common-view comparison involving ground based signals, the paths from the CVS to the two receivers will have unequal propagation delays that will need to be removed before sub-microsecond accuracy is possible. However, as previously noted in the section on geostationary satellites, if the MSCVDC was designed to use GNSS as its primary CVS and to “fall over” to a ground-based station when necessary, then its position should already be known, previously obtained from a GNSS position fix. And unlike a geostationary satellite, a ground-based transmitting antenna does not move appreciably, so the MSCVDC can simply use a constant value for the CVS coordinates.

If an LF radio signal was employed as a CVS today, it would add tremendous frequency diversity and resiliency to an MSCVDC whose primary links were based on GNSS or geostationary satellite signals. Unfortunately, all of the Loran-C stations in the United States were turned off in 2010. If a revitalized eLoran system returns to the United States (discussed in Section IV.B), eLoran would become a viable ground-based CVS option for an MSCVDC, if an additional receiver and antenna were added to the device.

The only LF signal currently available for common-view usage in the United States is NIST radio station WWVB (60 kHz). Although WWVB could be utilized as a CVS it has limitations that would make sub-microsecond accuracy difficult, including the cycle ambiguity involved in identifying its on-time marker (OTM), and the fact that it broadcasts from just one location (Fort Collins, Colorado). WWVB is better suited for use as a frequency holdover source, the topic of Section III.A.

II.C.1.5. Data transfer and processing for MSCVDC distribution system

This section discusses four topics – the current data transfer and processing methods employed by a NISTDC, the data transfer and processing methods recommended for widespread MSCVDC deployment, the frequency of data transfer needed to keep an MSCVDC locked, and methods for providing data transfer redundancy.

Data transfer and processing methods employed by a NISTDC - In the NISTDC implementation of an MSCVDC (Section II.C.1.2), data consisting of *NISTDC – CVS* measurements are uploaded to an Internet cloud every 10 minutes using the file transfer protocol (FTP). The server is hosted by Amazon Web Services (AWS) but controlled by and from NIST, and supplemented by a NIST firewall. All NISTDC units have synchronized time-of-day clocks, and thus upload data more or less simultaneously, and all files are stored on the cloud server. To prevent unauthorized access, each NISTDC has its own account, with a unique username and password and an IP address that NIST has previously safe listed.

Immediately after the upload, the NISTDC issues a request to the cloud server to obtain its common-view time difference. The request is made via the hypertext transfer protocol secure (https). The request is made by issuing a command that follows this basic format:

<NISTDC ID><UTCREF ID><#DATA><CV ID>

where

NISTDC ID is a unique code that identifies the NISTDC unit so its data can be retrieved,

UTCREF ID is a unique code that identifies the desired UTC reference time scale,

#DATA identifies the number of data points requested, and

CV ID identifies the CVS source and type of processing (for example, all-in-view or common-view processing can be used with a GNSS CVS source).

The cloud server immediately processes the requested common-view data and sends it back to the NISTDC via https in the form of a time tag followed by a time difference. For example, something similar to this:

59074.1944#5.49

where the first five digits are the Modified Julian Date (MJD), followed by the fractional part of the UTC day, a delimiter (#), and the time difference in nanoseconds.

Multiple requests can be issued by the NISTDC to get processed data for more than one server, or more than one CVS, providing the redundancy discussed in Section II.C.1.2. As was discussed in that section (Table II.C.1), the received time difference is converted to a frequency correction and applied to the local oscillator to keep the device locked to the desired UTC reference.

Data transfer and processing methods recommended for widespread deployment - Note that, in the NISTDC example, the server has copies of all of the data recorded at remote clock sites, and note also that the server performs all of the common-view data processing. This was necessary because NIST has implemented the system as a calibration service to paying customers. As such, NIST must have access to all customer data, to gain knowledge of how each NISTDC is performing. This allows NIST to validate each measurement and send periodic reports to each customer. The reporting is done by providing each

customer with a web portal, updated in real-time, where they can review their frequency and time offset with respect to UTC(NIST). In addition, formal calibration reports are issued to customers every month [47, 48].

An MSCVDC system slated for widespread deployment, for example a commercial product (discussed further in Section II.C.6), would handle its data transfer and processing in a different way, similar to a GPSDC. For example, a GPSDC disciplines its local clock with a signal freely provided by the satellites. Providing the signal is the only responsibility of the GPS system. The signal processing is done entirely by the GPSDC, and the GPS system, of course, has no access to data collected by a given GPSDC, and thus no knowledge of how well a given GPSDC is performing, nor does the GPS system issue calibration reports to its users.

The goal then, for a widely deployed MSCVDC system, is simply to provide “signals” in the form of *UTCREF – CVS_i* data. After receiving this data, each MSCVDC would perform its own data processing, by use of its own *MSCVDC – CVS_i* measurements. This eliminates the step of uploading (pushing) data to a cloud server, but retrieving (pulling) data is still necessary.

Each UTC reference station would provide “signals” by transferring their common-view data, obtained with multiple CVS sources for resiliency and redundancy, to a data repository, which would likely be an Internet cloud server. From there, the MSCVDC would obtain the data with a request that could be as simple as sending a message code, or for security purposes, sending a message code preceded by an authorization code. The number of possible messages would equal the number of UTC references multiplied by the number of CVS sources. If the primary UTC(NIST) time scale in Boulder and the two secondary time scales (Section I.A) were each compared to three CVS sources (GNSS, geostationary satellite, and LF radio, for example), the nine possible messages would be:

- <1>UTCREF₁ – GNSS
- <2>UTCREF₁ – GEO
- <3>UTCREF₁ – LF
- <4>UTCREF₂ – GNSS
- <5>UTCREF₂ – GEO
- <6>UTCREF₂ – LF
- <7>UTCREF₃ – GNSS
- <8>UTCREF₃ – GEO
- <9>UTCREF₃ – LF

The number of messages could be expanded; for example, if the three UTC(NIST) time scales provided data from all four major GNSS systems (GPS, Galileo, GLONASS, and BeiDou) instead of just one, the number of messages would expand to 18 (3 × 6). From there, if two additional time scales were added, there would be 30 messages (5 × 6). The MSCVDC could request as many “signals” as necessary to ensure resiliency and to match its own CVS reception capabilities. Requests could be made in a format such as https (compatible with nearly all firewalls) in the manner of the NISTDC, by using another existing format, or by using a new proprietary format.

Frequency of data transfer – In the NISTDC implementation, new common-view data from the UTC(NIST) time scales is made available every 10 minutes. This is sufficient, because the NISTDCs have internal rubidium oscillators that are more stable than the common-view time transfer link for short intervals, and can thus free run without requiring adjustment for more than 10 minutes. To illustrate this, Fig. II.C.9 shows a “crossover” graph, a tool often used to determine how long the oscillator in a disciplined clock can run without adjustment. The red line shows that the free running rubidium clock reaches a noise

floor of near 5×10^{-13} , as estimated with the Modified Allan deviation [8] at an averaging period of about 1000 s, or ~17 minutes. Due to the effects of frequency drift and aging, the clock's stability with respect to UTC(NIST) then begins to rapidly get worse. The blue line shows the stability of the same rubidium clock when disciplined by a NISTDC. The NISTDC control loop has already taken over before the rubidium noise floor is reached by adjusting every 10 minutes. This initially makes the stability slightly worse (a 20-minute steering interval would be more optimal), but NISTDC stability drops below rubidium stability after about one hour. For periods longer than one hour, the NISTDC is far more stable, reaching a stability of less than 1×10^{-14} after averaging for one day. This stability improves indefinitely, because any differences between UTC(NIST) and a NISTDC are always compensated for by the common-view corrections [48].




Fig. II.C.9. Frequency stability comparison of a free-running rubidium clock to the same rubidium clock disciplined via the MSCVDC method [48]. The common-view disciplining performed by the MSCVDC improves the rubidium clock performance for all averaging periods longer than one hour.

If MSCVDC units are widely deployed, either temperature-controlled or oven-controlled quartz crystal oscillators (TXCOs or OCXOs) are more likely to be included in their design than rubidium devices, and more frequent data transfer will be necessary. This is also the case with many low-cost GPSDC designs using inexpensive local oscillators where the control loop constant is usually much shorter than 10 minutes. A GPSDC design, however, has the advantage of a continuously available reference signal, with updates typically every 1 s (from a 1 PPS signal), or in some cases even faster.

Although not technically impossible, having a MSCVDC receive updates every second is impractical, and even if a system was designed this way, common-view data recorded at 1 s intervals is not especially stable. For this reason, making new common-view data available every minute is recommended. The frequency stability of a common-view link after one-minute of averaging is typically near 1×10^{-11} .

Figure II.C.10 shows stability measurements performed at NIST of three OCXOs and a poor quality TCXO. This graph does not come close to representing all low-cost local oscillator choices, but provides basic guidance. Many low-cost OCXOs, although they typically reach their noise floor in less than 5 s, are still stable to $< 1 \times 10^{-11}$ at one minute, so more frequent adjustments are not necessary. Even a poor quality

TCXO should be stable to $< 1 \times 10^{-9}$ at one minute. Therefore, while a one-minute loop constant will be too long to maintain the tightest possible TCXO lock, it should limit phase excursions to tens of nanoseconds, and maintain sub-microsecond time accuracy. A one-minute interval also provides more flexibility than a 10-minute interval for MSCVDCs that include more stable OCXOs or rubidium oscillators, as it allows averaging multiple 1-minute segments and choosing an optimal loop constant.




Fig. II.C.10. Frequency stability comparison of free-running quartz oscillators.

Methods for providing data transfer redundancy – As previously noted, common-view measurements from UTC reference sites are sent to a data repository, which will likely be an Internet cloud server. For redundancy, the same data could be made available from multiple cloud servers, allowing an MSCVDC to switch servers if one is unavailable, or even cross check data obtained from multiple servers. The servers themselves could be distributed amongst multiple cloud providers (for example AWS, Google Cloud, IBM, Microsoft Azure, and so forth).

Of course, if all local network access is blocked, an MSCVDC will likely need to utilize a CVS as its reference (Section II.C.1.2). This could be remedied with a “network-free” MSCVDC if common-view data were made available from a wireless source, which is feasible because only a small number of bytes need to be transferred. Radio signals covering the CONUS could modulate the messages described earlier on their carrier frequency, and even if the data transfer rate was slow, the messages could be broadcast via a defined schedule over a period of minutes or hours. Radio station WWVB (Section III.A), or a revitalized eLoran system (Section IV.B), would be two potential candidates for providing this service.

II.C.2 – Reliability and Redundancy of MSCVDC Distribution System

The inherent layers of redundancy in a well-designed MSCVDC include multiple UTC reference sources (Section II.C.1.3), multiple common-view comparison links (Section II.C.1.4), and multiple data transfer mechanisms (Section II.C.1.5), and should provide unmatched reliability. Widespread and sustained power outages represent perhaps the biggest threat to MSCVDC operation, but that is an unavoidable concern with every time realization and distribution system.

A reliability question that is certain to be asked when considering a MSCVDC distribution system involves complexity – an MSCVDC involves a number of “moving” parts compared to other disciplined clock systems. For example, a GPSDC, now that GPS receiver technology has matured and been commoditized, is relatively simple. As shown in Fig. II.C.11, it consists of a GPS receiver and antenna, an oscillator, a phase comparator that compares GPS signals to the oscillator, and a microcontroller that reads the results of the phase comparison, converts them to frequency and/or phase corrections, and sends them to the oscillator, either digitally or as a varying control voltage. This forms a continuously running control loop that locks the oscillator to GPS.

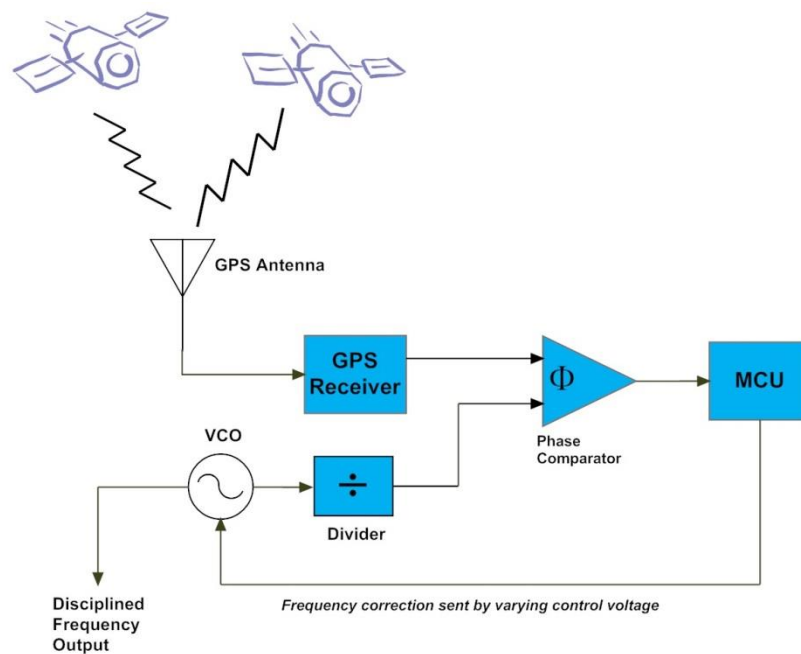


Fig. II.C.11. Block diagram of a GPS disciplined clock (GPSDC).

The electronics and software required to build a rudimentary GPSDC are often shared, and many do-it-yourself (DIY) designs constructed by hobbyists can be found on the Internet. In contrast, a MSCVDC involves systems at multiple locations that share data with each other (Fig. II.C.3), more complex software, as well as cooperation with the institutions that maintain UTC time scales.

When considering the impact that the additional complexity might have on reliability, it is important to note that a MSCVDC does not introduce any new technologies, nor does it require building any new infrastructure. It simply ties together, primarily with software, several mature technologies and resources that already exist. For example, the time scales that it references, such as UTC(NIST), are already continuously operating and providing time signals, regardless of whether or not those signals are being used to discipline clocks. The freely-available CVS signals are already being continuously broadcast for other

purposes, which may or may not be related to time distribution. The networks and network services used to transfer data, such as the Internet and associated cloud services, already exist for a multitude of other purposes.

While the control software of a MSCVDC is more complex than that of existing time distribution systems, it is still trivial when compared to other software systems, such as the e-commerce, news, weather, and messaging apps that we all routinely access on our smartphones. All of those systems have far more “moving parts” and are much more difficult to design and understand than a MSCVDC, and still maintain very high levels of reliability with only infrequent outages. Just as those systems offered a more modern approach to shopping, communications, and information retrieval systems, with benefits that far outweighed their complexity, the MSCVDC potentially offers a more modern approach to resilient time distribution systems than has previously existed.

There would, of course, still be a “learning curve” associated with the development and testing of a MSCVDC software platform. However, it seems likely that such a software platform could rapidly mature, and that any reliability concerns based on complexity could rapidly go away.

II.C.3 – Physical and Cyber Security of MSCVDC Distribution System

As noted in Section I.A.3.2., raw or processed measurement data products resulting from time measurements are not considered to be confidential or proprietary information. This type of data, while not always published in real-time, is routinely published by NIST and many other laboratories. While a secure, encrypted data channel, such as https, should be used for data transfer to avoid tampering, it is probably not necessary to require an MSCVDC requesting common-view data from an Internet server to provide an access or authorization code. The number of bytes transferred per request is small, comparable to the data transferred when accessing a public NTP server (Section V.A), and a standard protocol could be developed for common-view data that allowed a variety of MSCVDC designs to freely access the data and obtain synchronization to UTC(NIST), in much the same fashion that a variety of GPSDC designs freely access GPS signals and synchronize to UTC(USNO).

If cloud servers are utilized as common-view data repositories, two potential concerns are denial of service attacks caused by overloading a server with false requests from non-MSCVDC sources, or hacking into a server and modifying or corrupting the common-view data. Standard security practices routinely used by web servers, NTP servers, and other data providers can guard against these events, as can having an MSCVDC cross checking data obtained from multiple servers (Section II.C.1.5).

The largest security concern probably involves the physical and cyber security of the UTC time scales reference stations, since the time kept by a MSCVDC originates from these sources. The topics of physical and cyber security were addressed for the UTC(NIST) primary time scales in Section I.A.3, and in Section I.B.3 for the secondary time scales.

II.C.4 – Validation of MSCVDC as a Trusted Time Reference

In the NISTDC implementation of an MSCVDC (Section II.C.1.2), every instrument, with its associated antenna and antenna cable, is calibrated at NIST in Boulder, Colorado before it is shipped to the customer. The purpose of the calibration is to measure the differential delay between the customer’s MSCVDC; known as the system under test, and a reference system. The calibration is performed via the common-clock method, where the system under test and the reference system are both measuring the same clock, in this case a 1 PPS signal from UTC(NIST). The cables that connect both the reference system and the system under test to UTC(NIST) are carefully calibrated so that their differential delay is as close to zero as possible, and antennas for both systems are mounted on poles that have coordinates known to within 20 cm (Fig. II.C.12).

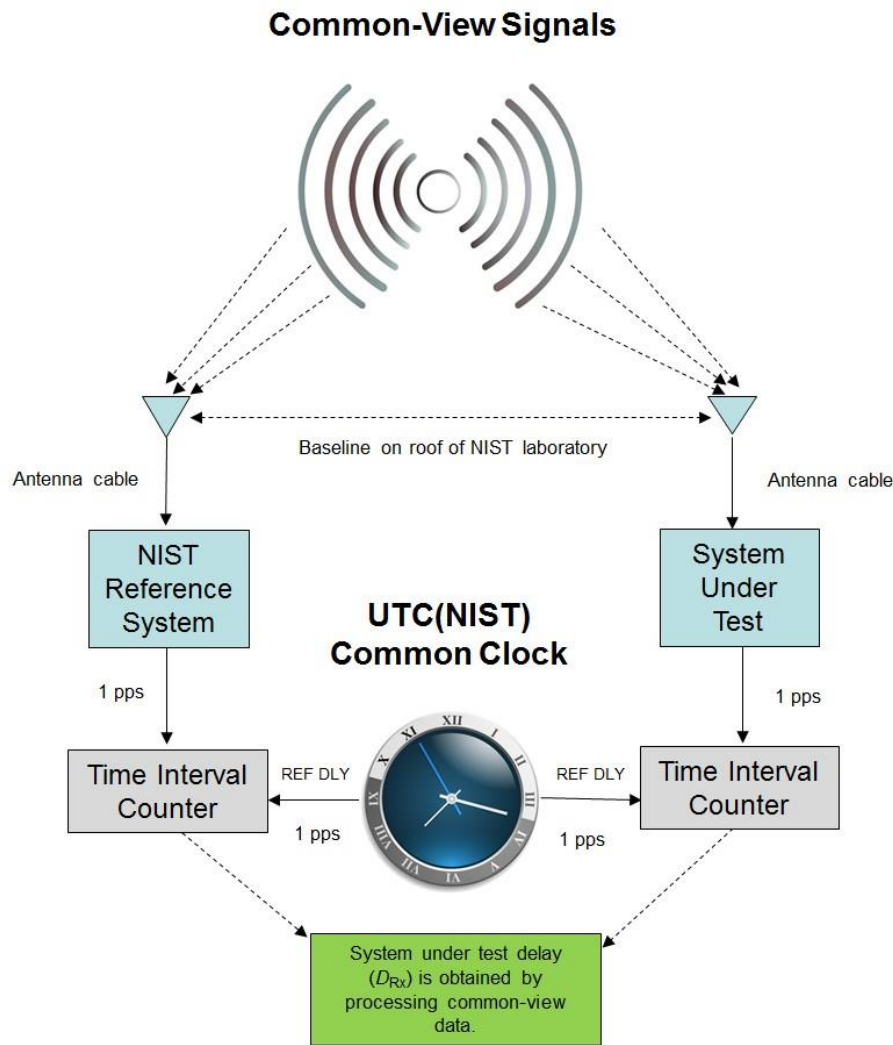


Fig. II.C.12. Configuration for a common-clock MSCVDC calibration at NIST.

The calibrations last for at least one week. When completed, the average time difference of the MSCVDC with respect to the reference system, D_{Rx} , is entered into the MSCVDC software to compensate for its internal delays [48]. If multiple CVS sources are used, common-clock calibrations should be performed for each CVS, with a separate value of D_{Rx} used for each common-view link.

When the common-clock calibration is finished and a delay value has been entered, the 1 PPS connection from UTC(NIST) is removed and the newly constructed NISTDC locks to UTC(NIST) via common-view comparisons. Then, a 1 PPS output from the NISTDC is directly compared to the 1 PPS output of UTC(NIST) with a TIC, allowing the results of the common-view comparison to be compared to the results of the direct comparison. This is done to ensure that the output of the new NISTDC is within 10 ns of UTC(NIST), which is the uncertainty of the calibration service. Typically, the difference is less than 5 ns. If the unit is outside of its stated uncertainty, it is not shipped to the customer until the uncertainty is brought within specifications.

To illustrate this, Fig. II.C.13 shows the results of a 60-day comparison between a NISTDC and UTC(NIST). The blue line shows the results of the direct comparison and the red line shows the results of the common-view comparison. For the 60-day interval, the average value of NISTDC – UTC(NIST) was -0.1 ns for the common-view comparison, and 3.4 ns for the direct comparison, a difference of 3.5 ns. This difference is primarily due to biases introduced by the uncertainty of the common-clock calibration, but other factors, including environmental changes since the time of the calibration, might have also contributed to the difference. The common-view comparison will always produce a time difference near 0 because the NISTDC control loop continuously adjusts its clock towards a set point of 0, providing compensation that can hide biases in the calibration. The direct comparison reveals these biases by measuring the true time difference and validates that the NISTDC is within its stated uncertainty [48].




Fig. II.C.13. Direct and common-view comparisons of a NISTDC.

This process provides the ultimate validation of a MSCVDC as a trusted time reference, because it actually verifies that the claimed measurement uncertainty is true. It is analogous to bringing a calibrated GPSDC inside the USNO laboratories and comparing it directly to UTC(USNO). However, for widespread deployment of a MSCVDC, it would be unrealistic to individually calibrate every unit, just as it is unrealistic to calibrate every GPSDC.

When GPSDCs are deployed, the premise is that sub-microsecond accuracy is expected without calibration, and this same premise could be applied to a MSCVDC to provide a realistic validation solution. This could be done, for example, by developing a standard configuration for a MSCVDC, one where each unit is constructed with identical hardware and software, and then calibrating a few of the units. All units could then be shipped with pre-inserted D_{Rx} values for each CVS source, based on delays for standard antenna cable lengths. As is the case with GPSDC products, the MSCVDC software would allow the D_{Rx} values to be changed when necessary, allowing users to enter new values to reflect, for example, if longer antenna cables were added.

The uncertainty assigned of the device could be expanded to conservatively cover all situations, again following the GPSDC model. For example, a 100 ns uncertainty (a typical GPSDC specification), should conservatively cover all of the GNSS CVS sources, even when the delay differences between units differ by tens of nanoseconds. Larger uncertainties may need to be assigned to non-GNSS CVS sources, but enough margin remains between 0.1 μ s (100 ns) and 1 μ s to validate sub-microsecond accuracy requirements for all configurations.

II.C.5 – Accuracy and Stability of MSCVDC Distribution System with respect to UTC(NIST)

A locked MSCVDC, referenced to UTC(NIST) and using GNSS as the CVS, will stay accurate to within a few nanoseconds of UTC(NIST) and report an average time offset near 0. To illustrate this, Fig. II.C.14 shows a 6-month (July to December 2019) comparison of a MSCVDC located at a major US stock exchange to UTC(NIST). The data points shown in the figure are one-hour averages. The peak-to-peak variation over the 6-month interval is ~25 ns, but most data points fall within a ± 5 ns range and the average time offset is less than 0.1 ns, or essentially 0.




Fig. II.C.14. Accuracy of MSCVDC at a major U. S. stock exchange, with respect to UTC(NIST).

As discussed in Section II.C.4, and revealed by the direct comparison data shown in Fig. II.C.13, the time difference will not actually be 0, but will have a small amount of bias introduced by the uncertainty of the MSCVDC calibration and other factors, such as temperature changes that can impact the local oscillator. These small biases are unknown to the MSCVDC control loop, which simply uses the common-view data to continuously adjust the clock towards a set point of 0. For these reasons, the time differences reported by the MSCVDC have an uncertainty of about 10 ns ($\pm 0.01 \mu\text{s}$) [48].

The accuracy of a MSCVDC will of course be affected by time steps that occur during time scale or CVS switchovers, none of which occurred in the data set shown in Fig. II.C.14. However, if these switchovers had occurred in that data set, they may not have been obvious, as their magnitude is typically a small number of nanoseconds. For example, when a MSCVDC switches to or from a primary to a secondary UTC reference, it will incur a time step equivalent to the difference between the two time scales. Thus, if the current time difference between the primary UTC(NIST) time scale in Boulder and the secondary time scale in Fort Collins is 10 ns, and the MSCVDC is forced to switch to the Fort Collins scale because data from Boulder is unavailable, the MSCVDC will incur a 10 ns time step. Switching between CVS sources will also introduce time steps, but if each CVS channel has been calibrated as described in Section II.C.4, those time steps should be small, at the level of a few nanoseconds.

The interval selected before a MSCVDC switches from one time scale to another or from one CVS source to another will depend on the stability of the its oscillator. For example, if a temporary network outage

occurs, and the MSCVDC cannot receive corrections from Boulder for 30 minutes, it may not be necessary to switch to Fort Collins if the local oscillator is a rubidium. This is the case with the current NISTDC implementation of the MSCVDC, where no switchovers occur in less than 30 minutes. However, if the local oscillator is a TCXO, the switchover will need to occur more quickly.

Figure II.C.15 shows the time deviation (stability) of the data shown in Fig. II.C.14 for averaging periods ranging from one hour to about one month. After averaging for one hour, the stability is about 1.5 ns, dropping below 0.4 ns after one day and below 0.2 ns after one week. This high level of stability is possible because the time differences between UTC(NIST) and the MSCVDC were always compensated for by the common-view corrections.




Fig. II.C.15. Stability of MSCVDC at a major U. S. stock exchange, with respect to UTC(NIST).

II.C.6 – Feasibility of a Commercially-Available MSCVDC Product

The introduction of a commercially-available MSCVDC product is entirely feasible. The hardware would be similar to existing GPSDC designs with multiple receivers/antennas added for the reception of multiple CVS sources, and network connectivity added for the retrieval of common-view data from UTC reference stations. If the device was designed to have just multi-constellation GNSS ability, just one antenna would be necessary. The cost of manufacturing the unit would likely be most heavily influenced by the quality and cost of the internal oscillator.

The device software, as discussed in Section II.C.2, would be slightly more complex than GPSDC software, but not prohibitively so, and certainly much less complex than the software found in many other commercial instruments and application platforms.

The manufacturer of a commercial MSCVDC product would need to work with NIST and/or other government laboratories that maintain UTC time scales, so that reference stations could be installed at those laboratories. The arrangement with the UTC reference laboratory would involve a signed agreement, most likely involving a legally-binding contract such as a cooperative research and development agreement (CRADA) or a memorandum of understanding (MOU). The reference stations would essentially be just another MSCVDC unit, modified to make its data available through a network.

The MSCVDC manufacturer would also need to provide one or more data repositories for common-view data collection, which would likely be hosted by an Internet cloud server provider. Once this basic infrastructure was provided, a MSCVDC time distribution system could support a nearly limitless number of users.

III. Methods of Distributing UTC(NIST) as a Holdover Frequency for GPS Disciplined Clocks

When GPS signals are unavailable, GPS disciplined clocks (GPSDCs) enter holdover mode, where their accuracy in the absence of GPS is entirely dependent on the quality of their local oscillator, and if present, the quality of their holdover software. In many cases, the holdover capability of a GPSDC is poor, and unable to maintain $\pm 1 \mu\text{s}$ for more than a few seconds or minutes. Even if a GPSDC has a rubidium local oscillator, maintaining $\pm 1 \mu\text{s}$ for more than 72 hours is unlikely.

This section describes two methods of distributing a redundant frequency source, disciplined to agree with UTC(NIST), that would potentially allow GPSDCs to maintain the sub-microsecond accuracy established by GPS throughout signal outages that last for multiple days. Clocks that implement this methodology would still have just one way to achieve time synchronization accurate to $1 \mu\text{s}$ (GPS), but would now have the benefit a backup disciplined oscillator that remains locked when the GPS disciplined oscillator is unlocked, and thus can maintain time synchronization indefinitely, unlike a free running local oscillator. The first method described in this section utilizes 60 kHz radio signals from NIST radio station WWVB. The second described method utilizes the precision time protocol (PTP). In addition to their potential integration into GPSDC products, both methods can potentially complement the MSCVDC method previously described in Section II.C.

III.A Distribution of UTC(NIST) Frequency to GPSDCs via NIST Radio Station WWVB

NIST radio station WWVB [79, 80, 81] continuously broadcasts standard time and frequency information via a 60 kHz signal. These signals reside in the portion of the radio spectrum known as low frequency (LF), which is much lower than the frequencies of the commercial AM radio broadcast band, which begins at 530 kHz. The signals have a wavelength of 5000 m, which allows them to pass through walls and buildings and be received indoors, often reaching basements or rooms without any windows. This gives WWVB an advantage over line-of-sight signals, such as those from the GPS satellites

The station's radiated output power is about 70 kW, placing it about midway between the legal limits for AM (50 kW) and FM (100 kW) transmissions. Signal attenuation is low in the LF portion of the radio spectrum – for example, a signal at 1 MHz, near the middle of the AM broadcast band, is attenuated by at least 20 dB more than a 60 kHz signal when both signals traverse over a 1000 km path. Thus, a power output of 70 kW is large enough to allow WWVB to reach the entire continental United States (CONUS) during both the daytime and nighttime hours. Reception is also possible in Hawaii, Central America, and most of the populated regions of Canada.

Due to their long wavelength, LF radio signals can diffract over obstacles and travel beyond the horizon, following the surface of the Earth. This type of propagation is known as ground wave, and has the advantage of relatively stable propagation delays, when compared to terrestrial based signals that are reflected from the ionosphere (sky wave). Stable propagation delays, even when the magnitude of the delay is unknown, are useful for providing an accurate and stable frequency. Thus WWVB, when used to discipline an oscillator, can potentially provide holdover frequency that is accurate enough for a GPSDC to continue to maintain sub-microsecond timing accuracy for multiple hours or days if GPS is unavailable.

Section III.A.1 provides a technical description of the WWVB architecture, including the station's transmission and antenna system, its broadcast format, and its coverage area. Section III.A.2 covers the station's reliability, its levels of redundancy, and the safeguards in place to ensure continuous operation. Section III.A.3 describes the physical and cyber security of the station. Section III.A.4 describes the frequency and time stability of the broadcast signals with respect to UTC(NIST), and the type of performance that could be expected from a WWVB disciplined oscillator (VBDO). Finally, Section III.A.5 discusses the feasibility of a commercially-available VBDO product for integration into GPSDCs.

III.A.1 – Technical Description of WWVB Architecture

NIST has continuously operated radio station WWVB since July 1963, broadcasting standard frequency and time signals from a 380 acre ($1.5 \times 10^6 \text{ m}^2$) site owned by the United States Department of Commerce (DOC). The site is shared with WWV, a NIST shortwave radio station that also broadcasts standard time and frequency signals on the high frequency (HF) bands known as shortwave. The station site is located about 10 km north of the city of Fort Collins, Colorado, and about 80 km north of the NIST laboratories in Boulder.

A block diagram of the WWVB station architecture is shown in Fig. III.A.1. The diagram shows the secondary UTC(NIST) time scale, previously described in Section I.B, that provides the reference for all of the frequency and time signals distributed by the station. It also shows the other elements of the station, which will be presented in the following sections (in this order) – the antenna array which consists of a north and south antenna, the three transmitters, the two helix houses, the broadcast control elements (including the broadcast control console and the RF switch matrix), and the time code generators (TCGs), which produce the distributed time codes by modulating the 60 kHz carrier frequency.

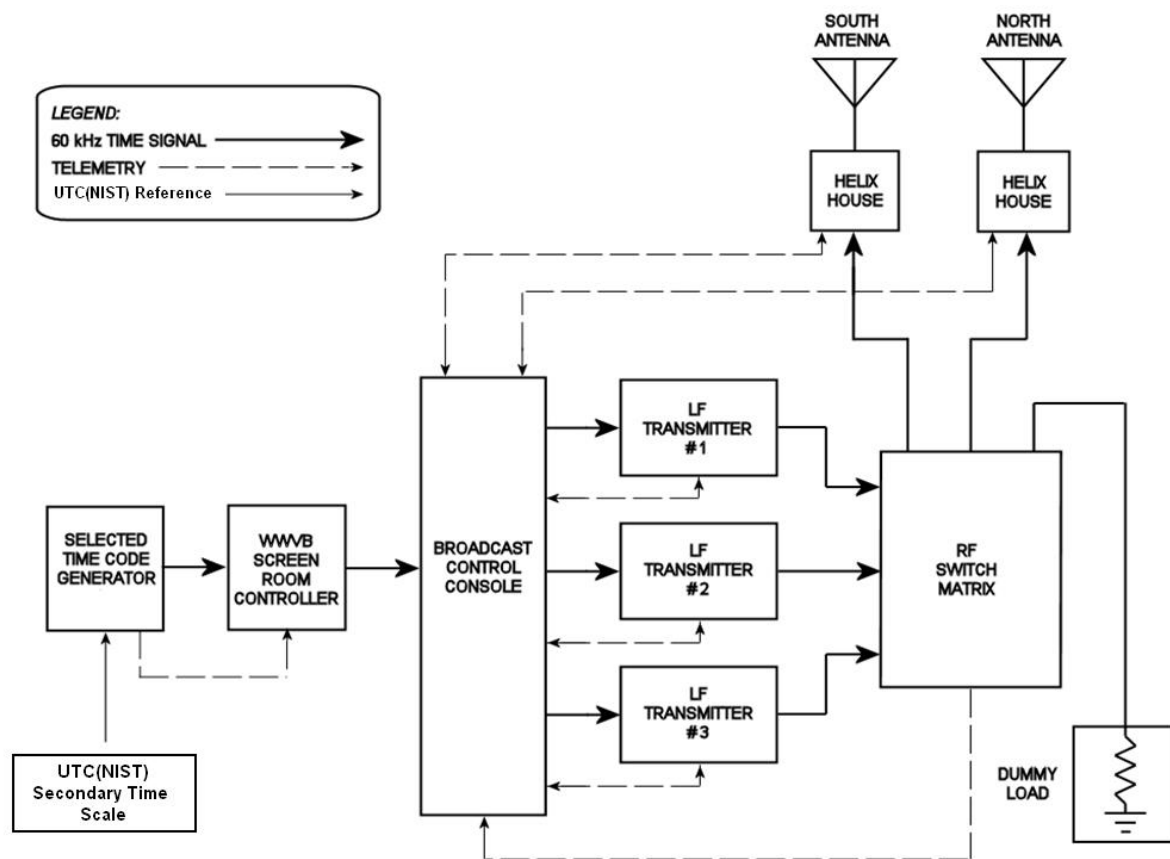


Fig. III.A.1. Block diagram of WWVB station architecture.

III.A.1.1 - WWVB antenna array

The most significant feature of WWVB is its large antenna array, which covers a ground area of about 30 acres (121 406 m²). The area consists of two transmitting antennas, situated to the northwest and southeast of the WWVB transmitter building, and known as the north and south antennas (Fig. III.A.2).

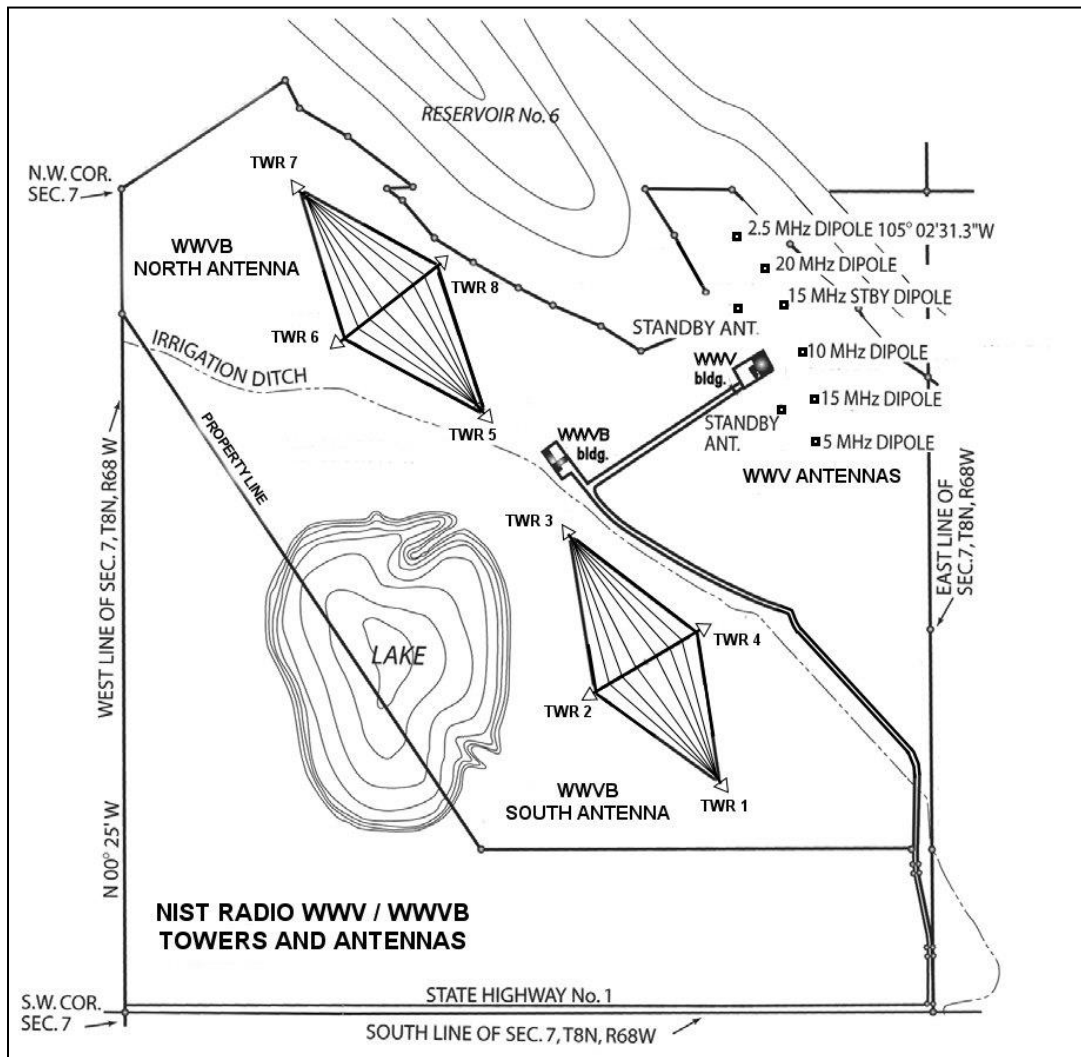


Fig. III.A.2. Site map of WWV/WWVB.

Each antenna system consists of four towers, a capacitance hat (or top hat) and a down lead, and a ground plane. The four towers in each antenna system are arranged in a diamond pattern. Each tower is held in place by three sets of guy cables, with three cables in each set. Seven of the eight towers are 122 m (400 ft) high; the exception is tower 2 in the south antenna which is 126.5 m (415 ft) high, due to its base being at a lower ground elevation than the bases of the other towers. The towers are electrically grounded at their bases and insulators are installed on the guy cables near the guy anchor.

The towers are not electrically “hot” as they are not used to radiate signals. Instead, they structurally support the parts of the antenna that do radiate signals, specifically the top hat and the down lead. The down lead of each antenna descends into a helix house, described later in this section; that contains equipment to

connect the antennas to the transmitters. The north antenna coordinates are $40^{\circ} 40' 51.308''$ N, $105^{\circ} 03' 00.011''$ W; the south antenna coordinates are $40^{\circ} 40' 28.298''$ N, $105^{\circ} 02' 39.531''$ W. The maximum distance between towers in either the north or south antenna is 573 m (1880 ft), and the distance between the north and south antenna down leads is 866 m (2840 ft). Figure III.A.3 is a diagram of the antenna array, and Fig. III.A.4 is an aerial photograph.




Fig. III.A.3. Diagram of the WWVB antenna array.



Fig. III.A.4. Aerial photograph of the WWVB antenna array.

Because ground waves should be vertically polarized to minimize attenuation by ground currents, meaning that the electrical field is vertical, while the magnetic field is horizontal, the two WWVB antennas are vertical monopoles. Even though the antenna array is physically large, it is electrically small for the LF signals being transmitted. Ideally, the length of the radiating element of an efficient transmitting antenna should be at least one-quarter wavelength, which would be 1250 m long in the case of 60 kHz, and impractical to build. The WWVB antennas partially compensate for this by placing the vertical element of the antenna as high as possible, and adding electrical “length” horizontally at the top of the vertical radiating element. This technique is known as top loading, and the top hat serves as the loading element.

The top hat consists of aluminum cables arranged horizontally in a diamond shape, with a tower at each corner of the diamond. The top hat is not fastened directly to the towers but instead is connected to porcelain insulators, which in turn are connected to steel cables that extend to 5000 lb (2268 kg) concrete counterweights near the tower’s base. This keeps the top hat electrically isolated from the towers, and allows it to “float” between the towers, with tension from the counterweights keeping it nearly parallel to the ground. Figure III.A.5 shows the antenna connections at the top of one tower.




Fig. III.A.5. Connection of the WWVB antenna to the top of a tower.

The down leads on each antenna consist of six aluminum cables that extend from the center of the top hat to a point about 93 ft (28.4 m) above the ground, where they connect through an insulator to a tensioning cable that keeps them nearly vertical. The tensioning cables are connected to counterweight structures located on the ground, equivalent in size to those on the towers. Opposite the tensioning cables, the down lead continues at an angle to a pass-through insulator on top of the helix houses. The electrical connections to the transmitter are made through the insulator.

The ground plane is necessary to minimize ground losses from the transmission, thus improving the antenna array’s efficiency. The ground plane consists of 300 wires per antenna, extending radially about 1300 ft (396 m) from the helix houses. The uninsulated wires are made of silver cladded copper braid, and about 3 mm in diameter. At the foundations of the helix houses, the ground wires are brazed to a buried 13 mm copper cable. This cable is soldered to a network of copper straps that extend out of the concrete foundation of the helix houses at 2 m intervals. These straps connect to stubs that protrude from the floor inside the

helix house, where the connection to the antenna is made. The efficiency of the antenna is also aided by the natural high soil conductivity at the Fort Collins site, measured at 14.5 millimhos per meter to a depth of 0.45 m [80].

The spectrum allocation for WWVB allows for ± 1 kHz of bandwidth (59 to 61 kHz). The time code information transmitted by WWVB requires very little bandwidth to transmit, typically less than ± 10 Hz. Thus, the antenna bandwidth is wider than necessary to transmit the time code, but still much narrower than the antennas of shortwave or commercial broadcast stations that transmit audio. The intrinsic bandwidth of the antennas was measured as 310 Hz for the north and 263 Hz for the south [82]. The combined bandwidth is estimated to be about 411 Hz. When loaded by the transmitter, this increases by a factor of about 1.5, resulting in an estimated bandwidth of slightly more than 600 Hz.

III.A.1.2 - WWVB transmitters

WWVB has three vacuum tube transmitters, designated as low frequency transmitters LFT-1, LFT-2, and LFT-3. They were originally manufactured by Continental Electronics for the United States Navy in the mid-1960s and known by their military model number as the AN/FRT-72. They were obtained by NIST as military surplus in the mid-1990s. Two of the transmitters are continuously “on the air” and the third serves as a standby.

The transmitters are large units consisting of six metal cabinets arranged side-by-side that house different elements of the equipment. Each transmitter is approximately $24 \times 3 \times 6.5$ ft ($7.3 \times 0.9 \times 2$ m). They are linear amplifiers, with four power amplifier (PA) tubes configured for Class AB₁ operation. A block diagram of a single transmitter is shown in Fig. III.A.6. As shown in the diagram, they are wired for 480 V input. When operating at an output level of ~ 50 kW with a 100% duty cycle, they require about 140 kVA of power, distributed as 70 kVA to each of two plate power supplies (one for each pair of PA tubes).

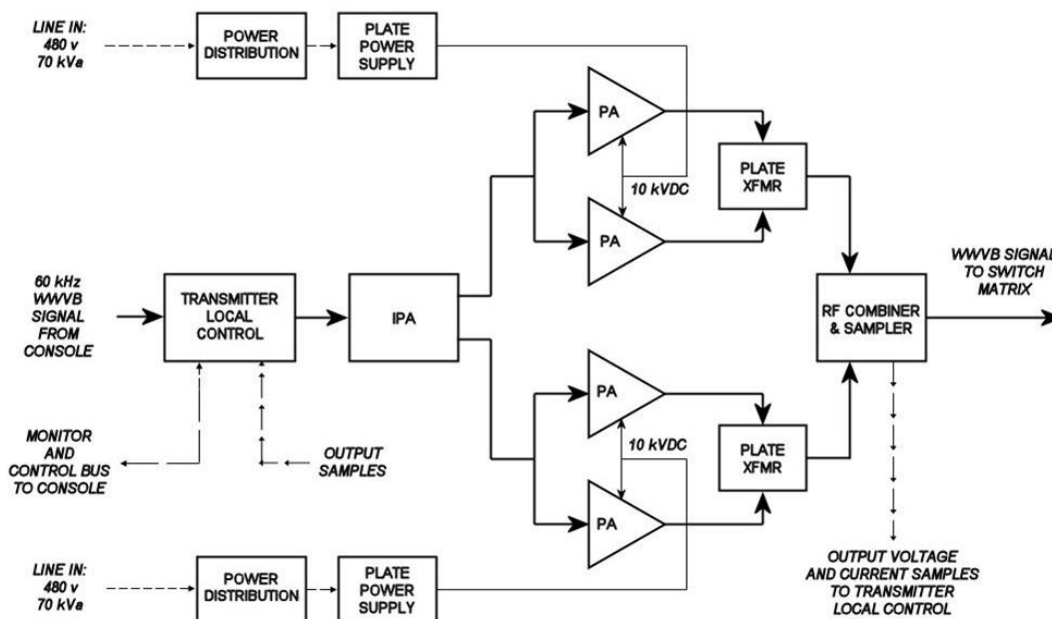


Fig. III.A.6. Block diagram of a single WWVB transmitter.

The input WWVB signal, with a nominal value of 1 V RMS, is fed from the control console through a local control unit inside the transmitter. The local control unit sends the low level WWVB signal to the solid

state intermediate power amplifier (IPA), which generates a 200 W drive signal, and sends it on to the four PA tubes.

The four PA tubes operate as two pairs. Each pair is configured as a push-pull amplifier, and the outputs of the two amplifiers are combined in a plate transformer. The outputs of the two plate transformers are then combined with a RF combiner/sampler unit. This unit connects to the RF switch matrix using rigid transmission line. Output voltage and current samples are sent to the transmitter local control unit [83].

III.A.1.3 - Helix houses

The two helix houses are metal buildings that reside underneath the north and south antennas. They are identical structures containing identical equipment. Each helix house is approximately 7.3 m (24 ft) tall, with a floor space of 6.1 m \times 9.4 m (57.3 m² or 620 ft²). The helix houses were named for historical reasons that no longer apply. When WWVB operated at lower power levels, they housed a large helical coil as part of the original antenna matching equipment. However, when the station began using higher power transmitters in the mid-1990s, the helical coils were replaced with a variometer and other coupling components that were capable of handling the higher power output. Figure III.A.7 is a photograph of one of the helix houses, and Fig. III.A.8 shows the rooftop connection to the antenna array.



Fig. III.A.7. Exterior view of helix house.



Fig. III.A.8. Roof of helix house and connection to antenna array.

The helix house equipment (Fig. III.A.9) matches the 50Ω impedance of the transmitters to the much lower impedance of the antennas. This is necessary because, as previously noted, the radiating elements of the antennas are extremely short when compared to their wavelength. A quarter-wavelength antenna for 60 kHz would be 1250 m tall, and the effective or electrical height of the WWVB antennas is about 86 m. If the radiating element were large enough to match the wavelength, its impedance could be made to be very near 50Ω , matching the transmitter. However, the impedance of the electrically small WWVB antenna is slightly less than 1Ω when operated in single mode, and just slightly more than 1Ω when operated in dual mode [82, 83, 84]. A large inductor cancels the capacitance in the antenna, and by doing so matches its input impedance to the output of the transmitter. A variometer, or variable inductor, tunes out any reactance in the antenna to maintain the impedance match.




Fig. III.A.9. Interior view of helix house.

A simplified diagram of the helix house equipment is shown in Fig. III.A.10. A 3 1/8 in (79 mm) flexible coaxial transmission line runs from the transmitters through an RF switch matrix (described later in this section) to each helix house. Other cables, including fiber optic, copper, and line voltage cables, run alongside the transmission lines, protected by a concrete-lined trench. The length of the transmission line is 417 m to the north helix house, and 423 m to the south helix house.

The transmission line is terminated in a fenced area inside the helix house, where it connects to a “tee” network. Coupling inductors are installed in parallel with the output of the tee; one coil is switched with an RF relay that changes the system impedance for single or dual antenna operation. Unless the station has had a malfunction or is undergoing maintenance, the station operates with dual antennas, so this relay is usually open. The antenna current is sampled by a current transformer, and sent to the control console in the transmitter building via fiber optic cables. Another coil, mounted next to the variometer, supplies phase and signal status information back to the transmitter building.

The variometer, pictured in Fig. III.A.9, is a large inductor that consists of a fixed horizontal winding called a stator. A rotor, with a smaller winding, can vary the inductance by rotating inside the stator. The rotor is automatically controlled by an auto tuning system that reacts to changes in antenna impedance. When the antenna impedance changes due to environmental conditions such as wind or ice, the auto tuning system activates an electric motor underneath the stator. This variometer drive motor turns the rotor in either direction as necessary to maintain an optimal transmitter/antenna impedance match [80].

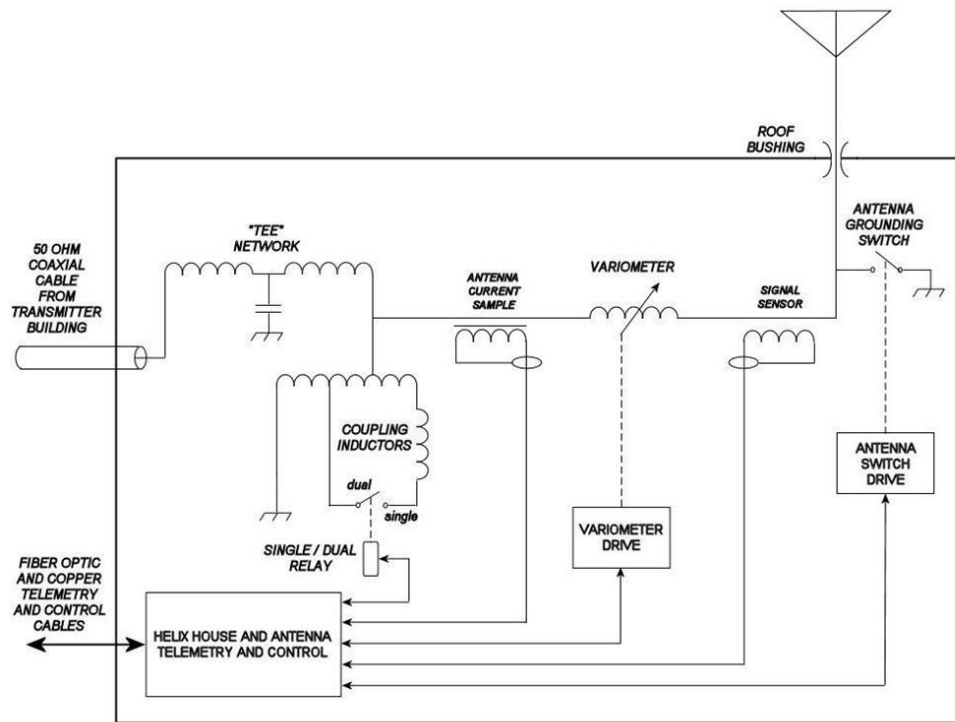


Fig. III.A.10. Simplified diagram of WWVB helix house.

Hazardous voltages are present inside the helix house, which is the reason for the fenced area inside the buildings. For safety and security reasons, the gate in the fence is always locked, and the key to the lock is stored in a secure area. Before opening the gate, an antenna grounding switch is activated. In addition, when the gate is opened, an interlock switch prevents the connected transmitter from operating. Some controls, such as access to the variometer drive motor and antenna ground switch, are located outside the fenced area and can be accessed while the transmitter is turned on [80].

III.A.1.4 - Broadcast control elements and station efficiency

The broadcasts are controlled from a control console and RF switch matrix (Fig. III.A.11) located near the transmitters in the WWVB building.

The control console uses a programmable logic controller (PLC) that constantly monitors the operating status of the broadcast, and prevents operating modes that would be unsafe to personnel or damaging to the equipment. For example, it does not allow more than one transmitter to operate into a single load, or allow a transmitter to operate without a load.

Unless maintenance or repairs are being performed on one of the antennas or helix houses, WWVB operates in dual antenna mode with two transmitters on the air that simultaneously broadcast the same signal at equal current levels. One transmitter feeds the north antenna, and the other feeds the south antenna. The efficiency of the WWVB broadcast, or the percentage of transmitter output power converted to radiated RF output power, is only about 55% on average when broadcasting from a single antenna (approximately 56% from the north antenna and 54% from the south antenna). However, this increases to nearly 69% in dual antenna mode. When the transmitters both output their typical level of about 51 kW, this results in a radiated output power of about 70 kW, as noted earlier.

The RF switch matrix (Fig. III.A.11) allows any transmitter to be connected to either antenna or to a dummy load. The switch is operated manually. It includes three sets of RF busses on either side of a grounded vertical baseplate. Transmitter and load connections are made with 3 1/8 in (79 mm) coaxial transmission line. The three transmitters are connected to the rear busses, and the three loads are connected to the front set of busses; with 50 Ω impedance maintained throughout the switch. Interlock switches are installed to track the switch configurations, and this information is sent to the control console.

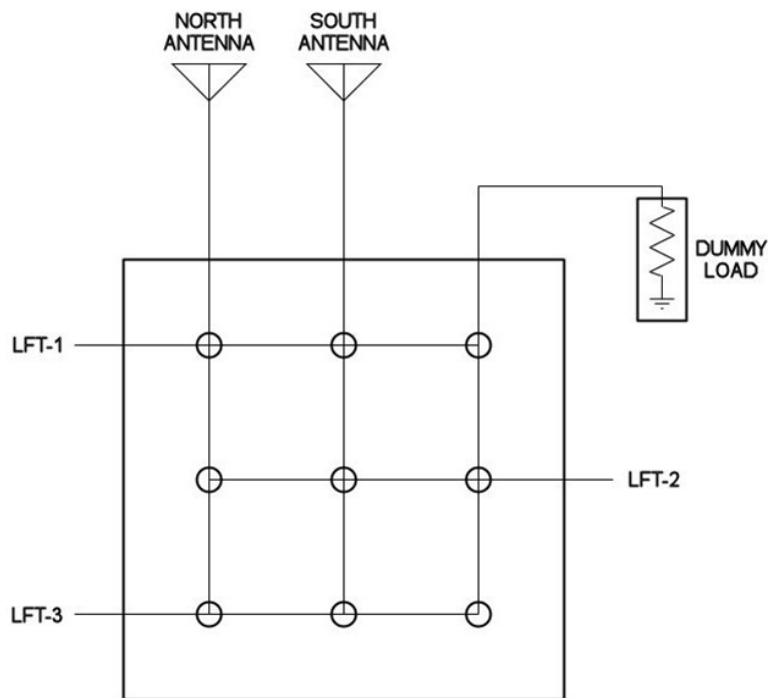


Fig. III.A.11. WWVB RF Switch Matrix.

As previously mentioned, the antennas are auto tuned with a variometer in the helix house to compensate for changes in their impedance. When the antennas move, due to environmental factors such as wind or ice loads, the impedance they present to the transmitters changes with their movement. This causes the reflected (reverse) power to increase, and the radiated (forward) power to decrease. A local control unit inside each transmitter (Fig. III.A.6) includes a phase comparator that monitors the phase difference of current and voltage samples taken from the output of the transmitter. When the antenna is properly tuned to the transmitter, the current and voltage are essentially in phase. When the antenna moves, it causes a change in the phase relationship between the voltage and current, and this change is detected by phase comparator. The transmitter control unit then issues a command to the control console, which in turn issues a command to the helix house to adjust the variometer and correct for the impedance mismatch [80].

III.A.1.5 - Time code generators and digital time code formats (legacy and phase-modulated)

The time code generators (TCGs) control all of the information transmitted by WWVB. The TCGs are rackmount instruments (Fig. III.A.12) located in the WWVB transmitter building. There are three TCGs, located in adjacent equipment racks; that are continuously inter-compared as described in Section III.A.2. A fourth TCG is installed as a spare, but not compared.



Fig. III.A.12. WWVB Time Code Generator.

Each TCG is enclosed in a 4U rack-mount chassis and is built around the Altera Stratix 3 FPGA (Field Programmable Gate Array) evaluation board which utilizes a factory integrated phase lock loop voltage-controlled oscillator (VCO) circuit that locks to the 5 MHz signal that originates from the secondary UTC(NIST) time scale (Section I.B). Interface circuits that surround the FPGA include an analog front end that accepts the 5 MHz reference, and an output driver that modulates a digital time code in two formats on to the 60 kHz signal (to be discussed in the next paragraph), three user configurable test ports used to monitor various timing attributes, and a touch screen display which provides the user interface. The user interface allows synchronizing the time codes with 1 ns of resolution, and entering values that are included in the time code, such as daylight saving time and leap second indicators, and other bits that are reserved for special purposes.

The two time code formats do not interfere with each other. The first format, known as the legacy time code, as it has been broadcast by the station since 1965. The second format is known as the phase modulated (PM) time code and was added to the broadcasts in October 2012.

The legacy time code [80] sends information in a format similar to binary coded decimal (BCD) format, where four binary digits (bits) are required to send one decimal digit. The bits are sent by temporarily dropping the power of the 60 kHz carrier, and then restoring the power to normal. The power drop is synchronized to UTC(NIST) and occurs at the beginning of each second, and the instant when the power drop occurs is known as the on-time marker (OTM).

The data rate of the legacy time code is slow, just 1 bit per second is transmitted, and a full minute is required to complete a time code. The duration that the power is held low determines whether the transmitted bit is a binary 0 or a binary 1. A binary 0 is sent when the power is held low for 200 ms and full power is restored for the next 800 ms. A binary 1 is sent when the power is held low 500 ms and full power is restored for the remaining 500 ms. Time code frame markers are also sent every 10 s, by holding the power low for 800 ms and restoring full power for the remaining 200 ms. The legacy time code format is shown in Fig. III.A.13.




Fig. III.A.13. The legacy WWVB time code format.

This simple form of amplitude modulation (AM), sometimes called amplitude shift keying, requires the receiver to distinguish between high and low power to properly decode each bit. One drawback of the method is that if signal conditions change during the period (200, 500, or 800 ms) when the carrier power is held low, or at any time during the 60 seconds needed to read the entire time code, then high power may be indistinguishable from low power and the time code will be unreadable. This becomes even more of a problem if the orientation of the receiving antenna is changing, in which case the received signal strength can change dramatically. To make the amplitude shift easier to detect, a large carrier power drop of 17 dB is utilized (it was 10 dB prior to 2006) [85].

The PM time code was added the WWVB broadcast in 2012 to make the time code more reliable and easier to receive, by allowing receivers experiencing different reception conditions to process the time code differently. For example, it allows for the accumulation of information received over multiple one-minute frames, which has the same effect as increasing the signal gain. This allows the PM code to be read more reliably than the legacy code when signal conditions are poor, such as during the daytime hours, or during

periods of high atmospheric noise. However, the normal PM time code has a 60 second frame like the legacy code, and thus still requires 60 seconds to decode. An even slower code, with a six-minute frame is also broadcast every 30 minutes, for use when signal conditions are very poor. The TCGs are capable of also generating a fast code that transmits the time code in just 630 ms once every 10 seconds, but the fast code has not yet been deployed [86].

The PM code was designed by Xtendwave, a Dallas, Texas company that also designed and manufactured the TCGs (Fig. III.A.12). The code utilizes a modulation scheme called binary phase shift keying (BPSK) that differentiates between a 0 and a 1 bit by shifting the phase of the 60 kHz carrier. A 0 bit is transmitted when there is no phase shift, and a 1 bit is transmitted by inverting the phase of the carrier by 180°. The phase transition between each bit and the next one in the 1 bps (bit per second) PM frame occurs 100 ms after the carrier amplitude drop, that serves as an OTM and indicates the start of a second. The PM time code includes a 13-bit synchronization word, a time and date word, and daylight saving time (DST) and leap second information. The time and date word, from which the year (excluding century), the date, and the hour and minutes may be extracted, is represented by a 26-bit minute counter that represents the number of minutes that have elapsed since this century began at 00:00 Coordinated Universal Time on January 1, 2000. The 26-bit time and date word is converted to a 31-bit Hamming code by the addition of five error correcting bits. This allows the time code receiver to correct one error and to detect up to two errors in the time word. Figure III.A.14 shows the frame structure of the normal PM time code, and how it compares to the legacy AM time code [86].

Bit index	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
AM	M	Minute		R	Minute		M	R	R	Hour	R	Hour		M	R	R	Day	R	Day		M									
PM	Sync word													Time and Date													R			
Bit index	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60
AM	Day			R	R	UT1		M	UT1		R	Year		M	Year		R	LY	LS	DST	M									
PM	Time and Date									R	Time and Date						D+L	R	D+L	DST Next			0							

Fig. III.A.14. Frame structure of the PM time code compared to the legacy AM time code. *M* denotes a frame marker, *R* denotes reserved, *D + L* denotes daylight saving time plus leap second information, *LY* denotes leap year information, and *LS* denotes leap second information.

III.A.2 – Reliability, Redundancy, and Safeguards Ensuring Continuous Operation of WWVB

Because WWVB is the synchronization source for many millions of consumer-oriented radio controlled clock products (Section III.A.5), reliability is of the utmost concern. The desired reliability metric, as stated in the NIST Time and Frequency division quality system, is for WWVB to be on the air at least 99.7% of the time, which allows about 1578 minutes of downtime during the average length of a year, or just slightly more than one day. All periods when the station is off the air are counted, including periods of scheduled maintenance that occur during normal working hours when the station is staffed, and unanticipated station failures that occur out of hours that require personnel to respond by returning to the station as quickly as possible. The 99.7% reliability metric was met during eight of the past 10 years (2010 to 2019), having been narrowly missed in 2013 and 2014, as shown in red in Table III.A.1. The percentage of time that WWVB was on the air in 2019 was 99.92%, the highest figure during the past ten years. The number of outages per year has ranged from 11 to 27, with the average outage length being less than one hour in all years except for 2009 and 2013.

Table III.A.1. Percentage of WWVB time on air (2010 to 2019).

Year	Time off Air (minutes)	Time on Air (%)	Number of Outages	Average Outage Length (minutes)
2019	414	99.92	12	34.5
2018	526	99.90	14	37.6
2017	811	99.85	14	57.9
2016	862	99.84	19	45.4
2015	1008	99.81	18	56.0
2014	1593	99.70	27	59.0
2013	1877	99.64	21	89.4
2012	979	99.81	17	57.6
2011	755	99.86	16	47.2
2010	654	99.88	12	54.5
2009	664	99.87	11	60.4

III.A.2.1 - Alarm systems and staff response time

The WWVB broadcast is continuously monitored by several parallel systems that alert the station staff of several types of failures. Staff members are on call 24/7, and immediately respond when an alarm message is received. The staff is well versed in troubleshooting and quickly repairing station malfunctions and much of the outage time consists of “drive time” to the station when failures occur out of normal working hours. There are two main types of WWVB alarms:

- 1) Time Code Generator (TCG) Alarms** – The three TCGs reside in an RF shielded room, known as the screen room, where a comparator/control system continuously monitors them for accuracy using majority vote logic. The time code output of each TCG, including both the legacy AM code and the PM code are continuously compared to the other two TCG units. If any TCG is found to be in error by more than 1.4 μ s, or if its output signal goes away, it is immediately taken off line. If the defective TCG was on the air, the controller connects another TCG to the transmitters. The controller then immediately sends an alarm to a commercial security company, which also monitors the station fire and burglar alarms. The company will notify station staff by telephone calls. The screen room controller also

sends an alarm signal to an on-site paging system which sends an alarm notification by text message to the cell phones of all staff members.

2) Transmitter Output Level Alarms - The output signal level of the transmitters is monitored by the screen room controller, and also by a computer running data acquisition software. If the WWVB output signal drops, the controller triggers and sends an alarm to the commercial security company and paging system mentioned previously. The data acquisition software, which monitors and logs the antenna current for each WWVB transmitter, sends text message notifications to staff members when the current drops below acceptable levels.

Note that the transmitters contain overload circuits that will shut off the 10 kV DC plate power supply (Fig. III.A.6) or the RF input, when equipment malfunctions or high reflected power levels occur. These overload events could be caused by several factors, such as high wind gusts, ice loading, lightning, or a defective tube or other component. Often the events are short in duration and the transmitter can be safely restored to normal operation in a few seconds. This is accomplished by an automatic overload recycling system, which resets the overload relay and restores the plate power supply. The system is set at three recycles; if after the third attempt an overload condition still exists, the plate power supply is disabled. In the case of high wind gusts or ice loading, the transmitter local control unit contains a Carrier Cutoff (CCO) circuit which monitors the reflected power levels and disconnects the transmitters for a few seconds if the reflected power levels are too high. After a preset number of CCO events within a preset time period, the drive level is locked out. When this happens, a computer running data acquisition and control software at the control console retunes the variometers in the helix house to a known-good position using the variometer control circuits (Fig. III.A.10), resets the CCO circuits in the transmitters and restores the RF input at half power to the transmitters for about 20 seconds. This allows the automatic tuning to regain control of the antenna/transmitter match. Full power is then restored to the transmitters and the broadcast resumes as normal.

If this effort fails station staff, notified by the alarms, can log into the computer remotely and manually retune the antenna/transmitter match. If for some reason one of the transmitters cannot be brought back online, for example due to overloads, staff can remotely change the transmitter configuration to single antenna operation to restore the broadcast as soon as possible with one transmitter. The staff then responds to the problems with the offline transmitter.

III.A.2.2 - Power failures and backup generator

Station outages due to power failures are extremely rare. The standby generator is an Onan model DFED-3370866. It is standby rated at 625 kVA, 500 kW at sea level, and has a 480 V three-phase output. The engine is a six-cylinder supercharged diesel rated at 755 horsepower at sea level, fed from a 1000 gallon (3785 L) above ground fuel storage tank which was installed in 1998. The generator was installed in 2000 [80] and is capable of carrying all station electrical loads, including the transmitters. When no power is provided by the utility company, the generator's diesel storage tank has enough fuel to keep WWVB on the air for approximately two weeks.

The generator is connected to an automatic transfer switch (ATS). When the commercial power is interrupted, the ATS starts the generator and transfers the building loads automatically within one minute. The transmitters are configured to restart automatically when this happens. After commercial power is

restored, the generators continue to run for 30 minutes to allow it to stabilize then transfer back to commercial power. This involves another brief power interruption. A sufficient number of UPS units are installed so that the station time scale, TCGs, and other distribution equipment are not affected by a commercial power outage.

III.A.2.3 - Age of WWVB transmitters

The age of the WWVB transmitters is an additional concern. So far, the transmitters have been exceptionally reliable. However, their vacuum tube design is energy inefficient, generates large amounts of heat and noise (from the necessary cooling devices), and requires regular maintenance. In addition, because few radio stations still utilize vacuum tube transmitters, replacement parts have become increasingly difficult to procure. For all of these reasons, NIST is currently exploring several funding avenues that would allow replacing the vacuum tube transmitters with solid state transmitters. The current transmitters consume about 1.8 million kWh per year, where 1 kWh equals power consumption of 1 kW sustained for one hour. We estimate that the installation of solid state transmitters would decrease power consumption by at least 33%, in addition to increasing the station reliability.

III.A.3 – Physical and Cyber Security of WWVB

Because so many Americans rely on the time signals of WWVB, and signals from the shortwave station WWV, which shares the same site, the physical and cyber security of the site is paramount. Only members of the station staff are authorized to enter the site without first receiving approval from the engineer-in-charge, who also serves as the site security manager. No tours of the station are given to the general public.

III.A.3.1 - Physical security

As previously noted, the station site is large, covering 380 acres ($1.5 \times 10^6 \text{ m}^2$), and is only accessible through a motorized gate. This gate is located about 1.5 km north of the highway leading to the station, where the fence line of the property crosses the entrance road, and the entrance road is unmarked. A manually operated gate is also located near the highway. This gate is open during normal operating hours, but chained and locked at all other times. A 2.4 m tall chain link fence, topped with three strands of barbed wire, has been installed on the perimeter of the property. A 21 acre ($84\,984 \text{ m}^2$) portion of the south end of the property is outside of the fence. The unfenced portion of the property is inaccessible from any road.

The gate is controllable from inside the WWVB building, where video surveillance equipment allows unrestricted views of the gate area. The gate is also controlled by a keypad mounted on a pedestal outside the fence line. Entrance codes are only available to station personnel who can reach the keypad from their vehicles, and to the electric power company. A telephone is also mounted near the keypad to allow authorized visitors to ring station personnel for entry. The main buildings are wired with a fire alarm and security system monitored by a local security company. Figure III.A.15 shows the WWVB gate, and Fig. III.A.16 shows a portion of the perimeter fence.



Fig. III.A.15. Motorized entrance gate at WWVB site.



Fig. III.A.16. Portion of perimeter fence at WWVB facility.

III.A.3.2 - Cyber security

All of the networked computer systems at WWVB are managed using the same security policies practiced on the NIST campus in Boulder, with operating systems and updates applied whenever necessary.

The computer systems that control the elements of the WWVB broadcast, including the TCGs and transmitters, cannot be accessed from any publicly accessible network. Therefore, it is not possible for an outsider to “hack” into the site and change any of the station parameters or to take the station off-line. The station, with the exception of a few password protected functions that are accessible remotely via point-to-point communications, can only be controlled by authorized personnel who are physically present at the facility. Thus, the most important aspect of the site’s cyber security is the physical security described above.

III.A.4 – Frequency and Time Stability of WWVB with respect to UTC(NIST)

The reference clock for WWVB is a secondary UTC(NIST) time scale (Section I.B) that is tightly synchronized, typically to within ± 25 ns of UTC. Despite the excellent accuracy at the source of its signal, WWVB unfortunately has limitations as a time distribution system that will likely prevent it from ever distributing time that meets $1 \mu\text{s}$ accuracy requirements. One huge limitation is that, unlike GPS, a navigation and positioning system that can determine the location of its clocks, WWVB clocks don't know where they are or where the station is, and thus cannot automatically perform the calculations necessary to compensate for the path delay between the transmitter and receiver. The WWVB signals are delayed as a function of distance from the transmitter, and these delays can be as large as 15 ms within the CONUS. Constructing a WWVB clock that compensated for path delay would first require the clock to obtain accurate coordinates for its location, something that could only be done if a positioning service, such as GPS, were used.

Assuming that a WWVB clock does have known coordinates, the next step would be to calibrate all other delays, including the transmission delay of the station, the delays associated with the receiving equipment, and to keep track and compensate in real time for the variations in propagation delay caused by skywave/groundwave interaction. Compensating for variations in propagation delay could potentially be done by constructing a network of monitoring stations that augment WWVB by measuring and accounting for differential delays along the transmission path, which would be an expensive and complex undertaking. Even then, the period of the 60 kHz carrier is $16.67 \mu\text{s}$, so obtaining an OTM accurate to within $1 \mu\text{s}$ requires much finer resolution than the LF carrier frequency provides. The best obtainable accuracy would likely involve identifying the correct carrier cycle, and then resolving the OTM to within one half cycle, or $8.3 \mu\text{s}$. In practice, even identifying the correct cycle is difficult, and WWVB has rarely been used for reliable time synchronization better than $100 \mu\text{s}$ (six carrier cycles).

A much more promising application of WWVB in critical infrastructure systems involves utilizing its stable LF groundwave signals to discipline an oscillator for use as a frequency holdover source for GPS clocks. Unfortunately, as will be discussed in Section III.A.5, no commercial products currently exist that utilize WWVB for this purpose. However, prior to the introduction of the PM time code, measurements conducted for many years by NIST in Boulder (~ 78 km from WWVB) indicated that the Allan deviation of a commercial VBDO was a few parts in 10^{12} after one day of averaging [80]. This level of frequency stability suggests that a clock already synchronized and locked to GPS could “fall back” to the WWVB frequency if GPS reception were lost, and potentially maintain synchronization indefinitely to within a few hundred nanoseconds (on average). More recently, measurements made at the University of Ohio with a PM code prototype receiver indicated holdover of better than $\pm 1 \mu\text{s}$ for short periods in the absence of GPS using a WWVB disciplined rubidium clock [87]. The technique of using an LF radio signal to provide frequency holdover to GPS clocks is not new, and was previously implemented with 100 kHz Loran-C signals (Section IV.B). Prior to the termination of Loran-C in 2010, Motorola had successfully deployed thousands of GPS clocks with Loran-C frequency holdover capability in telecommunication networks [88].

To better illustrate how this technique could work, Fig. III.A.17 shows phase data (normalized to 0) from two different VBDOs as received in Boulder during January and February, 2007. Each data point was obtained by dividing a 5 or 10 MHz VBDO output to 1 PPS, comparing the output to UTC(NIST) every second with a TIC, and averaging readings for one hour. Ferrite core loop antennas, about 25 cm in length and oriented broadside in the direction of the station, were used to receive the WWVB signal. The local oscillator in both devices was an oven controlled quartz crystal oscillator (OCXO) similar to those commonly found in GPSDCs. Their frequency offsets were measured to be about 4 to 5 parts in 10^8 , meaning that they would accumulate a $1 \mu\text{s}$ time error every 20 s to 25 s if allowed to free run. When disciplined by WWVB, they stayed within an approximate $\pm 1 \mu\text{s}$ range for the two-month interval.

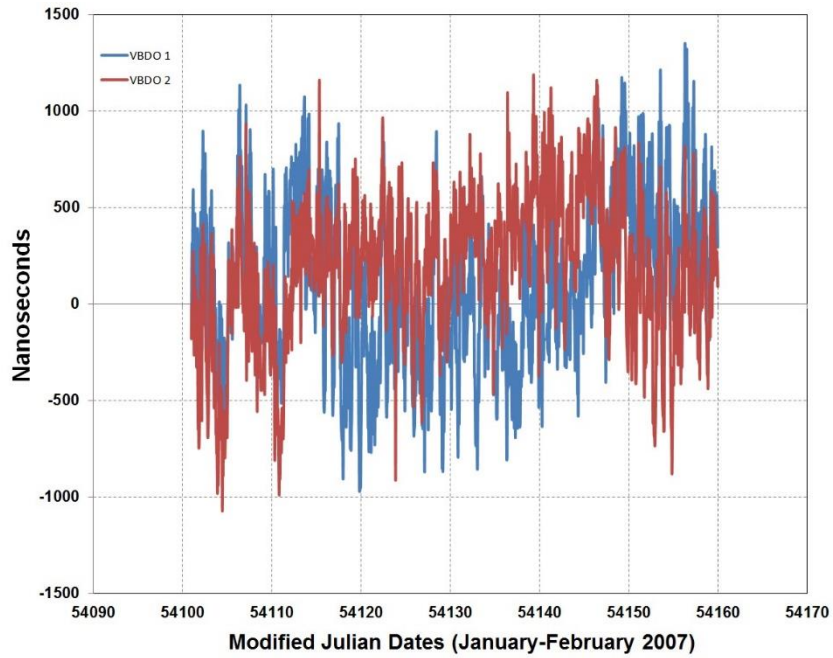


Fig. III.A.17 – Phase of two WWVB disciplined oscillators compared to UTC(NIST).

The frequency stability, estimated with the Allan deviation, of both devices was about 3×10^{12} at an averaging period of one day (Fig. III.A.18). The time stability, estimated with the time deviation statistic, was < 200 ns for the nearly all of the averaging intervals (Fig. III.A.19), indicating the ability of the devices to maintain sub-microsecond synchronization for a clock previously synchronized by GPS.

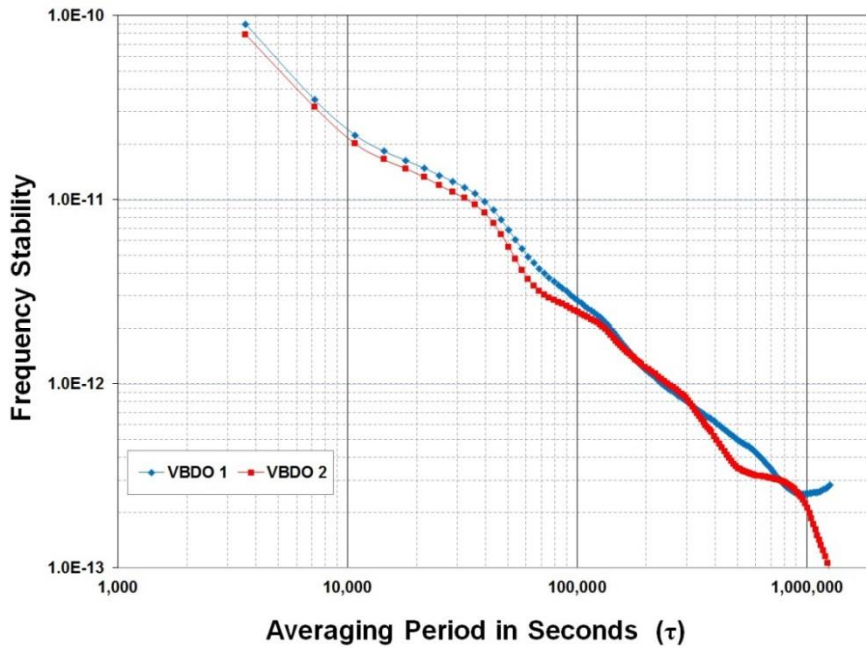


Fig. III.A.18. Frequency stability of two WWVB disciplined oscillators compared to UTC(NIST).




Fig. III.A.19. Time stability (time deviation) of two WWWB disciplined oscillators compared to UTC(NIST).

III.A.5 – Feasibility of a Commercial WWVB Disciplined Oscillator for Integration into GPSDCs

Commercial WWVB products are ubiquitous, with many millions of radio controlled clocks (RCCs) [89] found in businesses and households in the United States. The first WWVB RCC, manufactured by the German company Junghans, was introduced in 1996. Soon afterwards, semiconductor manufacturers such as Telefunken in Germany and Atmel (now part of Microchip) in the United States produced single-chip WWVB receivers, which led to a proliferation of low-cost RCCs [81]. The largest manufacturer of wall clocks is La Crosse Technology of Wisconsin and the largest manufacturer of wristwatches is probably Casio in Japan (Fig. III.A.20). Exact sales figures are not known, but La Crosse alone has sold tens of millions of units, and it seems plausible that the installed base of WWVB clocks and watches in the United States exceeds 100 million.



Fig. III.A.20. A WWVB radio controlled wall clock and wrist watch.

The low-cost WWVB RCCs are relatively simple devices, containing a quartz oscillator, a small receiver permanently tuned to 60 kHz that decodes the incoming time code; and a small ferrite bar or loop antenna that is wrapped with a coil of wire. The receiver is usually only activated once every 24 hours, typically at a time between midnight and 4 a.m. when signal conditions are the strongest. After reading the time code, it uses the information to synchronize the quartz clock, and then deactivates for the next 24 hours. During the period between synchronizations, the quartz clock free runs and accumulates a time error. Immediately after it is synchronized, a WWVB RCC is probably accurate to within < 30 ms (0.03 s), with the actual accuracy dependent upon its design and its distance from the transmitter. However, because the frequency of the quartz clock is never adjusted, and because its frequency offset typically ranges from about 1 to 5 parts in 10^6 , immediately prior to the next synchronization it may have accumulated a time error as large as 500 ms (0.5 s) [90]. Always being accurate to within 0.5 s is often sufficient for making the clock look correct to the human eye, but is insufficient for critical infrastructure timing systems. Such systems would need to access the WWVB signal far more often than once per day, and to use the received information to discipline the frequency of a local oscillator.

No manufacturer currently offers a WWVB disciplined oscillator product (VBDO), but such a product should be capable of the frequency stability indicated in Section III.A.4 (low parts in 10^{12} at an averaging period of one day for most parts of the CONUS, with performance degrading somewhat as a function of distance from the transmitter). Prior to GPS, beginning in the 1960s and continuing to about the year 2000,

several manufacturers sold VBDOs, including Fluke, Hewlett-Packard [91], Spectracom [92], Tracor, and True Time/Kinometrics. Many thousands of units were sold as frequency standards with either 5 MHz or 10 MHz outputs. These products were essentially driven from the market by GPS devices that typically had frequency uncertainties roughly $10\times$ smaller (low parts in 10^{13} at one day of averaging), could also provide sub-microsecond time accuracy, and which could be sold internationally, as they could work reliably anywhere on Earth and not just in the CONUS. Even so, a number of VBDOs remained in operation until in October 2012, when the introduction of the PM time code (Fig. III.A.14) made them obsolete. These legacy products depended upon a carrier sine wave that was only amplitude modulated, and the phase modulation of the carrier caused them to fail.

The development of a commercially-available VBDO for integration into GPSDCs is entirely feasible, even with the PM code. Such a device could also serve as an additional holdover frequency source for the MSCVDC previously described in Section II.C. A modern VBDO, compatible with the PM time code, could be designed by utilizing a WWVB receiver and antenna, a phase detector/comparator, a local oscillator, and control software. The phase comparator would measure the difference between the local oscillator and the incoming WWVB signal. Rather than phase locking to the carrier, a modern VBDO could compare its local oscillator once per second to any recurring element of the PM code, or to the OTM itself. Figure III.A.21 shows the OTM displayed on an oscilloscope. The green trace represents the 1 PPS signal from the secondary UTC(NIST) time scale at the station site. The yellow trace represents the modulated WWVB carrier, with the first cycle of reduced amplitude serving as the OTM. Note that the OTM as transmitted is within one cycle of UTC(NIST).





Fig. III.A.21. The WWVB on-time marker compared to 1 PPS from UTC(NIST).

The collected phase data, obtained from the local oscillator comparisons to WWVB, would be averaged over an appropriate interval, and then be converted, by a microcontroller and firmware, to a frequency correction that is applied to the local oscillator. This process would be repeated continuously to maintain a lock to WWVB. The optimal averaging interval would depend in part on the local oscillator stability – a rubidium oscillator, for example, allows for a longer averaging interval than a quartz oscillator.

A significant obstacle when designing a VBDO is dealing with diurnal phase shifts that occur at sunrise and sunset, typically when it is still daylight at the transmission site and dark at the reception site, or vice versa. These phase shifts rarely occur over a short path, such as the path from Boulder to WWVB, because the short path signal is predominantly groundwave. They primarily occur over longer paths when there is a change from groundwave to skywave propagation, or when the receiver is unable to distinguish between a mixture of skywave and groundwave signals. The phase shifts are typically one or more cycles in magnitude, and thus are multiples of about $16.67 \mu\text{s}$, the period of 60 kHz.

Figure III.A.22 shows a coarse model of the phase variations of WWVB as received in the eastern United States, over an approximate 2000 km path. Sunset and sunrise at the receiver site each occur about two hours earlier than they do at the transmitter site, and a cycle slip occurs during both of the two-hour intervals (the periods on the graph from approximately 0100 to 0300 UTC and 1000 to 1200 UTC). Note that during both the daytime and nighttime hours the phase usually is free from cycle slips, but the phase variations are typically larger at night.



III.B Distribution of UTC(NIST) Frequency to GPSDCs via the PTP Ring Method

The experimental PTP Ring method involves a network of IEEE 1588 Precision Time Protocol (PTP) [93] servers that can distribute time with sub-microsecond accuracy across a large geographic area, potentially across the entire CONUS. Normally, each server would be synchronized to its own GPSDC, which in turn is locked and synchronized to signals from the GPS satellites. However, if the GPS signal is unavailable or untrusted in the area where a given server is located, the server could still maintain synchronization by obtaining a holdover frequency from another PTP server on the network.

In many cases, the PTP server that provides the holdover frequency will derive that frequency from another GPSDC, but one that is located far enough away from the problem area so that it remains unaffected. This scenario is common, because GPS incidents such as jamming or spoofing are usually very local in nature and generally impact only small geographic areas. However, other servers located in the PTP Ring could be independent of GPS and directly referenced to a NIST primary (Section I.A) or secondary time scale (Section I.B), or indirectly referenced to UTC(NIST) via one of the methods described in Section II, such as a NISTDC (Fig. III.B.1). Regardless of where the “fall over” frequency source is obtained, it should provide sufficient holdover for PTP servers located in GPS-denied areas to continue to maintain sub-microsecond time accuracy for multiple days, perhaps indefinitely.

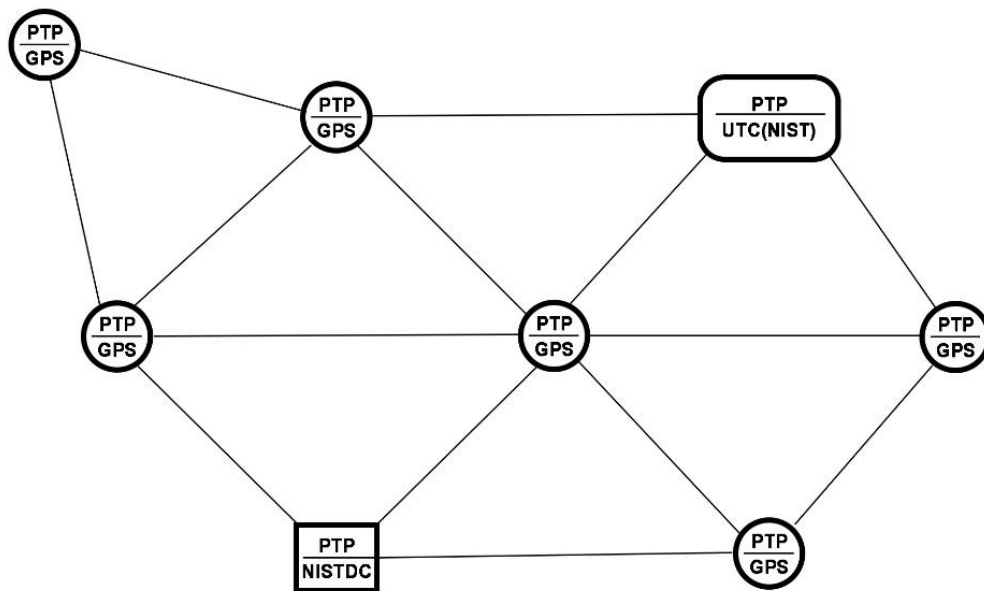


Fig. III.B.1. A section of a proposed PTP Ring, with nodes utilizing different types of reference clock sources.

III.B.1 – Conceptual and Experimental Description of PTP Ring Architecture

In typical ring architectures, each device (node) connects to two others in a circular topology and can acquire connection to other devices in the ring through adjacent hardware in either direction. The proposed PTP Ring method is a ring where each PTP server is a node in the center of a ring of other PTP server nodes, and can thus look at each of those nodes independently across a network, as a PTP client. Because GNSS signal outages are unlikely to cover a large geographic area, accessing the nearest node, even one only a short distance away, may be all a clock needs to maintain a frequency holdover source sufficient for maintaining $\pm 1 \mu\text{s}$ accuracy. However, access to concentric rings further away will also be available.

When a typical PTP client searches for the most favorable reference, it attempts to rank any available sources it sees by broadcasting multicast PTP messages by using several attributes in its data. This is often referred to as the best master clock algorithm (BMCA) [93]. The PTP ring implementation improves on this method by comparing only to known servers using unicast messaging [94]. The PTP servers that are “healthy”, with GPSDCs that are locked to GPS and operating normally, will compare their time and frequency signals to other PTP nodes in the ring. This will characterize the quality of the available PTP synchronization and can be used to calibrate the PTP offset/bias between nodes. It will also denote network changes and how often they occur. Each server can continually rank the PTP network paths with respect to the other servers around it, essentially acting as a self-calibrating network.

When a node cannot receive GPS, it will “fall over” to one of the best sources for PTP holdover, making sure that the source it selects is not also affected by loss of GPS. This affected node, now relying on PTP as its frequency holdover reference, can continue to operate as a local PTP server, and continue to compare itself to other nodes; which will aid in detecting and adapting to a network path change between itself and its reference.

The primary purpose of the PTP Ring method is to use PTP to immediately distribute frequency to GPSDCs during signal outages. PTP can, of course, also distribute time, but the GPSDCs affected by outages were already synchronized at sub-microsecond levels when the outage occurred. As previously noted in the WWVB discussion (Section III.A.4), maintaining frequency stable to within a few parts in 10^{12} after one day of averaging could potentially allow a previously synchronized clock to maintain sub-microsecond synchronization indefinitely.

Synchronous Ethernet (SyncE) methods [95] using the physical layer of a network for frequency transfer have been used to augment PTP for frequency transfer. These features could be implemented in the PTP ring to further improve the frequency holdover. The use of a White Rabbit network [29], which essentially combines PTP and SyncE, is a possible solution. In addition, the use of multiple network paths, which will be readily available in the PTP Ring configuration, could further improve upon the SyncE method for frequency synchronization [96].

When a GPS outage is over and normal reception has been restored, the affected node(s) can resume normal operation and go back to relying on their own GPSDC for both frequency and time. If implemented across a large enough area, the PTP Ring may also prove to be a useful diagnostic tool that can rapidly detect and report on GPS anomalies and where they occur.

III.B.2 – Reliability and Redundancy of PTP Ring Architecture

The PTP Ring method is based on having redundant frequency sources available, and uses this redundancy to increase reliability. When any node in the PTP Ring loses GPS reception, it can rely on any healthy node in the ring to maintain its synchronization. Increasing the number of nodes, and increasing the size of the geographic area covered by the PTP Ring, will in turn increase the level of redundancy and reliability.

The local network connection always represents a single point of failure for a PTP node. If network connectivity at the point of a node is lost, and if GPS reception is also lost, then the clock at that node will be free running and its performance will be no better than the stability of its local oscillator. However, if only parts of the network between two nodes are lost, it may be possible for the PTP packets to be re-routed. As long as this happens during normal operation, the new path will be part of the calibration around the ring. If there are network path changes during periods when one node is relying upon another node for holdover operation, another node in the ring where the path did not change can become the new reference.

Network congestion, even if it does not lead to a node failure, could still lead to the degradation of time and frequency transfer between nodes. Network path or asymmetry changes will affect the quality of the PTP time transfer, but is less likely to affect frequency transfer, especially if SyncE is used [95].

In GPS is unavailable and power is lost, the calibrated PTP links could be used to resynchronize a node when power is restored, but probably with less accuracy than GPS synchronization. Therefore, the use of a tiered power backup system, that includes UPS units and backup generators for the hardware at each node, would greatly improve the reliability of the entire system.

III.B.3 – Physical and Cyber Security of PTP Ring Architecture

Proposed sites for the PTP ring would be required to have secure physical entry for personnel and a reliable network with a high rate of uptime. The network should be monitored and secure, with only applicable ports open. PTP uses UDP ports 319/320, so other ports could be disallowed. The network devices may have the option to be controlled using https (port 443) or SSH (port 22). These connections should be password protected and, for additional security, ports other than the defaults can be used.

III.B.4 – Frequency and Time Stability of PTP Ring with Respect to UTC(NIST)

An initial experiment to prove the concept of the PTP ring was conducted using three PTP servers, each referenced to its own GPSDC. The three servers were connected with a local area network (LAN). A single GPS antenna supplied all three GPSDCs by use of an antenna splitter. The 1 PPS output from each of the three servers was then compared to UTC(NIST) with a multi-channel TIC that sequentially reads each channel (Fig. III.B.2).

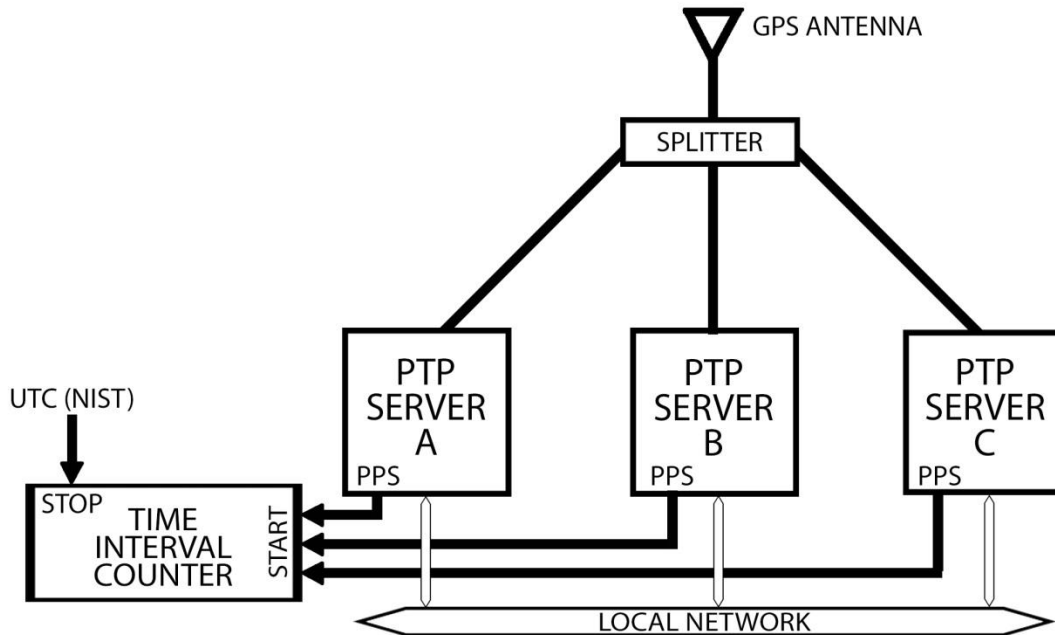


Fig. III.B.2. Block diagram for an experimental setup of a PTP Ring on a local network.

A self-surveyed GPS clock, with no other calibration than having cable delays removed, is often synchronized to within 50 ns of UTC(NIST). That was found to be the case in this experiment with all three servers. Then, the cable to the antenna splitter was removed from Server B, forcing GPS reception to be lost, and forcing Server B to become a PTP client of Server A. As previous experiments have shown, a PTP client, using the standard profile settings on a LAN, can be synchronized to within 100 ns of its reference [97]. That was confirmed here, as shown in Fig. III.B.3. The average time offset, with respect to UTC(NIST), is < 10 ns for Server A, which is still referenced to GPS. The Server B data, without GPS connected, is much noisier, but its average time offset is still < 50 ns.

Sub-microsecond time synchronization is, of course, more difficult when the nodes in a PTP Ring do not all reside on the same LAN. Subsequent experiments of the PTP Ring will involve moving at least one of the servers to a geographically separated but related network, for example, moving a server to another site that is connected via a virtual private network (VPN), as shown in Fig. III.B.4. In this example, the remote Server C is connected to its own GPS antenna, and Servers A and B are both configured as unicast clients, specifically using the IP address of Server C, so they do not use each other as their PTP reference. In this example, the network asymmetry that exists between Servers A and B and Server B may result in PTP synchronization errors that exceed 1 μ s, as shown in previous experiments with a similar configuration [97]. Therefore, using PTP as a frequency holdover source that preserves existing GPS synchronization, rather than relying upon PTP as a time synchronization source, is a key principle of the PTP Ring method.




Fig. III.B.3. Comparison of PTP servers (Server A is locked to GPS, Server B has no GPS but is locked to Server A in PTP client mode).

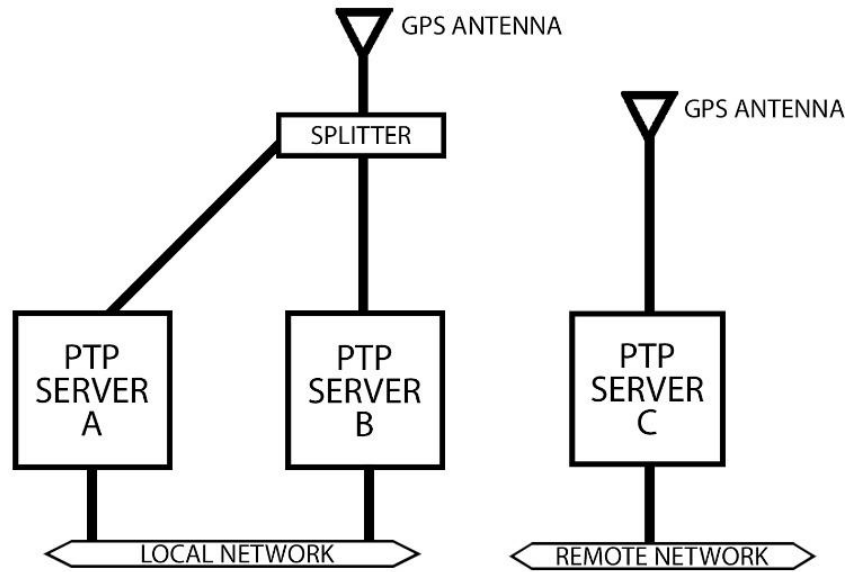


Fig. III.B.4. Experimental setup of a PTP Ring where one server resides on a remote network.

Comparisons between the NIST Boulder campus and the WWVB site in Fort Collins, a distance of 78 km, should provide a favorable location for future tests of the PTP Ring method. A virtual local area network (VLAN) with a 100 megabit circuit is already in place between the two sites, which could be readily used for unicast PTP transmissions, and will also allow trials using different types of network connectivity.

III.B.5 – Feasibility of a Nationwide PTP Ring Architecture

Although it has not yet progressed past the conceptual phase, the PTP Ring method can potentially serve as a key part of resilient time distribution architecture. The idea is simple - PTP can deliver a stable holdover frequency source over a network to GPSDCs that are currently experiencing a GPS outage. The stable frequency source can be supplied by clocks synchronized to UTC(NIST), but in most cases can be adequately supplied by any GPSDC that is located far enough away to be unaffected by the outage. If such a network were built to cover the entire CONUS, for example, only a massive GPS outage could cause all of the GPSDCs in a PTP Ring to lose their frequency lock.

As is usually the case with wired solutions that carry timing signals, costs are high when compared to wireless solutions. Subscriptions to private and fixed-path networks are available to be purchased, but they may be expensive. The ‘last mile’ to a node from a network backbone or data center would also need to be procured in order to utilize a non-typical network type. The number of nodes required would need to be studied, based on how symmetric and consistent the network connectivity is between all of the nodes, and how much redundancy is desired. Also, finding sites that meet the security criteria in Section III.B.3 may be challenging in some areas.

IV. Indirect Distribution of UTC(NIST) via non-GNSS, non-NIST Signals

Time signals that are not controlled by NIST, or directly referenced to UTC(NIST), can still provide UTC(NIST) indirectly. If these signals are GNSS independent, they can become an integral part of a resilient time distribution architecture and can contribute to and support the requirements of *Executive Order 13905* [1].

There are at least two methods that would allow a non-NIST time broadcast to indirectly distribute UTC(NIST). The first method involves applying a UTC(NIST) correction to the received time signal. The correction could be obtained from the received signal itself – for example, in a similar fashion to the UTC(USNO) correction contained in the GPS broadcast – or it could be obtained independently from another source, for example via a network connection. In either case, the generation of the corrections would require NIST participation in a signal monitoring/measurement program, so the time differences of non-NIST signals, with respect to UTC(NIST), are always known. The second method would be for NIST to actively participate in the synchronization of the non-NIST signal, by transferring time from the UTC(NIST) time scale to the place where the non-NIST signal originates.

Two systems that can potentially indirectly distribute UTC(NIST) are discussed in the following sections. Section IV.A discusses the Satellite Time and Location (STL) service operated from the Iridium satellites, and Section IV.B discusses the possible future implementations of the eLoran radio navigation system. Other non-NIST systems not discussed here could potentially serve as distribution sources if they were monitored and/or synchronized by NIST.

IV.A Indirect Distribution of UTC(NIST) via the Satellite Time and Location (STL) Service

The Satellite Time and Location (STL) service is operated by Satelles (www.satellitesinc.com), a company headquartered in Reston, Virginia. The STL service provides secure time and location signals that originate from low Earth orbit (LEO) satellites that are independent of GNSS. The service was designed to be immune to certain GNSS vulnerabilities. For example, it utilizes cryptographic techniques to protect the signal from spoofing attacks. In addition, the raw signal power of the STL signals as received on Earth is from $300\times$ to $2400\times$ stronger than GPS, making them less susceptible to jamming and interference, and unlike GNSS, usable with indoor antennas [98]. Commercial STL clock products are now available from both Satelles and their licensed business partners.

The STL signals originate from the Iridium constellation, which became operational in 1998. The constellation requires 66 satellites (11 satellites in six orbital planes) for worldwide coverage, but more than 70 satellites are in active service, with additional spare satellites available that can serve in case of failures. Like GNSS, the STL signals are broadcast in the L band, but at slightly higher frequencies, ranging from 1616 MHz to 1626 MHz. The use of the L band allows small, low-cost equipment to receive the signal.

The LEO satellites in the Iridium constellation are more than $25\times$ closer to Earth than GPS satellites, flying at an altitude of just 781 km, as opposed to the 20 180 km altitude of GPS. Their inclination angle is 86.4° , meaning that the satellites travel in a nearly polar orbit, and their orbital planes intersect near the poles. At the poles, between six and eight satellites should be visible at once, but typically one satellite is visible from the equator. The orbital period is about 100 minutes; therefore the satellites circle the Earth about $14\times$ per day.

Being closer to Earth is one reason that STL signals received on Earth are much stronger than GNSS signals, and provide STL with the advantage of being more capable than GPS of working inside buildings, in underground locations, and in urban canyons. However, LEO satellites have disadvantages with respect to coverage when compared to GPS. The diameter of the coverage area, or footprint, of an Iridium satellite is just 3 000 km, or about one-fourth the size of a GPS footprint. Therefore, even though there are more than twice as many Iridium satellites in orbit as GPS satellites, more GPS satellites will always be visible to receivers on Earth [99].

The Iridium constellation was originally designed, not for PNT usage, but rather to provide voice and data communications to satellite phones and pagers, with worldwide coverage. It did, however, support the equivalent of common “text messages” which can be transmitted during a single 90 ms signal frame. The STL service uses this capability to transmit single-frame messages, or “bursts”, that are specifically designed for PNT applications. To the satellite, the STL burst appears to be a simple text message but is actually a carefully formatted set of data that produces a spread-spectrum RF signal pulse in a proprietary format. The burst contains a pilot channel and a data channel, both highly encoded, that are received on Earth about once every second.

There is enough coding gain for an STL receiver to make time (pseudorange) and frequency measurements of each burst at attenuations of up to 39 dB relative to unobstructed reception. That allows the signal to penetrate buildings and other obstructions and work indoors. A receiver can determine the time with sub-microsecond accuracy from a single burst, typically in less than two seconds [98, 100].

IV.A.1 – Frequency and Time Stability of STL with respect to UTC(NIST)

The STL service is known to be a less accurate timing source than GNSS, but one that still easily demonstrates sub-microsecond accuracy. Previously published measurements in 2018 compared an STL clock with an indoor antenna to a GPSDC over a two-week interval. The average time offset of the STL clock with respect to the GPSDC was < 200 ns, with a peak-to-peak variation of about 600 ns [100].

During the summer of 2020, NIST conducted its own STL measurements by directly comparing the 1 PPS output from an STL clock, shown in Fig. IV.A.1, to the UTC(NIST) time scale using a time interval counter. The antenna was located indoors during the measurement, resting on top of the device, as shown in the figure.



Fig. IV.A.1. A prototype STL disciplined clock.

Measurements were continuously collected for a 50-day period (06/25/2020 to 08/13/2020, MJD 59025 to 59074), with the exception of a break in the data around MJD 59063 when the measurement system was temporarily stopped. During the measurement period, one-minute averages that compared a low-cost GPSDC to UTC(NIST) were simultaneously recorded, allowing us to compare the accuracy of the STL clock to the GPS clock. The two units each have local quartz oscillators that are believed to be of comparable quality. The GPSDC used a rooftop antenna, connected with ~ 25 m of RG-58 coaxial cable.

No calibration delay constant for the antenna cable was entered into the STL clock, as the short cable shown in Fig. IV.A.1 was only 15 cm in length and has a delay of < 1 ns. The GPSDC was also measured without delay compensation, to simulate the way they are normally used in industrial settings. However, the cable delay was measured at 123.7 ns and a separate data set was kept that compensated for this delay. The delays introduced by the cables connecting the STL clock, the GPSDC, and UTC(NIST) to the counter were each measured and removed from the readings to ensure a fair comparison.

As shown in Fig. IV.A.2, the uncalibrated GPSDC had an average time offset of -133 ns with respect to UTC(NIST), reduced to just -9 ns when compensated for the antenna cable delay. The STL clock had a larger time offset of -271 ns, but still demonstrated sub-microsecond accuracy without delay calibration. The GPSDC output was much quieter, with a peak-to-peak variation of just 50 ns as opposed to 382 ns for STL. Most of the STL readings, however, fell within a ~ 90 ns window, from about -220 to -310 ns.

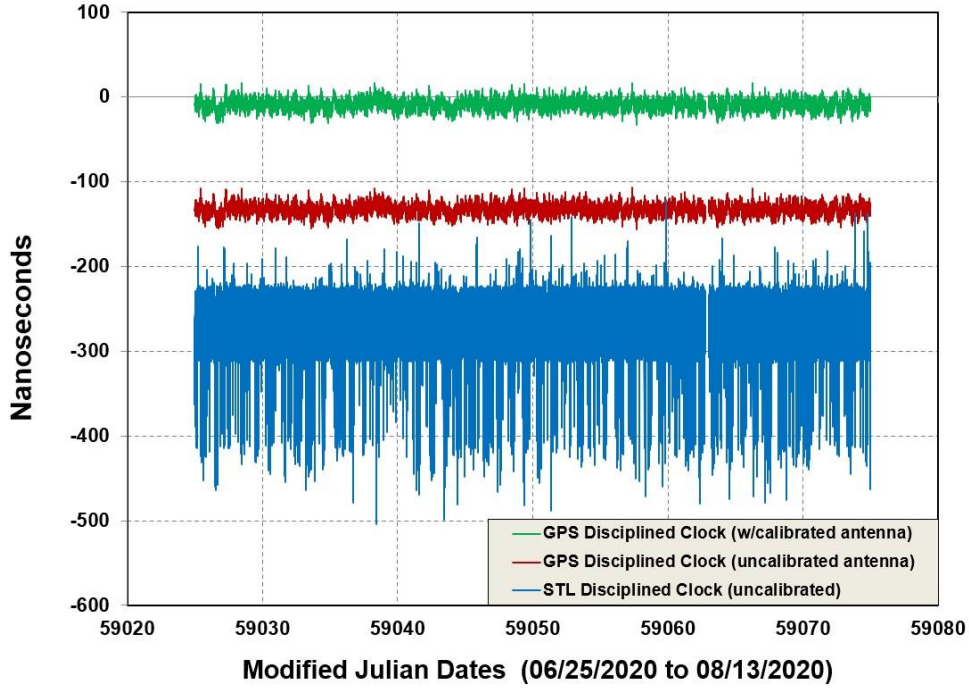


Fig. IV.A.2. Accuracy comparison of an STL clock to a GPSDC, each with respect to UTC(NIST).

Figure IV.A.3 shows the stability of the STL and GPSDC data as estimated with the time deviation (TDEV) statistic. The GPSDC instability never exceeded 4 ns at any averaging period and was < 2 ns after one day of averaging. The STL instability exceeded 20 ns at short averaging periods (out to about 10 minutes), but was < 3 ns after averaging for one day.




Fig. IV.A.3. Stability comparison of an STL clock to a GPSDC, each with respect to UTC(NIST).

The long-term stability of the STL clock is further evidenced by the data in Fig. IV.A.4, where the one-minute data points previously graphed in Fig. IV.A.2, were further averaged to produce one data point per day. The average time offset of course remains the same, -271 ns, but the peak-to-peak variation across the 50-day period is just 13 ns, demonstrating that sub-microsecond time accuracy from an STL clock is repeatable from day to day.

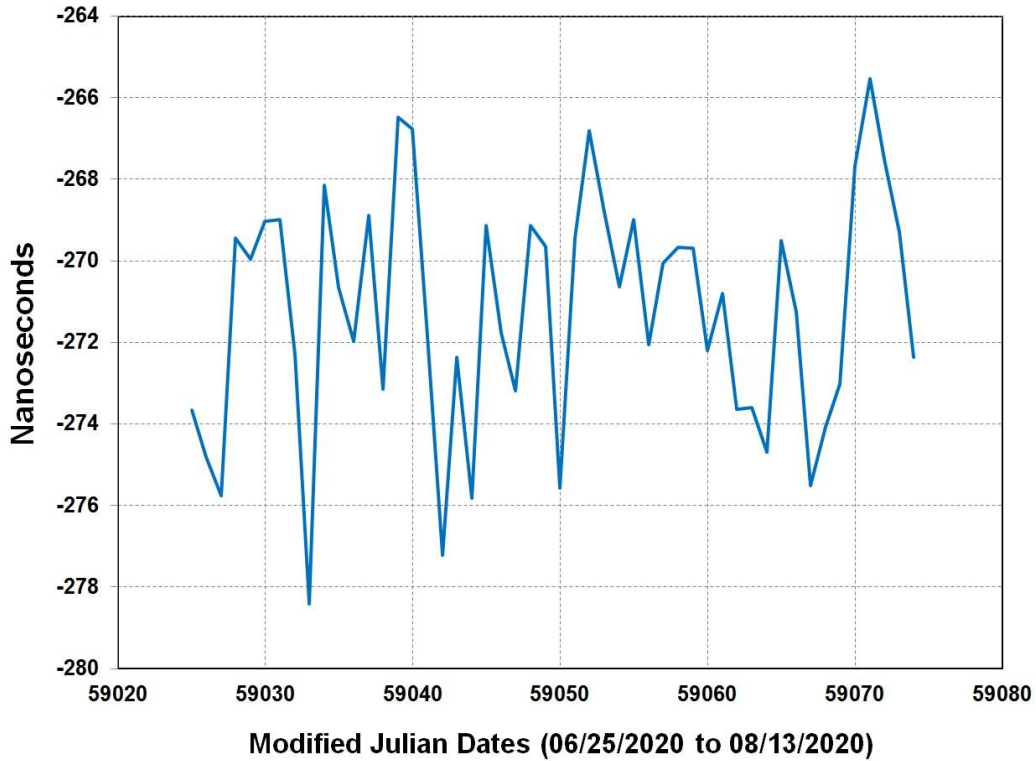


Fig. IV.A.4. Time differences of STL clock with respect to UTC(NIST), one-day averages.

IV.A.2 – Possible Establishment of a STL Ground Monitoring Station at NIST

At the time of writing (August 2020), NIST and Satelles are discussing options, which might include a signed memorandum of understanding (MOU) between the two organizations, that could lead to the establishment of an STL ground monitoring station in Boulder, Colorado, referenced to the primary UTC(NIST) time scale. If accomplished, this would allow Satelles to have continuous access to a completely GNSS independent time reference to assist in the UTC time synchronization of STL, and would provide additional validation and verification of the STL signals. It could also potentially allow for a UTC(NIST) time correction to eventually be included in the STL broadcast, analogous to the UTC(USNO) correction included in the GPS broadcasts.¹

¹ An agreement was signed and the ground station became operational at NIST in 2021.

IV.B Indirect Distribution of UTC(NIST) via eLoran

Loran-C (originally an acronym for LOnG RAnge Navigation), was a terrestrial based radio navigation system operated by the United States Coast Guard (USCG) that predated the space-based GPS system by several decades. The Loran-C system was turned off in the United States in February 2010 [101]. Prior to that date, 18 Loran stations were operational in the CONUS, and an additional six stations were in operation in Alaska, all broadcasting on a carrier frequency of 100 kHz.

At the time of its termination, the Loran-C system was undergoing a modernization effort in the hopes that it could be repurposed as a backup or alternative system to GPS, for both for navigation and time distribution. The modernized system, known as enhanced Loran, or eLoran, had a new modulation format that included, among other things, a digital time code [102]. Legacy Loran had long been a trusted source of frequency whose signals, like WWVB (Section III.A), resided in the LF part of the radio spectrum where they benefitted from stable groundwave propagation. However, Loran had many advantages over WWVB including much higher radiated power levels (as high as one megawatt), a large number of stations as opposed to just one, and a modulation format that made it easier for receivers to stay locked to the same cycle and thus eliminate cycle slips. The additional features added by eLoran, especially the addition of a digital time code and its ability to accept time corrections from differential stations, made the system very attractive as a time distribution source.

As a result, even after its demise in the United States in 2010, and after stations were subsequently turned off in Canada, Europe, Russia, and elsewhere, the interest in eLoran for time distribution remains high. Station tests and demonstrations are periodically conducted and the possibility remains that eLoran could eventually resume operation. The *National Timing Resilience and Security Act*, signed into law in December 2018 [103], directed the Secretary of Transportation to establish a wireless, terrestrial, difficult-to-disrupt, wide-area coverage system that could be used underground and indoors and that was synchronized to UTC. These requirements all pointed to the revitalization of eLoran. In addition the act stated that if the Secretary of Transportation needed any existing Loran property, equipment, or radio spectrum allocations currently held by the USCG to establish the new system, that those properties would be transferred to the Department of Transportation.

At this writing (August 2020), the exact status of the eLoran revitalization effort and the probability that eLoran may return are unknown to NIST. With that in mind, Section IV.B.1 discusses what the architecture of a future eLoran system might look like, and Section IV.B.2 discusses how NIST could potentially contribute to the eLoran effort by enabling the system to indirectly distribute UTC(NIST).

IV.B.1 – Possible Architecture of a Future eLoran System

Prior to its 2010 shut down, Loran-C provided positioning and navigation services to most of North America. In addition to the aforementioned 24 stations in the U. S., an additional five stations resided in Canada. The 29 station network was divided into 12 units called chains, with most stations “dual rated” and belonging to two chains. Each chain covered a specific geographic area, the West Coast chain, the Great Lakes chain, the Northeast chain, and so forth.

Even if eLoran reemerges, it is probably unrealistic to expect that a network of stations that large will be constructed again, as many of the former Loran-C transmitter sites were dismantled or repurposed. The reemergence of a smaller system, dedicated solely to timing, is a more realistic scenario. In keeping with that reasoning, UrsaNav, a company heavily involved in the development of eLoran transmitters and receivers, as well as other LF radio technologies, has proposed a UTC distribution system that could cover the entire CONUS with just four eLoran transmitters, each radiating one megawatt of

power. The four transmitter sites would be Wildwood, New Jersey, where the U. S. Coast Guard Loran-C Support Unit was formerly located, and three former Loran-C station sites: Dana, Indiana; Boise City, Oklahoma; and Fallon, Nevada. Tentatively known as the Basic eLoran Timing Service (BeTS), this system, with just four transmitters and no differential reference stations, could potentially provide $1 \mu\text{s}$ synchronization throughout the CONUS, as shown in Fig. IV.B.1.

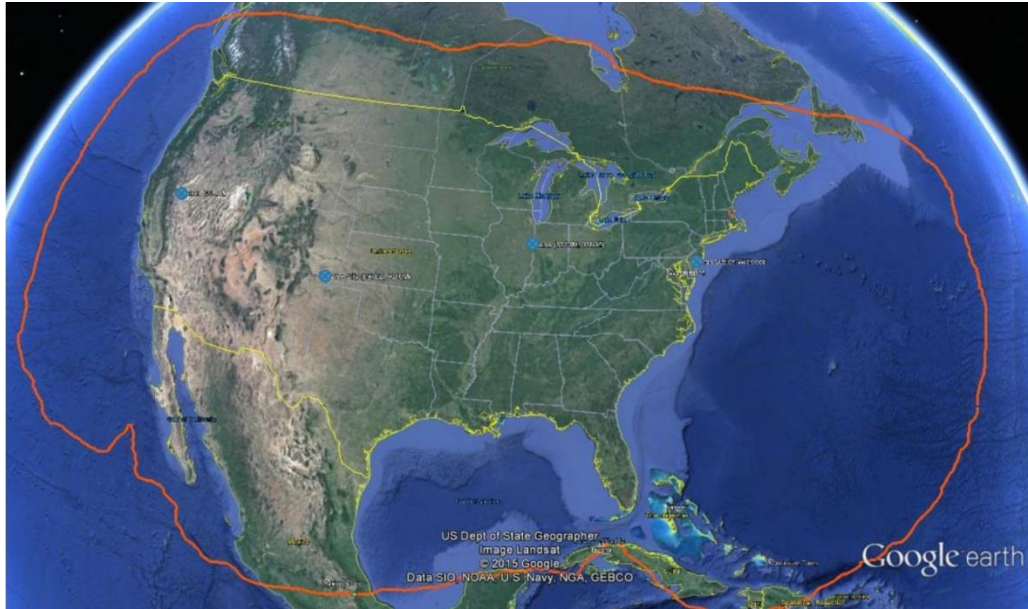


Fig. IV.B.1. Estimated coverage area of eLoran BeTS service with four one megawatt transmitters [103].

Ursanav has also proposed a Precision eLoran Timing Service (PeTS) that would add 71 differential reference stations that are strategically located to be near the top 50 major metropolitan areas, ports, harbors, and airports (Fig. IV.B.2). With the addition of these differential corrections, which would be included in the Loran Data Channel (LDC) messages, the potential synchronization accuracy of the service drops to below $0.1 \mu\text{s}$ (100 ns) with respect to UTC [104].

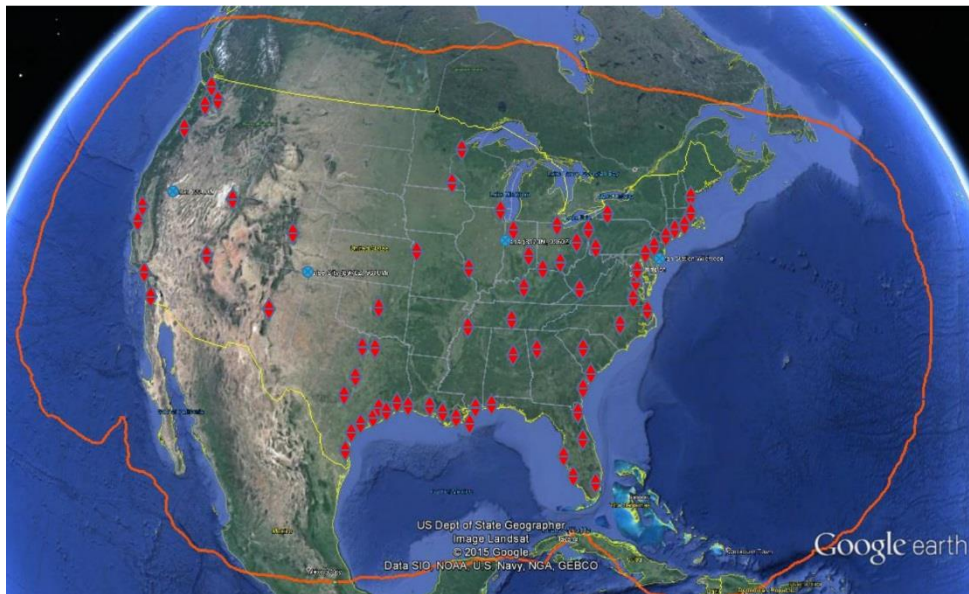


Fig. IV.B.2. The eLoran PeTS service, supplemented by 71 differential reference stations [103].

IV.B.2 – Potential Methods for Indirectly Distributing UTC(NIST) via eLoran

The synchronization of the Loran-C stations to UTC was required by an agreement between the USCG and the USNO, beginning in 1971 and continuing until the demise of the system. During all of this period, the USNO published daily Loran-C time differences all stations in a document known as the *Series 4* [105]. During nearly all of this same period, NIST also monitored Loran-C and published daily time differences for two or three Loran-C stations in its monthly *Time and Frequency Bulletin* [106].

Beginning in the early 1990s, the Loran-C stations were synchronized with GPSDCs, with the Hewlett-Packard GPS Smart Clock used as the primary reference at many station sites after about 1997. This direct synchronization to GPS did, of course, invalidate the premise that Loran-C, and the proposed eLoran system that followed, were GNSS independent, which made them unsuitable as a GPS backup system. As a result, the USCG and Symmetricom (now owned by Microchip) began work on a proposed system that featured two control centers that would be linked via TWSTFT (Section II.B) to either UTC(USNO) or UTC(NIST), and preferably to both time scales [107]. Descriptions of the TWSTFT system, which would also allow for the comparison of cesium clocks at geographically separated transmitter sites and for the formation of a Loran time scale, were presented [108] but remained unfinished when the system was terminated.

The idea that TWSTFT or another time transfer mechanism could link eLoran to NIST or other national timing laboratories with nanosecond-level accuracy, paved the way for the system to indirectly distribute any source of UTC. Originally, fields were added into the eLoran message structure that identified the UTC time source that supplied corrections to the signals [109]. The current international eLoran signal standard, *SAE 9990/2* [110], has evolved in recent years but contains many of the original elements. Among other things it allows (as part of message 0) phase corrections where the reference station can identify itself and indicate the level of its time base quality, with the highest level, Level 1, being < 15 ns RMS. An eLoran station can transmit corrections from multiple reference stations, potentially allowing time from several UTC scales to be obtained from the same broadcast.

To briefly examine how this could work, consider that an eLoran station is synchronized by a combination of a local time scale and one or more remote time scales. The local time scale is normally an ensemble of three cesium clocks [104] that are synchronized and/or disciplined to the remote time scale.

In the case of the lower accuracy BeTS service (Fig. IV.B.1), all that is required of the remote time scale is to synchronize the ensemble time scale at the eLoran transmitter sites. Receivers, assuming that they know their position; can record the time of arrival (TOA) of the eLoran signal, estimate and compensate for the path delay, and then synchronize their clocks to the time broadcast by the transmitters, typically to within 1 μ s.

The higher accuracy PeTS service (Fig. IV.B.2) would require all of the differential reference stations, in addition to the transmitters, to contain both an eLoran receiver and a GNSS-independent reference clock. The reference clock is used to measure the time of arrival (TOA) from the eLoran transmitter of interest, and to compare that measurement to the signal's expected TOA. The expected TOA is modelled with a range calculation, based on the distance of the differential station from the transmitter, that is improved by applying three delay terms known as the primary factor (PF), the secondary factor (SF), and the additional second factor (ASF). The PF and SF are both known and modelled as constants. The PF is the difference between the propagation delays of the signal in the Earth's atmosphere, as opposed to the delays in free space, and the SF accounts for the propagation delay of the signal as it travels over sea water. The ASF is obtained from a grid like map of the area around the transmitter. It is the additional delay of the signal as it passes over land and elevated terrain as opposed to over sea

water, and depends on numerous factors including the soil type and ground conductivity [109].

The difference between the measured and expected TOA is converted to a correction that is returned by the differential station to the transmitter (Fig. IV.B.3). There, the differential corrections are added to the LDC and broadcast to position, navigation, and timing users. Timing receivers can then be programmed to automatically apply the corrections to 1 PPS and other output signals, so that errors in the transmitted timing information are corrected at the receiver site, enabling synchronization of 0.1 μ s (100 ns) or less with respect to UTC.

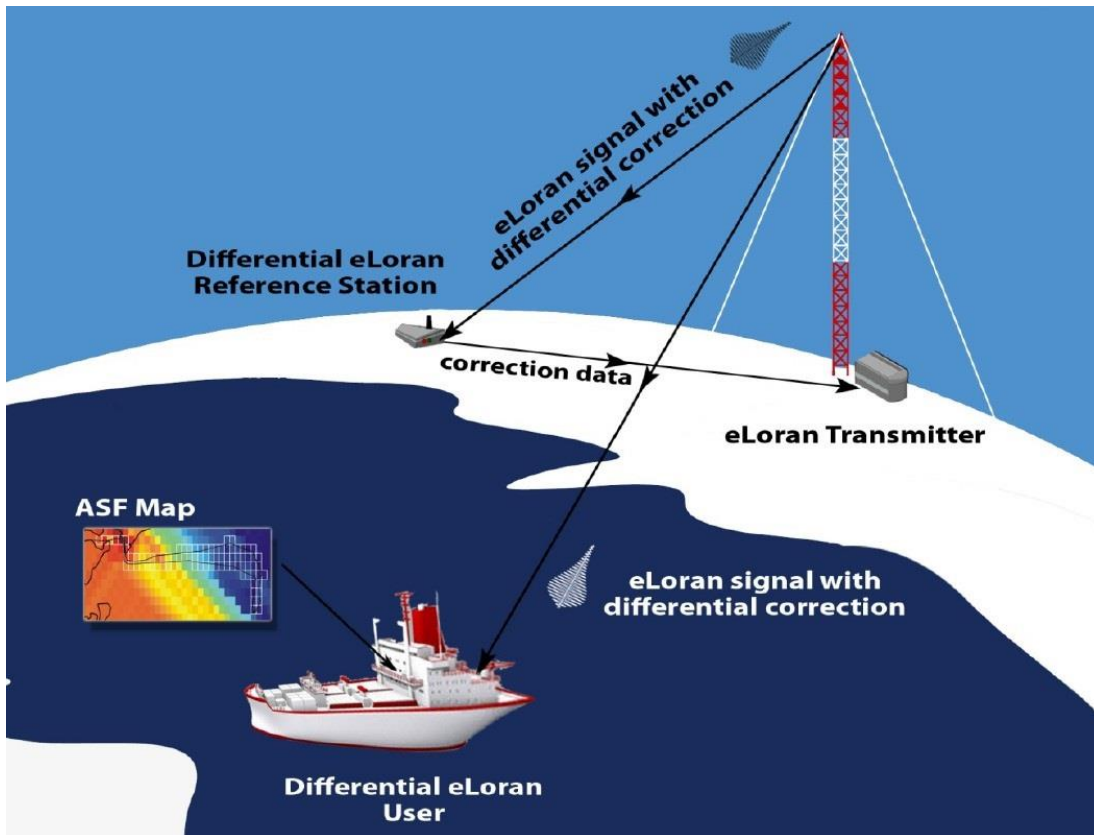


Fig. IV.B.3. A differential eLoran reference station sends corrections back to the transmitter [111].

The indirect distribution of UTC(NIST) via eLoran could be accomplished by having the remote time scale consist of one or more GNSS independent inputs that NIST could potentially supply, such as time over fiber (Section II.A), TWSTFT (Section II.B), or some variation of the MSCVDC method (Section II.C). This would in turn result in a UTC(NIST) correction that would be included in the LDC [110] and broadcast to timing users.

V. Reception of UTC(NIST) Time Codes for GNSS Time Validation

The resilient time distribution architecture described in this report is focused on systems that can deliver signals providing $\pm 1 \mu\text{s}$ accuracy, an accuracy level that, as noted in the Introduction, is required to reduce or alleviate the dependencies that critical infrastructure systems have on GNSS. However, systems referenced to UTC(NIST) that are not capable of microsecond-level accuracy can still contribute to a resilient time distribution architecture, by providing time-of-day redundancy checks for GNSS receivers, and thus validating that their output is correct.

The time received from GNSS satellites is normally regarded as the “truth” and unquestionably accepted by most systems. However, there are at least three ways that a GNSS clock can produce incorrect time. The first way is if the satellites simply broadcast the wrong time. These errors are exceptionally rare, especially in the case of GPS, but can happen. One example is the approximate $13 \mu\text{s}$ error in the UTC offset parameters that were broadcast by half of the GPS satellites on January 25-26, 2016 [112]. A second way is spoofing, which is an attack by a radio transmitter located near the target that interferes with the legitimate GNSS signal and delivers false information. The third, and most common way, is jamming, which is intentional or unintentional RF interference that prevents the receiver from receiving satellite signals. In this case, the GNSS clock will go into holdover and begin to accumulate a time error at a rate usually determined by the accuracy and stability of its local oscillator.

In the case of jamming, a GNSS clock will usually provide some type of alarm or indication that the time it produces is incorrect, because it is easy enough to determine when no satellites are being received and the local oscillator is unlocked. However, in the case of a broadcast error or spoofing, the GNSS clock will be completely “tricked” unless it has some mechanism for detecting that the time it is receiving is wrong. One simple way to do this is to compare the received GNSS time to a GNSS-independent time code. The January 2016 incident produced a sufficiently small error that could only be detected by comparisons to time sources that had microsecond-level resolution and accuracy, such as UTC(NIST) [112]. However, time errors caused by spoofing are likely to be much larger, and the freely accessible time codes described in this section are capable of detecting millisecond-level and larger errors. Their integration into GNSS clock systems can help with the detection of broadcast errors or spoofing incidents, and by doing so help validate GNSS time.

Section V.A discusses the reception of UTC(NIST) time codes via the Internet, and Section V.B discusses the reception of UTC(NIST) time codes via radio stations WWVB, WWV, and WWVH.

V.A Reception of UTC(NIST) Time Codes via the Internet

The NIST Internet Time Service (ITS) provides time codes in the NTP format [28] from groups of servers located on the NIST campuses in Boulder, Colorado and Gaithersburg, Maryland, from the WWV/WWVB station site in Fort Collins, Colorado, and from JILA, a joint institute of the University of Colorado and NIST, which is also located in Boulder. The five NIST-Boulder servers are directly connected to the primary UTC(NIST) time scale (Section I.A), and the Gaithersburg and Fort Collins servers (five at each site) are directly connected to the secondary UTC(NIST) time scales. The two JILA servers are connected to a cesium clock that is regularly compared to the primary UTC(NIST) time scale. The ITS servers handles tens of billions of timing requests per day [113].

Table V.A.1 provides a list of the publicly accessible ITS servers, including their names, IP addresses, and locations. An NTP request sent to the address *time.nist.gov* will automatically be directed to one of the 17 servers on the list. The most current server list can be obtained from [114].

Table V.A.1. List of publicly accessible servers maintained by the NIST Internet Time Service (ITS).

Public Server Name	IP Address	Location
time.nist.gov	Global address (load balanced)	Boulder, Fort Collins, Gaithersburg
time-a-g.nist.gov	129.6.15.28	NIST, Gaithersburg, Maryland
time-b-g.nist.gov	129.6.15.29	NIST, Gaithersburg, Maryland
time-c-g.nist.gov	129.6.15.30	NIST, Gaithersburg, Maryland
time-d-g.nist.gov	129.6.15.27 2610:20:6f15:15::27 (IPv6)	NIST, Gaithersburg, Maryland
time-e-g.nist.gov	129.6.15.26 2610:20:6f15:15::26 (IPv6)	NIST, Gaithersburg, Maryland
time-a-b.nist.gov	132.163.96.1	NIST, Boulder, Colorado
time-b-b.nist.gov	132.163.96.2	NIST, Boulder, Colorado
time-c-b.nist.gov	132.163.96.3	NIST, Boulder, Colorado
time-d-b.nist.gov	132.163.96.4 2610:20:6f96:96::4 (IPv6)	NIST, Boulder, Colorado
time-e-b.nist.gov	132.163.96.6 2610:20:6f96:96::6 (IPv6)	NIST, Boulder, Colorado
time-a-www.nist.gov	132.163.97.1	WWW, Fort Collins, Colorado
time-b-www.nist.gov	132.163.97.2	WWW, Fort Collins, Colorado
time-c-www.nist.gov	132.163.97.3	WWW, Fort Collins, Colorado
time-d-www.nist.gov	132.163.97.4 2610:20:6f97:97::4 (IPv6)	WWW, Fort Collins, Colorado
time-e-www.nist.gov	132.163.97.6 2610:20:6f97:97::6 (IPv6)	WWW, Fort Collins, Colorado
utcnist.colorado.edu	128.138.140.44	JILA, University of Colorado, Boulder
utcnist2.colorado.edu	128.138.141.172	JILA, University of Colorado, Boulder

In addition to the publicly accessible servers, the ITS offers a group of four servers that provide authenticated time via NTP. The time messages are authenticated using symmetric-key encryption in a manner that is fully compatible with the published NTP documentation (autokey and asymmetric key modes will not be used.) The service is free, but requires registration with NIST [115]. Each registered user is assigned a unique encryption key, which will be linked to the IP address of the user's system. A registered user will be able to communicate with the authenticated server using this assigned encryption key or using a default key of 0, which is equivalent to disabling the encryption algorithm. Table V.A.2 provides a list of

the authenticated ITS servers, including their names, IP addresses, and locations. The most current server list can be obtained from [114].

Table V.A.2. List of authenticated NTP servers maintained by the NIST Internet Time Service (ITS).

Authenticated Server Name	IP Address	Location
ntp-b.nist.gov	132.163.96.5	NIST, Boulder, Colorado
ntp-www.nist.gov	132.163.97.5	WWW, Fort Collins, Colorado
ntp-c.colorado.edu	128.138.141.177	JILA, University of Colorado, Boulder
ntp-d.nist.gov	129.6.15.32	NIST, Gaithersburg, Maryland

V.B Reception of UTC(NIST) Time Codes via WWVB, WWV, and WWVH

UTC(NIST) time codes, independent of GNSS, are also broadcast from three NIST radio stations, the LF radio station WWVB, and the high frequency (HF) stations WWV and WWVH. The details of the WWVB time code, both the legacy AM code and the newer PM code, were previously provided in Section III.A.1. WWVB and WWV are referenced to the secondary NIST time scale in Fort Collins described in Section I.B, and WWVH is referenced to a cesium clock disciplined to the primary NIST time scale in Boulder, with the ability to automatically switch to the secondary time scale in Fort Collins by use of the MSCVDC method (Section II.C).

WWV and WWVH provide time codes that are similar to each other and documented in [80]. WWV continuously broadcasts time codes from transmitters on six different frequencies, four of which (2.5, 5, 10, and 15 MHz) are also utilized by WWVH. Because propagation on the HF bands varies according to time-of-day, distance from the transmitter, and atmospheric conditions; broadcasting from multiple frequencies was done to help ensure that at least one frequency is usable at all times.

A digital time code is provided by WWV and WWVH on a 100 Hz subcarrier on each of their allocated frequencies that can be decoded with appropriate receiving equipment in much the same fashion as the WWVB time code. In addition, both WWV and WWVH distribute an audio time code, which is often useful for spot checks and troubleshooting of timing equipment, and for use as a “soundtrack” on audio or video recordings. The audio time code includes seconds pulses heard every second (except for the 29th and 59th seconds of each minute). The top of each minute is preceded by a voice announcement of time – a man’s voice is used on WWV and a woman’s voice on WWVH – followed by an audio tone of 800 ms in duration (1000 Hz on WWV and 1200 Hz on WWVH) that marks the UTC minute. The UTC hours are marked with tones of 1500 Hz on both stations. Table V.B.1 lists the locations, frequencies, and radio power of the NIST time signal radio stations.

Table V.B.1. Locations, frequencies, and radiated power of the NIST time signal radio stations.

Station	Location	Frequency (MHz)	Radiated Power (kW)
WWVB	Fort Collins, Colorado	0.06 (60 kHz)	70
WWV	Fort Collins, Colorado	2.5	2.5
		5	10
		10	10
		15	10
		20	2.5
		25	2.5
WWVH	Kauai, Hawaii	2.5	5
		5	5
		10	10
		15	10

VI. Acknowledgments

The authors thank and acknowledge the Department of Homeland Security for their sponsorship of this effort and for being given the opportunity to write this report. Special thanks are due to Brannan Vilee, Rachel Wallner, and Ernest Wong of the United States Department of Homeland Security, Science and Technology Directorate (DHS S&T); and James Platt of the Cybersecurity and Infrastructure Security Agency's National Risk Management Center (CISA/NRMC).

This report was partially funded by DHS in their continued effort to support more secure and resilient infrastructure.

In addition, the authors thank the following people for their contributions: former NIST staff member Stefania Romisch, current NIST staff members Jonathan Hardis, James McCarthy, Jim Spicer, and Katy Stephenson; the staff at NIST radio station WWVH – Dean Okayama, Dean Takamatsu, Chris Fujita, and Adela Ochinang; Monty Johnson of OPNT; Michael O'Connor and Christina Riley of Satelles; John Clark, Kevin Franzen, Dave Eckhard, and Demetrios Matsakis of Masterclock, Inc., and Frank van Grass and Kevin Croissant of the University of Ohio.

Note that the manufacturer and model names found in this report, as well as the references made to private entities, were included for the sake of technical completeness only, and do not express or imply any endorsement by NIST. Other products may be found to work equally as well or better, and other private entities may be found that offer similar services or products.

VII. References

- [1] Title 3, The President, “Executive Order on Strengthening National Resilience through Responsible Use of Positioning, Navigation, and Timing Services,” *Executive Order 13905*, Federal Register, vol. 85, no. 3, pp. 9359-9361, signed February 12, 2018, published February 18, 2020.
- [2] M. Lombardi, “An Evaluation of Dependencies of Critical Infrastructure Timing Systems on the Global Positioning System (GPS),” *NIST Technical Note 1289*, 57 p., October 2021 (original report submitted to DHS Science and Technology (S&T) Directorate, February 28, 2020).
- [3] U. S. Code, Title 15, Chapter 6, Subchapter 9, Section 261 (<https://www.law.cornell.edu/uscode/text/15/261>).
- [4] Edited by Bureau International des Poids et Mesures (BIPM), “Proceedings of the 26th meeting of the General Conference on Weights and Measures (November 2018),” p. 491 (English version), December 2019. The current SI definition of the second was first published as *Resolution 1 of the 13th Conference Generale des Poids et Mesures (CGPM)*, 1967.
- [5] F. Arias, Z. Jiang, W. Lewandowski, and G. Petit, “BIPM comparison of time transfer techniques,” *Proceedings of the 2005 IEEE International Frequency Control Symposium and Exposition*, 4 p., August 2005.
- [6] BIPM, Circular-T archive: <https://www.bipm.org/en/bipm-services/timescales/time-ftp/Circular-T.html>
- [7] T. Heavner, E. Donley, F. Levi, G. Costanzo, T. Parker, J. Shirley, N. Ashby, S. Barlow, and S. Jefferts, “First accuracy evaluation of NIST-F2,” *Metrologia*, vol. 51, no. 3, pp. 174-182, 2014.
- [8] W. Riley, “Handbook of Frequency Stability Analysis,” *NIST Special Publication 1065*, 136 p., July 2008. Also see the IEEE standard: “Standard definitions of physical quantities for fundamental frequency and time metrology—Random instabilities,” *IEEE Standard 1139*, 2008.
- [9] S. Stein, D. Glaze, J. Levine, J. Gray, D. Hilliard, D. Howe, and L. Erb, “Automated high-accuracy phase measurement system,” *IEEE Transactions on Instrumentation and Measurement*, vol. 31, no. 1, pp. 227–231, 1983.
- [10] J. Levine, “The statistical modeling of atomic clocks and the design of time scales”, *Review of Scientific Instruments*, vol. 83, pp. 021101/1-28, 2012.
- [11] See, for example, the text file maintained at: <https://www.ietf.org/timezones/data/leap-seconds.list>
- [12] Unpublished; observed with a GuideTech GT668PCIe-1 measuring UTC(NIST) PPS signals, valid for averaging intervals through ~100 s. Even among more moderately-priced modern TICs (e.g. the Pendulum CNT-91), we have observed noise performance $\sigma_y(\tau) \leq 25$ ps ($\tau/1s$), which, when given 60 s of averaging, is adequate as a microwave clock measurement instrument.
- [13] J. Sherman and R. Jördens, “Oscillator metrology with software defined radio,” *Review of Scientific Instruments*, vol. 87, pp. 054711/1-11, 2016.
- [14] J. Valls, T. Sansaloni, A. Perez-Pascual, V. Torres, and V. Almenar, “The use of CORDIC in software defined radios: a tutorial,” *IEEE Communications Magazine*, vol. 44, no. 9, pp. 46–50, 2006.
- [15] J. Savory, J. Sherman, and S. Romisch, “White Rabbit-Based Time Distribution at NIST,” *Proceedings of the 2018 IEEE International Frequency Control Symposium*, 5 p., May 2018.
- [16] E. Arias, G. Panfilo, and G. Petit, “Timescales at the BIPM,” *Metrologia*, vol. 48, no. 4, S145, 2011.

- [17] W. McGrew, X. Zhang, H. Leopardi, R. Fasano, D. Nicolodi, K. Beloy, J. Yao, J. Sherman, S. Schäffer, J. Savory, R. Brown, S. Römisch, C. Oates, T. Parker, T. Fortier, and A. Ludlow, “Towards the optical second: verifying optical clocks at the SI limit,” *Optica*, vol. 6, no. 4, pp. 448–454, 2019.
- [18] R. Nelson, D. McCarthy, S. Malys, J. Levine, B. Guinot, H. Fliegel, R. Beard, and T. Bartholomew, “The leap second: its history and possible future,” *Metrologia*, vol. 38, no. 6, pp. 509-529, 2001.
- [19] IERS Bulletin C, published electronically: <https://www.iers.org/SharedDocs/News/EN/BulletinC.html>
- [20] J. Levine, “Realizing UTC(NIST) at a Remote Location,” *Metrologia*, vol. 45, pp. S23-S33, December 2008.
- [21] NIST Service ID 78100S, “Time over Fiber Special Test (Gaithersburg),” (<https://shop.nist.gov>).
- [22] NIST Service ID 78110S, “Time over Fiber Special Test (Boulder),” (<https://shop.nist.gov>).
- [23] “Optical power loss (attenuation) in fiber access – type, value, and sources,” *Ad-Net Technology web site* (<https://www.ad-net.com.tw/optical-power-loss-attenuation-in-fiber-access-types-values-and-sources/>), December 28, 2008.
- [24] OFS optical fiber catalog (www.ofsoptics.com), accessed May 2020.
- [25] IEEE, “Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems,” *IEEE Standard 1588-2019*, November 2019.
- [26] J. Eidson, *Measurement, Control, and Communication Using IEEE 1588*, Springer, 2006.
- [27] Software and documentation about NTP is available from the Network Time Foundation (www.ntp.org).
- [28] D. Mills, *Computer Network Time Synchronization: The Network Time Protocol*, CRC Press, 2006.
- [29] J. Serrano, M. Cattin, E. Gousiou, E. van der Bij, T. Wlostowski, G. Daniluk, and M. Lipinski, “The White Rabbit Project,” *Proceedings of the 2nd International Beam Instrumentation Conference*, Oxford, UK, pp. 936-942, September 2013.
- [30] E. Dierikx, A. Wallin, T. Fordell, J. Myyry, P. Koponen, M. Merimaa, T. Pinkert, J. Koelemeij, H. Peek, and R. Smets, “White Rabbit Precision Time Protocol on Long-Distance Fiber Links,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 63, no. 7, pp. 945-952, July 2016.
- [31] Civil Laser web site, (<https://www.civillaser.com>), accessed May 2019.
- [32] J. Levine, “The Statistics of Computer Clocks and the Design of Synchronization Algorithms,” *Journal of Research of the National Institute of Standards and Technology*, vol. 125, pp. 125008/1-33, 2020.
- [33] J. Barnes and D. Allan, “Effects of Long-Term Stability on the Definition and Measurement of Short-Term Stability,” *Proceedings of IEEE-NASA Symposium on the Definition and Measurement of Short-Term Frequency Stability*, pp. 119-123, November 1964.
- [34] D. Allan and J. Barnes, “A Modified ‘Allan Variance’ with Increased Oscillator Characterization Ability,” *Proceedings of the 35th Annual Frequency Control Symposium*, pp. 470-475, May 1981.
- [35] W. Diener and G. Lutes, “Thermal coefficients of delay for various coaxial and fiberoptic cables,” *The Telecommunications and Data Acquisition (TDA) Progress Report 42-99*. (NASA, Jet Propulsion Laboratory, Pasadena, CA), pp. 43-59, November 1989.

- [36] P. Merritt, R. Tatam and D. Jackson, “Interferometric chromatic dispersion measurements on short lengths of monomode optical fibers,” *Journal of Lightwave Technology*, vol. 7, no. 4, pp. 703-716, 1989.
- [37] E. Laier-English, P. Whibberley, C. Langham, D. Hicks, and L. Lobo, “Update on the NPLTime® and Future Developments with White Rabbit,” *Proceedings of the 2017 Precise Time and Time Interval (PTTI) Meeting*, pp. 80-87, January 2017.
- [38] OPNT press release, (<https://www.prnewswire.com/news-releases/opnt-achieves-network-timing-synchronization-exceeding-gps--announces-back-up-to-gps-and-launches-timing-as-a-service-300861470.html>), June 2019.
- [39] D. Goward, “GPS backup demonstration projects explained,” *GPS World*, January 9, 2020.
- [40] D. Kirchner, “Two-Way Satellite Time and Frequency Transfer (TWSTFT): Principle, Implementation, and Current Performance,” *Review of Radio Science 1996-1999*, Oxford University Press, New York, pp. 27-44, 1999.
- [41] ITU Radio Communication Sector, “The operational use of two-way satellite time and frequency transfer employing PN codes,” *Recommendation ITU-R TF.1153-4*, 2015.
- [42] BIPM, *Circular T 389*, <ftp://ftp2.bipm.org/pub/tai/Circular-T/cirhtml/cirt.389.html>, May 2020.
- [43] Z. Jiang, V. Zhang, Y. Huang, J. Achkar, D. Piester, S. Lin, W. Wu, A. Naumov, S. Yang, J. Nawrocki, I. Sesia, C. Schlunegger, Z. Yang, M. Fujieda, A. Czubla, H. Esteban, C. Rieck, and P. Whibberley, “Use of software-defined radio receivers in two-way satellite time and frequency transfer for UTC computation”, *Metrologia*, vol. 55, no. 5, pp. 685–98, August 2018.
- [44] M. Fujieda, D. Piester, T. Gotoh, J. Becker, M. Aida, and A. Bauch, “Carrier-phase two-way satellite frequency transfer over a very long baseline,” *Metrologia*, vol. 51, no. 3, pp. 253-262, May 2014.
- [45] BIPM, “TWSTFT Calibration Guidelines for UTC Time Links,” Version 3.0, <ftp://ftp2.bipm.org/pub/tai/publication/twstft-calibration/guidelines/>, 2016.
- [46] M. Lombardi, “Disciplined clock for providing a disciplined time and a disciplined frequency synchronous with a reference clock,” *United States Patent 10,466,363 B2*, granted November 5, 2019.
- [47] M. Lombardi, “The Reach and Impact of the Remote Frequency and Time Calibration Services at NIST,” *NCSLI Measure: The Journal of Measurement Science*, vol. 12, pp. 30-37, June 2019.
- [48] M. Lombardi, A. Novick, B. Cooke, and G. Neville-Neil, “Accurate, Traceable, and Verifiable Time Synchronization for World Financial Markets,” *Journal of Research of the National Institute of Standards and Technology*, vol. 121, pp. 436–463, 2016.
- [49] M. Lombardi and A. Dahlen, “A common-view disciplined oscillator,” *Review of Scientific Instruments*, vol. 81, no. 5, pp. 055110/1-6, 2010.
- [50] M. Lombardi, “A NIST Disciplined Oscillator: Delivering UTC(NIST) to the Calibration Laboratory,” *NCSLI Measure: The Journal of Measurement Science*, vol. 5, pp. 46-54, December 2010.
- [51] J. Bernard, A. Charbonneau, B. Hoger, H. Pham, and M. Gertsvolf, “NRC Remote Clock – A Secure and Traceable Time Source,” *Proceedings of the 2017 Precise Time and Time Interval (PTTI) Meeting*, pp. 73-79, January 2017.

- [52] L. Kun, Y. Hang, Z. Fei, Z. Side and Z. Yuan, "Disciplined oscillator system by UTC(NIM) for remote time and frequency traceability," *Proceedings of the 2014 European Frequency and Time Forum (EFTF)*, pp. 468-471, 2014.
- [53] R. Chen, Y. Liu, Y. Yang, and X. Li, "Research on the Method of Remote Reproduction of UTC(NTSC) and its Realization," *Proceedings of the 2018 IEEE Advance Information Technology, Electronic and Automation Control Conference (IAEAC)*, pp. 118-122, October 2018.
- [54] N. Ozdemir, P. Defraigne, G. Ceretto, E. Cantoni, and P. Tavella, "GNSS Disciplined Oscillators: an approach based on real-time steering over the Internet with certification," *Proceedings of the 2016 European Frequency and Time Forum (EFTF)*, 1 p., April 2016.
- [55] M. Imae, Y. Fujii, Y. Mitamoto, T. Suzuyama, T. Kawakami, H. Yoshida, H. Hurukawa, and M. Susumu, "Development of a GPS Common-view Terminal for Time and Frequency Remote Calibration Service," *Proceedings of 2010 International Symposium on GPS/GNSS*, 5 p., October 2010.
- [56] A. Mitchell, "Frequency Comparison of Atomic Standards by Radio Links," *Nature*, vol. 198, pp. 1155-1158, June 22, 1963.
- [57] J. Tolman, V. Ptacek, A. Soucek, and R. Stecher, "Microsecond Clock Comparison by Means of TV Synchronizing Pulses," *IEEE Transactions on Instrumentation and Measurement, IM-16*, no. 3, pp. 247-254, September 1967.
- [58] D. Allan and M. Weiss, "Accurate Time and Frequency Transfer During Common-View of a GPS Satellite," *Proceedings of the 1980 Frequency Control Symposium*, pp. 334-346, May 1980.
- [59] Note that the only two GNSS systems legally usable in the U. S., without a license from the Federal Communications Commission's (FCC) International Bureau) are GPS and Galileo (see <https://www.fcc.gov/document/fcc-approves-galileo-global-navigation-system-0> for the November 2018 approval of Galileo). However, it is not clear if this precludes the use of GLONASS and BeiDou as common-view signal sources, nor is it clear if it applies only to U. S. government installations and not to U. S. industry, as multi-constellation GNSS receivers are now embedded into nearly every smartphone.
- [60] A. Hashemi, S. Thombre, N. Ferrera, M. Bhuiyan, and M. Pattinson, "STRIKE3 – Case Study for Standardized Testing of Timing-grade GNSS Receivers against Real-world Interference Threats," *Proceedings of 2019 International Conference on Localization and GNSS (ICL GNSS)*, 8 p., July 2019.
- [61] N. Ferrera, M. Bhuiyan, S. Soderholm, L. Ruotsalainen, and H. Kuusniemi, "A new implementation of narrowband interference detection, characterization, and mitigation technique for a software-defined multi-GNSS receiver," *GPS Solutions*, vol. 22, article 106, pp. 1-16, August 2018.
- [62] H. Elghamrawy, M. Karaim, M. Tamazin, and A. Noureldin, "Experimental Evaluation of the Impact of Different Types of Jamming Signals on Commercial GNSS Receivers," *Applied Sciences*, vol. 10, pp. 4240/1-15, June 2020.
- [63] L. Scott, "Spoofing Incident Report: An Illustration of Cascading Security Failure," *Inside GNSS*, October 9, 2017.
- [64] M. Psiaki and T. Humphreys, "GNSS Spoofing and Detection," *Proceedings of the IEEE*, vol. 104, no. 6, pp. 1258-1270, June 2016.
- [65] K. Zhang and P. Papadimitratos, "Secure Multi-Constellation GNSS Receivers with Clustering-Based Solution Separation Algorithm," *Proceedings of 2019 IEEE Aerospace Conference*, 9 p., March 2019.
- [66] P. Banerjee, M. Saxena, and B. Mathur, "A study on the accuracy of time comparison via geostationary satellite in a common view mode," *Indian Journal of Radio & Space Physics*, vol. 24, pp. 333-339, December 1995.

- [67] A. Sen Gupta and B. Mathur, "Standard Time and Frequency Broadcast Via INSAT – Accuracy Improvements Using Differential Mode," *IEEE Transactions on Instrumentation and Measurement*, vol. 46, no. 3, pp. 212-215, April 1997.
- [68] D. Hanson and D. Howe, "Industrial Time Service Study," *NBSIR 86-3042*, 189 p., February 1986.
- [69] J. Yoon, K. Lee, B. Lee, B. Kim, K. Choi, Y. Chang, Y. Chun, and S. Ra, "Geostationary orbit determination for time synchronization using analytical dynamic models," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 40, no. 4, pp. 1132-1146, October 2004.
- [70] J. Walker and M. Genova, "Experimental and Simulation Study for a Time Transfer Service via a Commercial Geostationary Satellite," *NCSLI Measure: The Journal of Measurement Science*, vol. 8, no. 2, pp. 62-69, June 2013.
- [71] J. Davis and P. Pearce, "Time Dissemination in the United Kingdom Using Signals from Geostationary Television Satellites," *Proceedings of the 7th British Electromagnetic Measurements Conference (BEMC)*, paper 19, pp. 1-4, November 1995.
- [72] F. Meyer, "One-Way Time Transfer Using Geostationary Satellite TDF2," *IEEE Transactions on Instrumentation and Measurement*, vol. 44, no. 2, pp. 103-106, April 1995.
- [73] X. Liu, Y. Hua, Y. Xiang, and H. Liu, "Method of common view based on digital satellite television," *Applied Mechanics and Materials*, vol. 721, pp. 666-669, 2015.
- [74] M. Lombardi and D. Hanson, "The GOES Time Code Service, 1974-2004: A Retrospective," *Journal of Research of the National Institute of Standards and Technology*, vol. 110, no. 2, pp. 79-96, March-April 2005.
- [75] S. Leschiutta, "The definition of the 'atomic' second," *Metrologia*, vol. 42, pp. S10-S19, June 2005.
- [76] Bureau International de L'Heure (BIH), *BIH Annual Report for 1981*, Table 13, p. B-24, 1982.
- [77] See index of Bureau International de L'Heure (BIH) Annual Reports for lists of time transfer links used before 1981. The index is available at: <ftp://ftp2.bipm.org/pub/tai/annual-reports/bih-annual-report/>
- [78] T. Celano, K. Carroll, C. Biggs, and M. Lombardi, "Common-View Loran-C as a Backup to GPS for Precise Time Recovery," *Proceedings for the 2003 Precise Time and Time Interval (PTTI) Meeting*, pp. 81-92, December 2003.
- [79] United States Code of Federal Regulations, "WWV-WWVH-WWVB Broadcasts," *15 CFR § 200.107*, 2019.
- [80] G. Nelson, M. Lombardi, and D. Okayama, NIST "Time and Frequency Radio Stations: WWV, WWVH, and WWVB," *National Institute of Standards and Technology Special Publication 250-67*, January 2005.
- [81] M. Lombardi and G. Nelson, "WWVB: A Half Century of Delivering Accurate Frequency and Time by Radio," *Journal of Research of the National Institute of Standards and Technology*, vol. 119, pp. 25-54, March 2014.
- [82] P. Hansen and D. Gish, "WWVB Antenna and Antenna Tuning System: Baseline Measurements," Naval Command, Control and Ocean Surveillance Center, RDT & E Division, *Technical Report 1693*, February 1995.
- [83] P. Hansen and J. Birkett, "WWVB Year 2000: Post Upgrade Baseline," Space and Naval Warfare Systems Center (SSC San Diego), *Technical Report 1833*, December 2003.
- [84] P. Hansen and D. Gish, "WWVB Post Repair: Antenna Performance Measurement," Space and Naval Warfare Systems Center (SSC San Diego), *Technical Report 1834*, December 2003.
- [85] J. Lowe and K. Allen, "Increasing the Modulation Depth of the WWVB Time Code to Improve the Performance

- of Radio Controlled Clocks,” *Proceedings of the 2006 IEEE Frequency Control Symposium*, pp. 615-621, 2006.
- [86] Y. Liang, O. Eliezer, D. Rajan, and J. Lowe, “WWVB Time Signal Broadcast: An Enhanced Broadcast Format and Multi-Mode Receiver,” *IEEE Communications Magazine*, vol. 52, pp. 210-217, May 2014.
- [87] K. Croissant, S. Ugazio, and F. van Graas, “Feasibility of Microsecond-Level Timing with a WWVB disciplined oscillator,” *Proceedings of 2018 Precise Time and Time Interval (PTTI) Meeting*, Reston, Virginia, January 2018.
- [88] W. Walsh, “Loran as an Independent Source of UTC in Telecommunications Applications,” *Proceedings of the 2002 International LORAN Association (ILA) and Technical Symposium*, October 2002.
- [89] M. Lombardi, A. Novick, J. Lowe, M. Deutch, G. Nelson, D. Sutton, W. Yates, and D. Hanson, “WWVB Radio Controlled Clocks: Recommended Practices for Manufacturers and Consumers (2009 edition),” *NIST Special Publication 960-14*, September 2009.
- [90] M. A. Lombardi, “How accurate is a radio controlled clock?,” *The Horological Journal*, vol. 152, no. 3, pp. 108-111, March 2010.
- [91] T. Osterdock, “Timekeeping and Frequency Calibration,” *Hewlett Packard Application Note 52-2*, November 1975.
- [92] R. Hesselberth, T. Danaher, and J. Sandahl, “Disciplined Oscillator System with Frequency Control and Accumulated Time Control,” *United States Patent 4,525,685*, June 25, 1985.
- [93] IEEE, “IEEE Standard for a Precision Clock Synchronization for Networked Measurement and Control Systems,” *IEEE Standard 1588-2019*, 2019.
- [94] International Telecommunication Union (ITU), “Precision time protocol telecom profile for frequency synchronization,” *ITU-T G.8265.1/Y.1365.1*, 2014.
- [95] International Telecommunication Union (ITU), “Packet delay variation network limits applicable to packet-based methods (frequency synchronization),” *ITU-T G.8261.1/Y1361.1*, 2012.
- [96] S. Rinaldi, P. Ferrari, A. Flammini, D. Fontanelli and D. Macii, "Improving syntonization of synchronous ethernet nodes using multiple paths," *Proceedings of the 2014 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS)*, pp. 1-6, 2014.
- [97] A. Novick, M. Weiss, K. Lee, and D. Sutton, “Examination of Time and Frequency Control across Wide Area Networks using IEEE1588v2 Unicast Transmissions,” *Proceedings of the 2011 Joint Meeting of the IEEE International Frequency Control Symposium and the European Frequency and Time Forum (EFTF)*, pp 670-675, 2011.
- [98] Satelles White Paper, <https://www.satellesinc.com/wp-content/uploads/2016/06/Satelles-White-Paper-Final.pdf>, 2019.
- [99] T. Reid, A. Neish, T. Walter, and P. Enge, “Leveraging Commercial Broadband LEO Constellations for Navigation,” *Proceedings of the 29th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2016)*, 15 p., September 2016.

- [100] G. Gutt, D. Lawrence, S. Cobb, M. O'Connor, "Recent PNT Improvements and Test Results Based on Low Earth Orbit Satellites," *Proceedings of the 2018 Precise Time and Time Interval (PTTI) Meeting*, pp. 72-79, January 2018.
- [101] Department of Homeland Security – Coast Guard, "Terminate Long Range Aids to Navigation (Loran-C) Signal," *United States Federal Register*, vol. 75, no. 4, p. 998, January 7, 2010.
- [102] S. Lo, B. Peterson, P. Enge, and P. Swaszek, "Loran Data Modulation: Extensions and Examples," *IEEE Transactions on Aerospace and Electronics Systems*, vol. 43, no. 2, pp. 628-644, April 2007.
- [103] Public Law 115-282, Section 514, *National Timing and Resilience and Security Act of 2018*, 49 USC § 312, "Alternative Timing System," December 4, 2018.
- [104] G. Offermans, S. Bartlett, and C. Schue, "Providing a Resilient Timing and UTC Service Using eLoran in the United States," *Proceedings of the 2016 Precise Time and Time Interval (PTTI) Meeting*, 11 p., January 2016.
- [105] L. Charron, "Relationships between U. S. Naval Observatory, Loran-C and the Defense Satellite Communications System," *Proceedings of the 1981 Precise Time and Time Interval (PTTI) Meeting*, pp. 201-26, December 1981.
- [106] A partial archive of *NIST Time and Frequency Bulletins* is maintained at: <https://www.nist.gov/pml/time-and-frequency-division/services/time-scale>
- [107] K. Montgomery and M. Lombardi, "PTTI Capabilities of the Modernized Loran System," *Proceedings of the 2008 Precise Time and Time Interval (PTTI) Meeting*, pp. 507-526, December 2008.
- [108] K. Shmiluk, "LSU Timing Update (State of LDC and TWSTFT)," *Proceedings of the 2006 International Loran Association (ILA) Meeting*, October 2006.
- [109] USCG, "Draft version of eLoran Broadcast Standard," *COMDTINST M165XX.X*, unpublished, March 2008.
- [110] SAE International, "Transmitted Enhanced Loran (eLoran) Signal Standard for 8th Pulse Modulation," *SAE 9990/2*, 2018.
- [111] A. Helwig, G. Offermans, C. Stout, and C. Schue, "eLoran System Definition and Signal Specification Tutorial," *Proceedings of the 2011 International Loran Association (ILA) Meeting*, November 2011.
- [112] J. Yao, M. Lombardi, A. Novick, B. Patla, J. Sherman, and V. Zhang, "The effects of the January 2016 UTC offset anomaly on GPS-controlled clocks monitored at NIST," *Proceedings of the 2017 Precise Time and Time Interval (PTTI) Meeting*, pp. 155-163, January 2017.
- [113] J. Sherman and J. Levine, "Usage Analysis of the NIST Internet Time Service," *Journal of Research of the National Institute of Standards and Technology*, vol. 121, pp. 33-46, March 2016.
- [114] The current status of the NIST Internet Time Service is displayed at: <https://tf.nist.gov/tf-cgi/servers.cgi>
- [115] "NIST Authenticated NTP Service," (<https://www.nist.gov/pml/time-and-frequency-division/time-services/nist-authenticated-ntp-service>)