Synchronization between Remote Sites for the MINOS Experiment

S. Römisch[1], S. R.Jefferts[1], V. Zhang[1], T. E. Parker[1], N. Ashby[1], P. Adamson[2], G. Barr[3], A. Habig[4], J. Meier[4], C. James[2], R. Nicol[5], R. Plunkett[2], C. Rosenfeld[6], R. Bumgarner[7], M. Christensen[7], J. Hirschauer[7], B. Fonville[7], S. Mitchell[7], A. McKinley[7], E. Powers[7], J. Wright[7], and D. Matsakis[7]

[1] NIST – Time and Frequency Division, Boulder, CO – USA

[2]Fermilab, Batavia, IL – USA

[3]University of Oxford -UK

[4]University of Minnesota, Duluth, MN - USA

[5] University College, London – UK

[6]University of South Carolina, SC – USA

[7]USNO, Washington, DC – USA

romisch@boulder.nist.gov

Abstract—In the context of time-of-flight measurements, the timing at the departure and arrival locations is obviously critical to the outcome of the experiment. In the case of neutrino time-of-flight experiments, the locations are many hundreds of kilometers apart with synchronization requirements of nanoseconds for several months at a time. In addition to the already stringent set of requirements outlined above, the locations of the origin of the particle beam and the detector need to be precisely determined. NIST and USNO have provided the MINOS (Main Injector Neutrino Oscillation Search) collaboration with both hardware and expertise to synchronize the two sites of the experiment, the accelerator at Fermilab in Batavia, IL and the Soudan Mine in northern Minnesota. Two GPS receivers are installed at each location where the local clocks are commercial Cesium clocks. Two more GPS receivers are constantly traveling between locations (including NIST in Boulder, CO) to provide multiple differential calibrations of the fixed receivers. The availability of the TWTFST equipment from USNO allowed for one comparison between the GPS and TWSTFT for the link between the locations, providing an independent means of determining the accuracy of the synchronization. Several months of continuous GPS data are now available, including the two-way calibration instance and several differential GPS calibrations. The results of data processing yielded synchronization stability below one nanosecond with accuracy at the nanosecond level over several months.

I. BACKGROUND

In the fall of 2011 the unexpected results of the OPERA collaboration's experiment were announced [1], prompting the scientific community to try to reproduce the detection of superluminal neutrinos. Members of the MINOS (Main Injector Neutrino Oscillation Search) collaboration at Fermilab in Batavia, IL, contacted NIST and USNO to discuss the timing setup for their facilities.

After an initial informal meeting, a workshop and a visit to the facilities in Batavia, IL a fruitful collaboration was initiated, continuing to this day.

The MINOS Time-Of-Flight (TOF) experiment involves three locations:

- MI60: it is one of the two locations in Batavia, IL, and it is the closest to the origin of the particle beam that eventually generates the neutrinos whose speed is measured. More importantly, it houses the resistive wall current monitor that is triggered by each "bunch" of particles that transits towards the detectors.
- Near Detector (ND): it is the second location in Batavia, IL. It is a copy of the detector at the Soudan Mine and it is situated approximately 1.2 km from MI60 with the purpose of allowing a calibration of the latency times in the Far Detector.
- Far Detector (FD): it is situated approximately 700m underground in the historical iron mine called Soudan Mine, in Minnesota, near the Canadian border. It is the arrival point for the beam of neutrinos and it is situated at a distance of approximately 735km. A full discussion of the geodetic implications of determining the actual distance travelled by the neutrinos can be found in [2].

In Figure 1 is shown a simplified representation of the experiment setup, with the indication of the three locations: MI60, ND and FD.

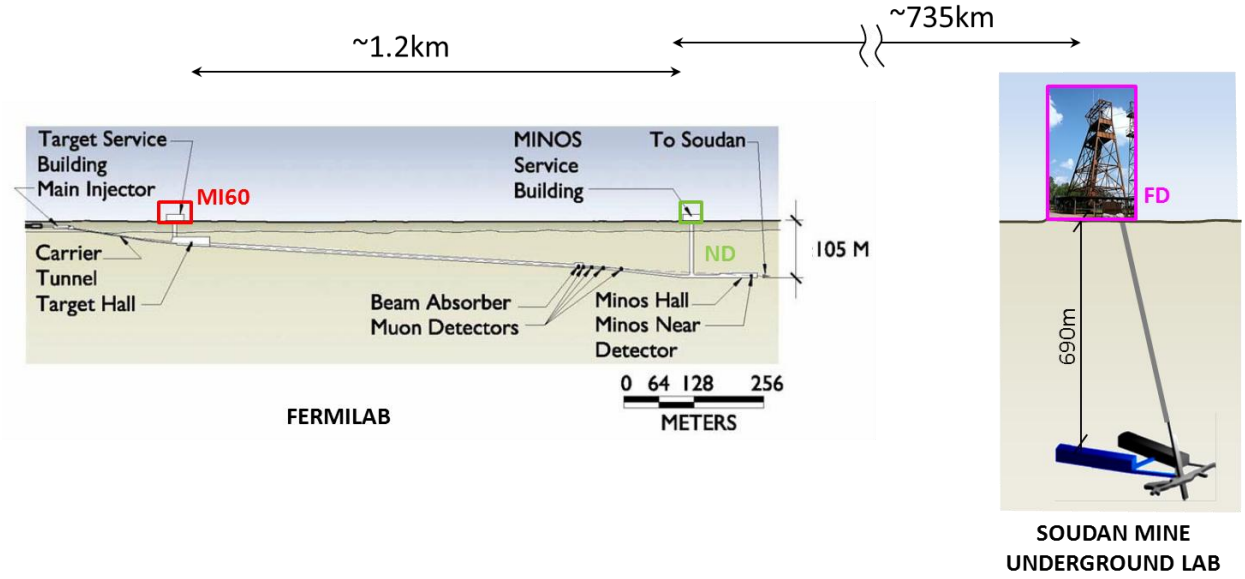

Figure 1. Simplified representation of all locations involved in the MINOS Time-Of-Flight experiment.

This paper limits its scope to the description of the synchronization setup serving the experiment and the discussion of its results. A more complete description of the Time-Of-Flight (TOF) experiment and of the timing systems at each of the three locations can be found in [4], [4], [5].

A. Timing requirements

The particle beam generated at the Fermilab facilities has a complicated structure, but for the purpose of this discussion it is enough to say that is it made of "bunches" of particles, separated by 18.83 ns and amounting to a full-width of approximately 3.5ns, as shown in Figure 2 below where 162 bunches were overlaid to generate it. The profile of the "bunches" is detected using a current monitoring device implemented around the particle beam at the MI60 location.

The width of each "bunch" is indicative of the right order of magnitude for the timing requirement of the synchronization between the three locations; moreover, the Far Detector's present position uncertainty is approximately 70 cm, amounting to an equivalent timing uncertainty around 2 ns.

Given these considerations, a synchronization goal of about 1 nanosecond, to be maintained over the duration of the experiment (months to years) is reasonable. Synchronization means knowledge of the relative time difference between the time references at two or more locations, and the nanoseconds-level requirement refers both to stability and accuracy. Throughout this paper all uncertainties are listed as 1- σ uncertainties (68% confidence) although for the purpose of identifying neutrino events, the users are most interested in 3-σ uncertainties (99.5% confidence).

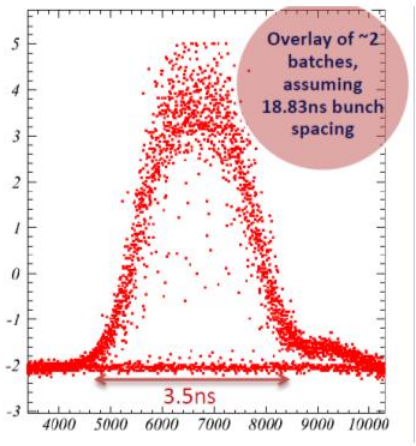

Figure 2. Overlay of approximately 162 "bunches" of particles, showing the full-width of a "bunch" to be approximately 3.5 ns. A "batch" is a sequence of 81 "bunches". The units on both axis are arbitrary. The profile of the "bunches" is detected using a current monitoring device implemented in the wall around the particle beam.

Finally, it is important to remark that these requirements don't refer to time-of-day accuracy, which is only needed at a much coarser level.

B. Synchronization setup

There are at least four possible ways to provide synchronization between clocks at two remote locations:

- GPS link, with either C/A (code) or carrier-phase processing;
- Two-Way Satellite Time and Frequency Transfer (TWSTFT);
- Repeated shipments of a stable clock between locations (clock trips);
- Fiber-based Two-Way Time Transfer (TWTT).

The simultaneous use of more than one synchronization link allows a direct comparison of multiple measurements of the time difference between two locations. In particular, the presence of several methods deemed to be independent allows for the calculation of the differences between the time measurements resulting from each method. With three independent techniques, a "three-cornered hat" approach is possible, allowing the estimation of the uncertainty of each link separately. The three uncertainties could then be treated as statistically independent and appropriately combined to obtain the final uncertainty of the synchronization between two locations.

It was not possible to implement the optimal scenario described above, but a reduced version of it, as shown in Figure 3. The "clock trip" option is still open and will likely be implemented in 2013.

The GPS link was easily deployable with 2 fixed receivers for each of the three locations of the MINOS experiment, and two travelling units to perform periodic calibrations (the travelling units are drawn using dashed lines). The TWSTFT link has been deployed only between the ND and FD (longer baseline) for a single instance lasting a few days, and it was used as a tool to verify the accuracy of the GPS link with a single-point comparison in April 2012.

The fiber-based TWTT link is currently in place between the two locations in Illinois (MI60 and ND, shorter baseline) and will be also used as a comparison tool for the GPS link between the same locations. All fiber-based two-way time transfer systems are implemented with two unidirectional fibers within the same bundle.

USNO and NIST are also shown in Figure 3 because they were used to supply supplementary information for, respectively, the TWSTFT link and the GPS link. In the case of TWSTFT, the USNO terminal provides a supplementary semi-independent indirect means to calibrate the direct link between ND and FD. As for the GPS link, the periodic presence at NIST of a travelling receiver constitutes additional redundancy in the calibration process and allows monitoring of the clocks at the various MINOS locations against the NIST maser ensemble. The data comparing the MINOS clocks with the NIST time scale is not included in this paper.

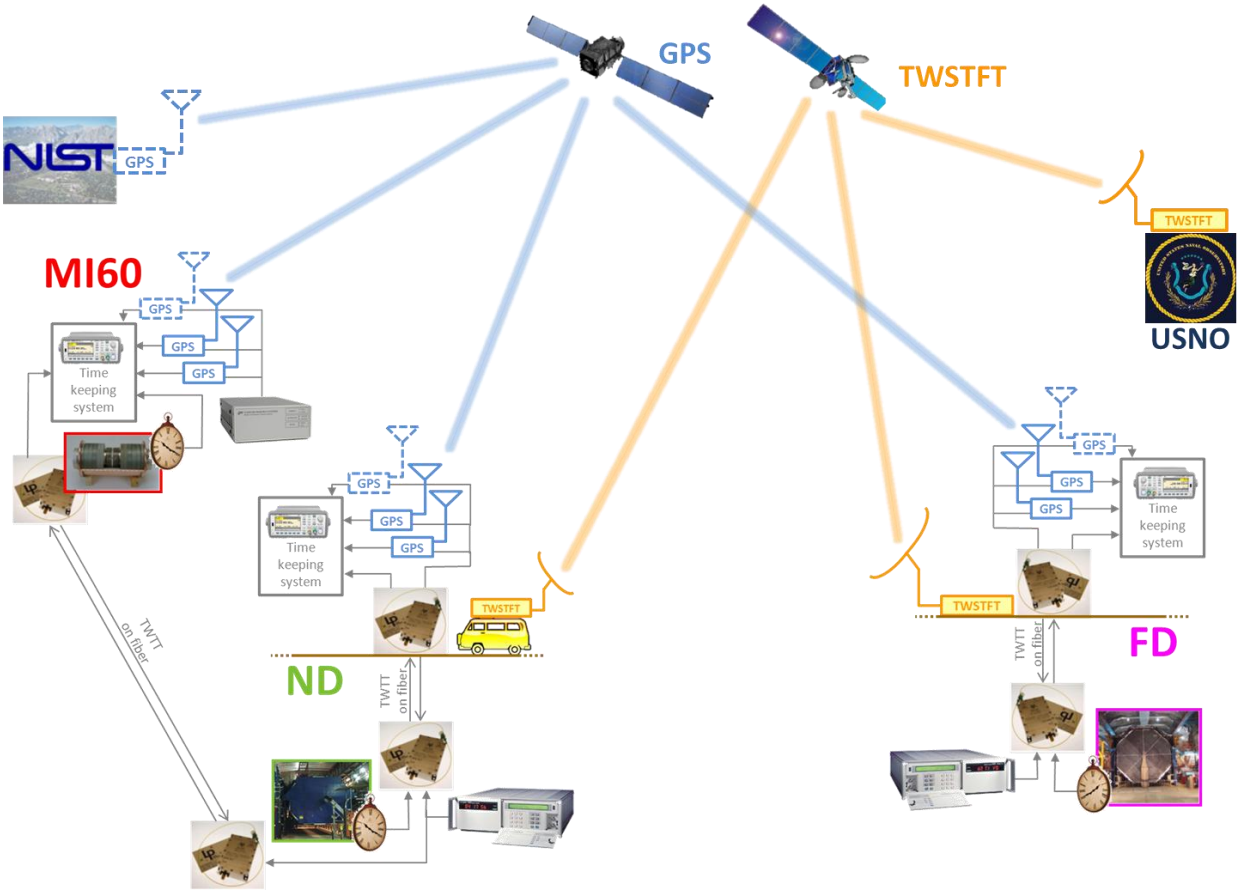

Figure 3. Complete synchronization setup, including the TWSTFT and the fiber-based TWTT. USNO and NIST are also depicted as calibration locations for, respectively, the TWSTFT link and the GPS link.

At the ND and FD facilities the GPS receivers are on the surface, while the detectors and the reference clocks are underground (approximately 100 m at ND and 690 m at FD): in both cases there is a fiber-based two-way time transfer system to relate the underground timing signals to the surface.

The data relative to these two-way connections are not shown in this paper, whose scope is to describe and evaluate the synchronization among the local time reference planes (located on the surface) at the MI60, ND and FD.

The local reference clocks are standard performance commercial cesium Symmetricom 5071A[2] (courtesy of USNO) at the ND and FD facilities, while at MI60 there is a commercial rubidium FS725 from SRS[2].

In order to have meaningful comparisons of time in different places, a time reference plane is defined at each location of interest. A time reference plane, sometimes also called time reference point, is a specific physical place where a timing signal (pulse-per-second) is present: all the existing timing signals at that location are defined in term of delay with respect to that place.

In the specific case of the MINOS experiment, a time reference plane was defined at each location (MI60, ND and FD) as well as at NIST and USNO. The term *time reference plane* will be used throughout this paper according to the above definition: all time differences are the difference between timing signals referred to any spot in the local time reference planes, ignoring of course extremely minor relativistic considerations.

II. THE GPS LINK

We deployed eight modern, multi-frequency GNSS receivers, of which six are virtually identical (Novatel OEMV)[2] and two are a newer version of the same kind of device (Novatel OEM6)[2] with a 100-MHz internal oscillator in place of the 20-MHz one used in the OEMV. They are named GPSn, where n=1..8, including the travelling units used for periodic differential calibrations.

Figure 4 shows a schematic representation of the location of each GPS unit, depicting also the two travelling units GPS5 and GPS6 with the approximate relative distance of their location with respect to the fixed units.

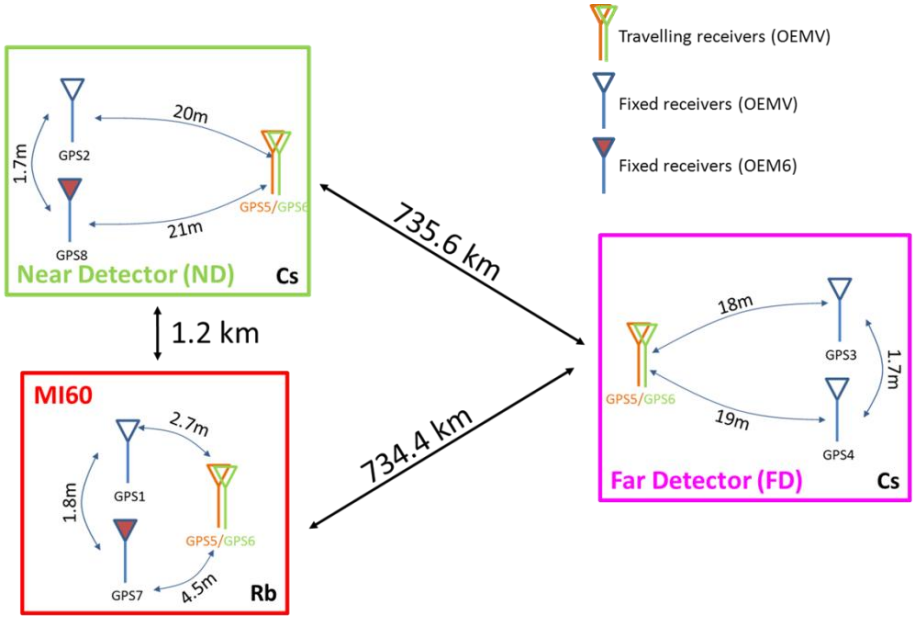

Figure 4. Schematic representation of the baselines of the GPS units deployed for the MINOS TOF experiment.

All antenna positions were obtained using the online service CGSR-PPP (Canadian Spatial Reference System - Precise Point Positioning) offered by the Canadian Geodetic Service of Natural Resources

[2] This information is provided for technical completeness. As a matter of policy, neither the authors nor their institutions can endorse any commercial product or make generalized evaluations of product performance.

Canada [6]. Each 24-hour data segment produces a set of spatial coordinates for the receiver's antenna. The day-to-day variations in the coordinate values for all locations produce timing errors smaller than 50 ps: the positions were therefore generated using the first few days of data, and assumed fixed thereafter. In the case of the travelling receivers, the position of each antenna is re-determined at each calibration instance.

Each receiver is embedded in an instrument, designed and built at NIST that (among other things) dynamically records the time difference between the local time reference and the receiver's time base, thereby increasing the confidence of link calibrations. More details regarding the instrument and its advantages are available in [7].

Within the timing community slightly different approaches to differential GPS calibrations are used, so we prefer to risk stating the obvious and refer to the simple schematic shown in Figure 5 to illustrate the meaning of the crucial quantities CABDLY, REFDLY and INTDLY in the computation of the difference between the time reference planes at two locations. In particular, if we consider location A with receiver 1 and location B with receiver 2, we would write the time difference between A and B as:

$$\mathrm{REFplane_A} - \mathrm{REFplane_B} = \mathrm{CV_{1,2}} - \Delta\mathrm{CABDLY_{1,2}} + \Delta\mathrm{REFDLY_{1,2}} - (\mathrm{INTDLY_1}\text{-}\mathrm{INTDLY_2}) \qquad (1)$$

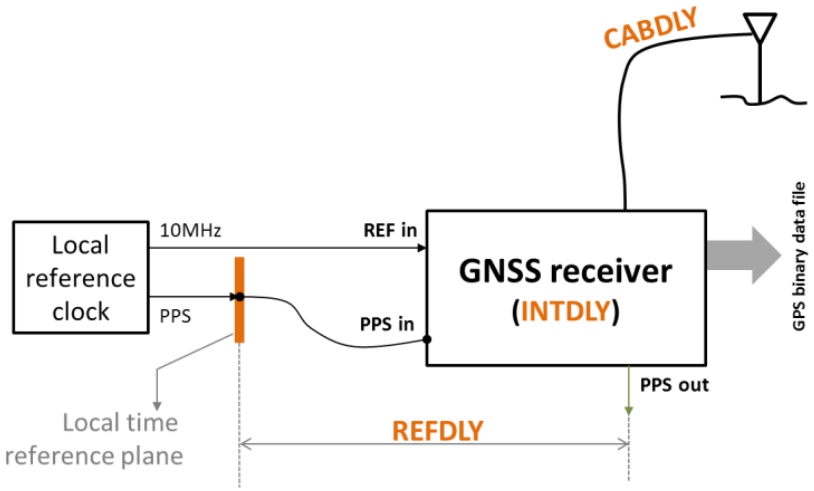

Figure 5. Simple schematic illustrating the definition of CABDLY, REFDLY and INTDLY used throughout this paper.

In Eq. (1), with reference to Figure 5:

$\mathrm{CV_{1,2}}$ represents the common-view difference between the time at receiver 1 and 2, as obtained from the CGGTTS files produced by the receivers.

CABDLY is the time delay associated with the cable used to connect the antenna to the GPS receiver. Prior to the GSP units' installation the group delay of each antenna cable was measured for both L1 (1575.42 MHz) and L2 (1227.60 MHz) frequencies. The results were also confirmed by delay measurements using Time-Domain Reflectometry. It is a single number associated to each antenna cable. $\Delta\mathrm{CABDLY_{1,2}}$ is the difference between the delays of the antenna cables used for receivers 1 and 2.

REFDLY is the delay between the best available representation of the receiver's internal time base and the local time reference plane. It is associated with each receiver in a specific location and it is not a simple number, but is a dynamically measured quantity, with a measurement every 100 s executed by the custom electronics described in [7] and logged on a data file. All GPS receivers used in this experiment are in fact receiver *systems* that include the custom electronics described in [7]. $\Delta\mathrm{REFDLY_{1,2}}$ is the difference between REFDLY for receivers 1 and 2.

INTDLY is a delay that is unique to each receiver and includes possible internal delays that are not part of the previous definitions. It can be determined only through an absolute calibration of the receiver. On the other hand, when using GPS receivers to synchronize remote sites, the goal is to measure the time difference between two locations. This only requires the knowledge of the *difference* between the INTDLY of the pair of receivers constituting the link. For this work, measurements of INTDLY and CABDLY are indistinguishable and fully correlated, however the distinction is maintained to cover the possibility that receivers and antennas can be interchanged.

Using a travelling receiver in turn at each location, the difference between INTDLYs of the local fixed receiver and the travelling one is determined through a common-clock measurement. The reference plane is common to both receivers ($REFplane_A$ - $REFplane_B$ = 0), and the common-view data ($CV_{1,2}$), REFDLY and CABDLY are known quantities. The difference between the INTDLYs is a number obtained by averaging the data collected during each calibration instance.

Each link calibration requires two common-clock calibration instances, one at each location, as shown below, where the subscript T indicates the travelling GPS system.

$$\begin{array}{cc} \text{Location A} & \text{Location B} \\ (INTDLY_1 - INTDLY_T) & (INTDLY_2 - INTDLY_T) \\ & \downarrow \\ & INTDLY_1 - INTDLY_2 \end{array} \qquad (2)$$

A. Evaluation of the GPS link uncertainty

There are several imperfect tools to be used to evaluate the stability and uncertainty of the GPS link between the different locations.

Common-clock measurements are generally performed with co-located receivers and eliminate both the local clock noise and most of the iono/troposphere effects. To the extent that the system and component sensitivities are identical, they also reduce shared environmental effects and possibly systematic multipath effects as well. Because of these cancellations, these measurements mostly evaluate the noise contributions to the link made by the receivers and antenna assemblies, which can be associated with the $CV_{1,2}$ term in Eqn. (1). Differential calibrations with travelling receivers at each location are also common-clock measurements with co-located receivers and antennas, but this time the comparison is done with a travelling unit that will be transported to the other location and act as a transfer device to allow the determination of the difference between INTDLY for the two receivers constituting the link, as shown in Eqn. (2). Each calibration instance is represented by several days of data, of which the mean is the ($INTDLY_1$-$INTDLY_2$) number we are looking for. The stability of each data set is the short-term uncertainty of the calibration, but does not say anything about the stability over months or years of the calibration technique itself. The repetition of several calibration instances at all locations allows the evaluation of the long term behavior of the calibrations themselves, which, although difficult to estimate given the small size of the data set, ultimately constitutes the calibration uncertainty, associated with the term $INTDLY_1$-$INTDLY_2$ in Eqn. (1).

The presence of two synchronization links between the same pair of locations allows direct comparison of multiple measurements of the time difference between two locations. If the difference between the two measurements is larger than allowed by the stated uncertainty for each measurement, the results are deemed metrologically inconsistent for the stated uncertainty. Because in this case simultaneity ensures that the measured quantity hasn't changed between measurements, the overall synchronization uncertainty is appropriately increased so that all results are metrologically consistent. The difference between synchronization links is referred to as "double difference."

The mean value of a double difference between calibrated links will be, in principle, zero at the time of calibration. Its growth over time and the associated total time deviation are both indicators of the links uncertainty. It needs to be kept in mind that, in reality, the receivers at the two locations constituting the link are not calibrated at the same time.

Moreover, because the total time deviation would be insensitive to a constant rate of growth of the mean value, variations in the mean value must be considered in concert with the total time deviation to derive the stated uncertainty associated with each link. In particular, a "large" mean value is likely an indication of systematic effects contributing to the total uncertainty.

While the above statements are always true, in the case when the two synchronization links are of the same kind (i.e. two GPS links) the presence of common-mode systematic effects may reduce the overall estimate of uncertainty. In the case of two GPS synchronization links, common-clock GPS data can provide an equivalent measure to double-differences over the link. Conversely, the double difference between two links using dissimilar synchronization techniques provides the best available estimate of the link uncertainty and variations. In this experiment, the double difference over the shorter baseline can be calculated using the GPS link and the fiber-based TWTT, providing a comprehensive measure of that link accuracy.

The longer baseline doesn't have a permanent alternative to the GPS-base synchronization link. Nonetheless, the point-calibration performed using a TWSTFT system (courtesy of USNO) allows a validation of the GPS link and the quantification of the agreement between the two techniques.

The interplay of the results provided by the measurements outlined above will be discussed in detail in the following sections.

All data processing that follows involves the calculation of common-view data from CGGTTS files, averaged over all satellites in view for each point (every 16 minutes) and corrected for ionosphere using the IGS measured ionosphere model and IGS Final orbits; all relativistic corrections associated with this work are incorporated via such standard products.

B. Common-clock measurements

The two receiver units at each location (MI60, ND and FD) are referenced to the same clock, providing common-clock, short-baseline measurements. In Figure 6 are shown the results of the measurements: the mean of the data represents simply the difference between the internal delays (INTDLY) for the two units and it is not of interest at this time, except for noting that at MI60 and ND the two units are different (OEMV and OEM6) and the difference between their INDLY is larger (around 20ns) than for FD, where both receivers are OEMV.

The interruptions in the data for both MI60 and ND are due to construction work that started in both locations around June 2012 (MJD56075). The third location, FD, didn't have any disruptive activities and has continuous data except for a few days around MJD56125 when the two antennas were moved approximately 18 m and 16 m due west respectively. The difference between the mean of the common-clock measurements before and after the relocation is approximately 120 ps. This variation could be due to systematic multipath, the effects of moving the antenna, or errors in the position determinations; however it is small enough to be ignored in this work.

Figure 7 shows the total time deviations [8] calculated from the data in Figure 6. For both MI60 and ND, only a subset of the data was used to compute the time deviation to avoid the frequent interruptions due to construction work after approximately MJD56080.

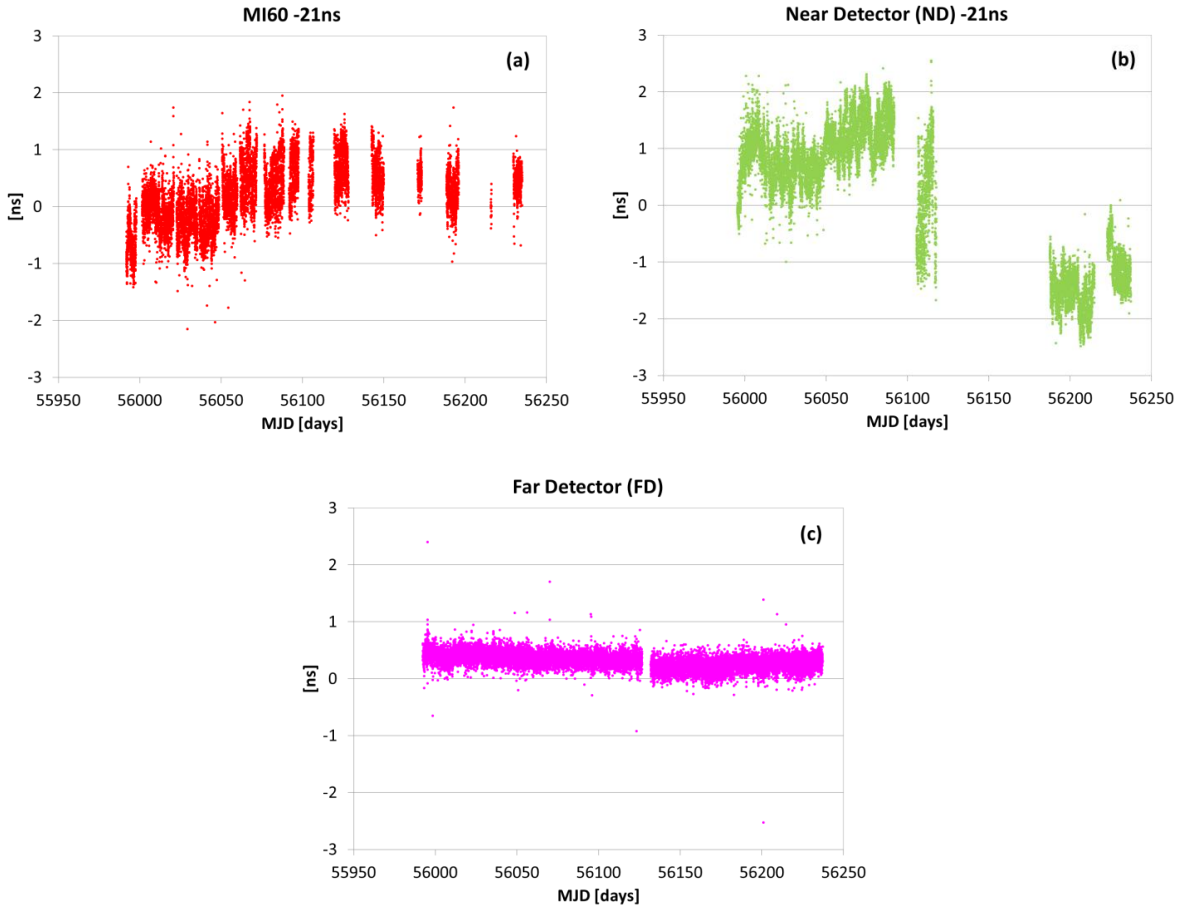

Figure 6. Common-clock, short-baseline measurement results. (a) The location is MI60 and the receivers are called GPS1 (OEMV) and GPS7 (OEM6). The data are shifted by 21 ns for clarity. (b) The location is ND and the receivers are called GPS2 (OEMV) and GPS8 (OEM6). The data are shifted by 21 ns for clarity. (c) The location is FD and the receivers are called GPS3 and GPS4 (both OEMV).

The receivers at FD, together with the amplifiers distributing the timing reference to them, are inside a tightly controlled environmental chamber that holds the temperature to ±1 K, which, together with a relatively multipath-free environment, may explain time stability twice as good as in the other locations. The stability at FD was confirmed by an independent analysis at the USNO, which computed the RMS of the Code-Phase difference of PPP residuals. The residuals are sensitive to the differential sensitivities of Code and Phase to both multipath and environmental variations.

All locations show that the intrinsic receivers' stability is below 200 ps, and can potentially be reduced to less than 80ps if a better controlled environment is provided.

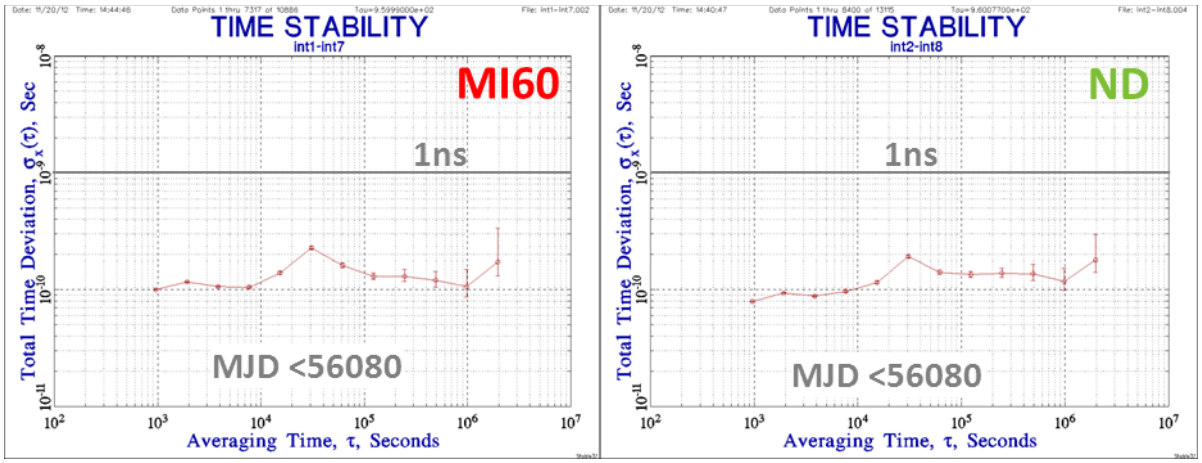

(a) (b)

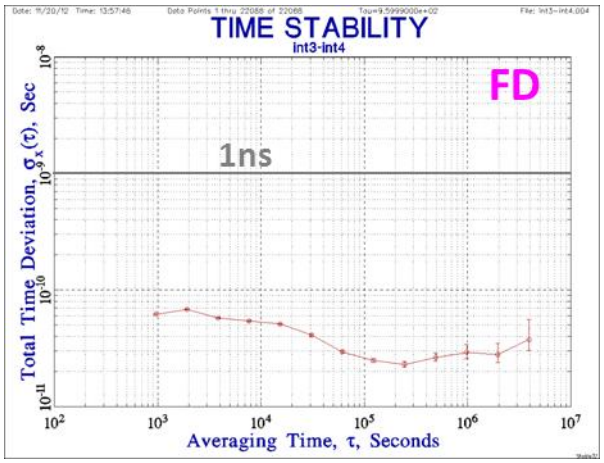

(c)

Figure 7. Time deviation calculated for the data shown in Figure 6. (a) MI60: only data before beginning of construction work were used (MJD < 56080). (b) ND: only data before the beginning of construction work were used (MJD < 56080). (c) FD: all data are included.

C. Differential calibration with travelling receiver

These are also common-clock, short-baseline measurements exactly of the same kind as those described in the section above. The measurements are utilized in pairs, each one performed at a different location, to calibrate the link between them by providing the number corresponding to the difference between the INTDLY of the two receivers constituting the link.

Figure 8 shows the results of the measurements used to calibrate the longer baseline link between ND and FD. On the left side there are the results of the common-clock measurements performed at ND starting on MJD56034 between the travelling receiver GPS5 and the local one GPS2: the mean is the difference between $INTDLY_2$ and $INTDLY_5$, while the total time deviation is used to estimate the uncertainty associated with that difference. Likewise, on the right side are shown the data from the common-clock measurement between the same travelling receiver GPS5 and the fixed GPS4 at FD, starting on MJD56005.

The difference between the two means is $INTDLY_2$-$INTDLY_4$, with an associated uncertainty that is the RSS (Root Square Sum) of the uncertainties for each measurement:

$$INTDLY_2 - INTDLY_4 = 0.6 \pm 0.2 \text{ ns} \qquad (3)$$

Similar measurements are performed at all three locations: now, for each link, all the terms in (1) are known quantities, allowing for the computation of the time difference between the time reference planes at each pair of locations.

The uncertainty of these measurements represents the stability of a single calibration instance, telling about the noise and unwanted variations that occur during the approximately week-long measurement, but it doesn't answer the question regarding the repeatability of such calibration instances.

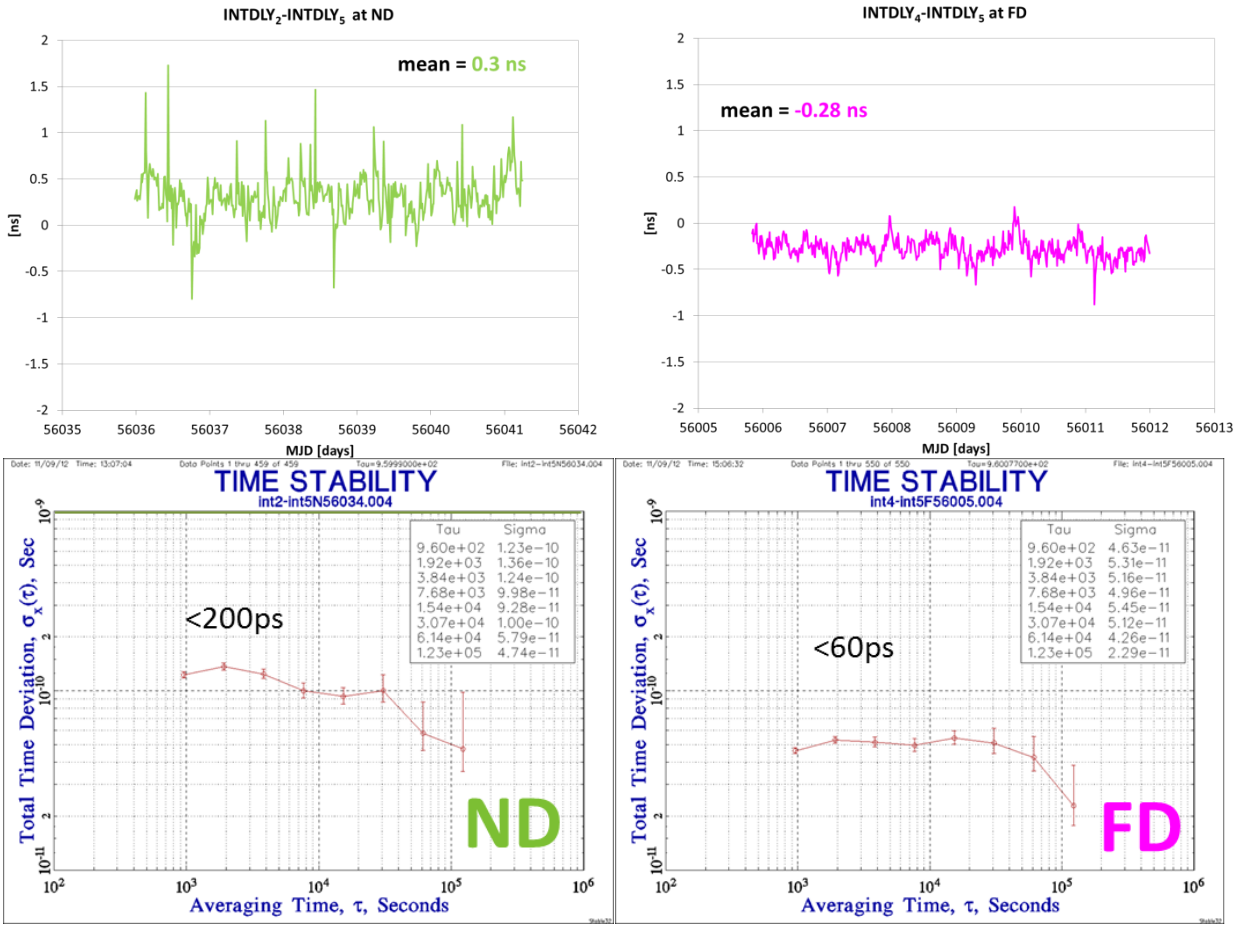

Figure 8. Example of two calibration instances used to calibrate the long-baseline link between ND and FD. (left) Data from a common-clock measurement between GPS2 and the travelling receiver GPS5, starting on MJD56034. (right) Data from a common-clock measurement between GPS4 and the travelling receiver GPS5, starting on MJD56005.

D. Calibrated double-difference for each link

Although with three locations there are three possible links, we will mostly work on two of them: the short-baseline one between ND and MI60 (approximately 1.2 km) that can be compared with the same one implemented with the fiber-based TWTT, and the long-baseline link between ND and FD, that can be compared with its point calibration performed by the TWSTFT equipment.

Because there are two fixed receivers at each location we have four highly correlated (virtually identical) GPS synchronization links for each pair of locations: GPS2-GPS1, GPS2-GPS7, GPS8-GPS1 and GPS8-GPS7 for the link between ND and MI60 and GPS2-GPS4, GPS2-GPS3, GPS8-GPS3 and GPS8-GPS4 for the link between ND and FD.

Either link could be computed using either GPS antenna at each site, or their average. The travelling antennas through repeated calibrations provide tie-breaking data, in the case of large deviations between the pair at each site. The variations within the pairs at each site provide a lower estimate to the uncertainties.

We chose to work with a subset of the possible links: in Figure 9 (right side) is shown the result of the double difference between (GPS2-GPS1) and (GPS8-GPS7) for the short-baseline link between ND and MI60.

In Figure 9 (left side) is shown the double difference between (GPS2-GPS4) and (GPS8-GPS3) for the long-baseline link between ND and FD.

The mean of the double-difference data and its total time deviation both contribute to the uncertainty of the link: the total time deviation represents the stochastic component of such uncertainty, while a non-zero value for the mean may indicate variations in of the systematic effects (likely in the calibration process).

Figure 9 shows that the ND-FD and ND-M160 double-differences jumped around 2 ns at around MJD 56200. Figure 10 shows that the travelling receiver GPS5 measured a 2 ns variation in GPS8 before and after this time. We conclude that GPS 8 suffered a 2 ns calibration variation, and this is also consistent with common clock differences of GPS2-GPS8.

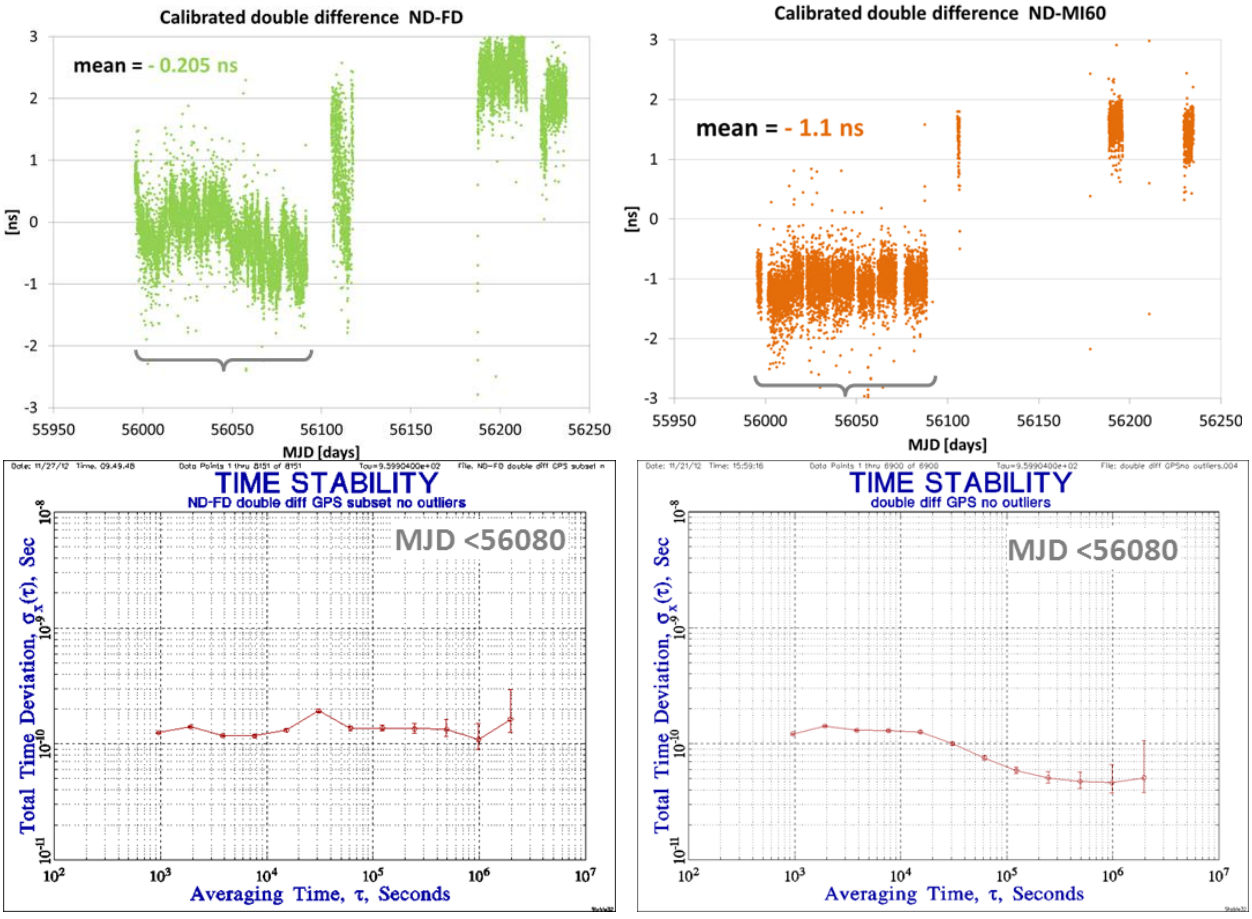

Figure 9. Double-differences for two links. The total time deviations were calculated using a subset of the data sets for MJD < 56080. The discarded data were collected with ongoing contruction work at ND and MI60. (left) Long-baseline synchronization link between ND and FD. (right) Short-baseline between ND and MI60. A small number of outliers (> 5 sigma) was removed from the data series before computing the total time deviation.

Some considerations can be made for the data subset indicated in Figure 9, corresponding to the period of time during which the particle beam for the MINOS experiment was actually running. The accelerator facility was, in fact, turned off around MJD 56090 for a scheduled upgrade involving construction work that caused the frequent interruption of data shown for MJD larger than 56090.

Concerning the long-baseline link between ND and FD, the mean value shown in Figure 9 (left side) is consistent with zero according to the statistical uncertainty estimated from the total time deviation.

In the case of the short-baseline link between ND and MI60, the mean value of the double difference shown in Figure 9 (right side) is larger than the total time deviation. We have chosen to be conservative and increased the uncertainty of the link to 1.1 ns.

More information regarding the overall uncertainty of the link will be gleaned from the comparison between different synchronization techniques.

E. Repeated differential calibrations

Over the first nine months of duration of this synchronization experiment, and because of the availability of two simultaneously travelling receiver units, it was possible to perform several differential calibrations at all three locations and at NIST, used as a control location.

The travelling receiver units are named GPS5 and GPS6 and travel with their antennas and dedicated antenna cable. The results of all calibration instances are shown in Figure 10 when the travelling receiver is GPS5 and in Figure 11 when the travelling unit is GPS6.

Different colors indicate specifically different units at each location, identified by their number: for example *int1-int5* refers to the difference $INTDLY_1 - INTDLY_5$ measured at MI60 for GPS1 and GPS5; *intNIST-int6* refers to a calibration instance at NIST where $INTDLY_{NIST} - INTDLY_6$ is measured for our primary receiver (called NIST) and GPS6; etc. Different symbols indicate different locations: ▲=NIST, ■=FD, ●=ND, ◆=MI60.

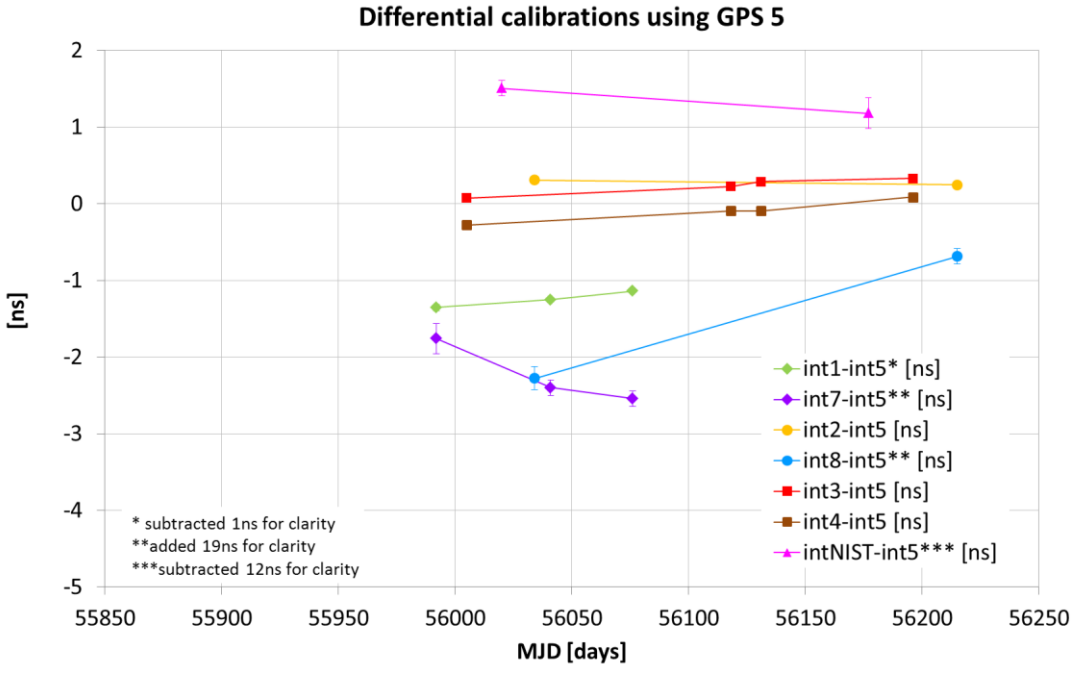

Figure 10. Complete collection of calibration instances performed so far at all locations using travelling receiver GPS5. Each point is the mean of several days of common-clock data. Different colors refer to different receiver units, the symbols indicate different locations: ▲=NIST, ■=FD, ●=ND, ◆=MI6

Each point shown in the plots is the mean of several days of common-clock data, while the error bars represent the stability of that calibration instance, with some caveats. The data shown in Figure 8 are representative of the statistical properties of all calibrations: the total time deviation is mostly constant for shorter averaging times, indicating non-stationary processes and then tends to decrease. The error bars shown in Figure 10 represent the stability of each calibration instance, estimated using the total time deviation at the end of the measurement duration. With reference to Figure 8, for example, the value used for the error bars in Figures 10 and 11 would be 50 ps for *int2-int5* and 30 ps for *int4-int5*.

The first information gleaned from Figure 10 is that there are clear differences in term of short-term stability (error bars associated with each calibration instance) among different instances of calibration at different locations. We know that the environmental situations are of different quality, with the best performance out of FD where the receivers and the distribution electronics are in a thermally controlled chamber where the temperature is kept within 1 K.

The second information is that the distance between mean values for different calibration instances for the same receiver is larger than the error bars associated with any of the measurements. In other words, the calibration values are not consistent with the same Gaussian distribution, but point towards non-stationary characteristics of the calibration process, i.e. variations in the system calibrations over time. It is not yet clear if those variations will constitute a clear trend that could be eliminated from the data (i.e. annual term) or if they will remain stochastic in nature.

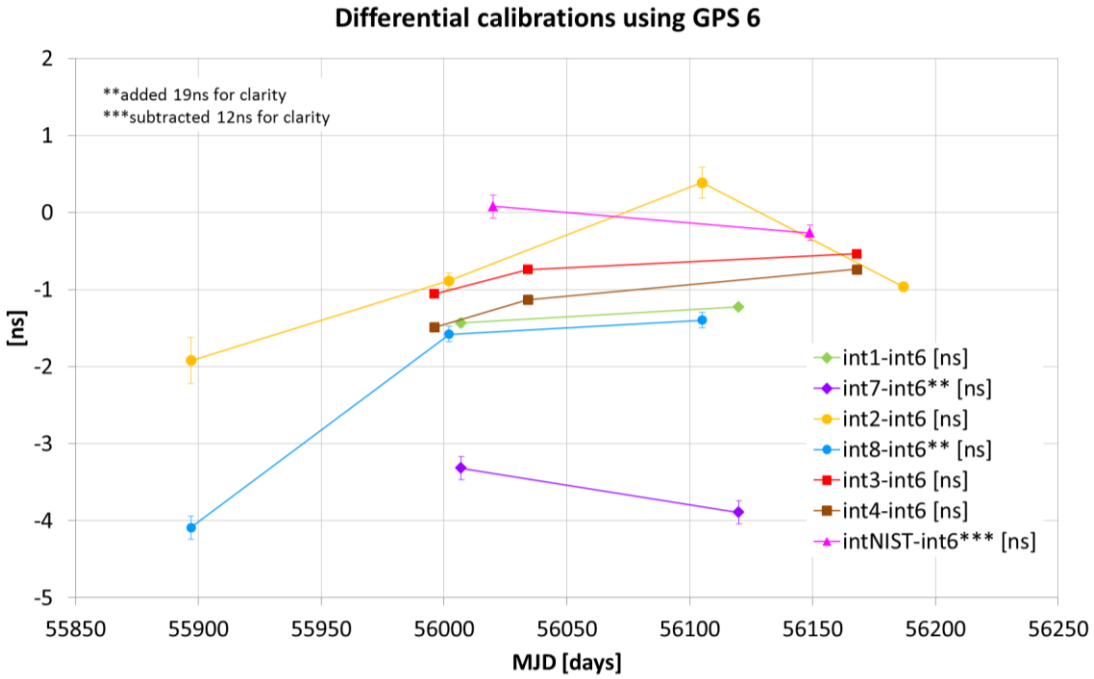

Figure 11. Complete collection of calibration instances performed so far at all locations using travelling receiver GPS6. Each point is the mean of several days of common-clock data. Different colors refer to different receiver units, the symbols indicate different locations: ▲=NIST, ■=FD, ●=ND, ◆=MI6.

The uncertainty associated with all calibrations is then increased to include the variations from one calibration instance to the next. Notwithstanding what we just said regarding the non-stationary nature of such distribution, we still choose to represent it by calculating the standard deviation of all the calibration values for each receiver. The reason for this choice is the too-small size of our samples (sometimes limited to only two occurrences) that doesn't allow the use of other statistical tools.

Table 1 and Table 2 summarize the calibrations shown in Figure 10 and 11. The mean and the standard deviation are calculated over all calibration instances for any one receiver as calibrated by a specific

travelling unit: $INTDLY_n$-$INTDLY_5$ and $INTDLY_n$-$INTDLY_6$ are computed for *n* = GPS1, GPS2, GPS3, GPS4, GPS7 and GPS8.

Table 1. Summary of all differential calibrations performed with GPS5.

GPS5		**$INTDLY_n$ –$INTDLY_5$ [ns]**
MI60 (◆)	***GPS1***	-0.25 ± 0.09
	GPS7	-21.2 ± 0.34
Near Detector (●)	***GPS2***	0.28 ± 0.03*
	GPS8	-20.5 ± 0.8*
Far Detector (■)	***GPS3***	0.2 ± 0.1
	GPS4	-0.1 ± 0.13
NIST (▲)	***NIST***	13.3 ± 0.16*

*Only two calibration instances

Table 2. Summary of all differential calibrations performed with GPS6.

Calibration with GPS6		**$INTDLY_n$ –$INTDLY_6$ [ns]**
MI60 (◆)	***GPS1***	-1.3 ± 0.1*
	GPS7	-22.6 ± 0.29*
Near Detector (●)	***GPS2***	-0.8 ± 0.82
	GPS8	-21.3 ± 1.2
Far Detector (■)	***GPS3***	-0.8 ± 0.2
	GPS4	-1.1 ± 0.31
NIST (▲)	***NIST***	11.9 ± 0.17*

* Only two calibration instances

The uncertainties estimated from repeated calibrations contribute to the uncertainty for each synchronization link implemented by GPS receivers, as shown in Table 3, where $INTDLY_n - INTDLY_m$ is determined for (n,m) = (2,1), (8,7), (2,4), (8,3).

Table 3. Results from calibrations for the GPS-based synchronization links.

GPS link uncertainty (from differential calibrations)		**$INTDLY_n$ –$INTDLY_m$ [ns]**	
		GPS5	**GPS6**
ND-MI60 (~1.2 km)	***GPS2-GPS1***	0.53 ± 0.1	0.5 ± 0.8*
	GPS8-GPS7	0.7 ± 0.9	1.3 ± 1.2
ND-FD (~735 km)	***GPS2-GPS4***	0.38 ± 0.1*	0.3 ± 0.9
	GPS8-GPS3	-20.7 ± 0.8*	-20.5 ± 1.2

* Only two calibration instances

III. COMPARISON WITH THE FIBER-BASED TWTT

The fiber-based TWTT technique is available only for the short-baseline link between ND and MI60 (1.2 km). There is a link that transfers 5MHz signals and one for the PPS signal. Each link is implemented using two optical fibers in the same bundle, one for each direction.

As previously discussed, the comparison between calibrated synchronization links using different techniques highlights the presence of eventual systematic effects inherent with each technique and may cancel when the double difference is calculated between similar links (i.e. two GPS links).

The easiest way to compare two synchronization links is to calculate the double difference, shown in Figure 12 (a) between each of the GPS links between ND and MI60 and the fiber based-TWTT. Figure 12 (b) shows the associated total time deviations. The strange bimodal distribution that appears in both curves indicates that it is not a feature of the GPS links, but it is part of the fiber-based TWTT.

Table 4 summarizes the mean and time deviation numbers for the double differences shown in Figure 12, together with the same values for the double difference between the two co-located GPS links. The conclusion is that there are significant systematic effects in all measurements and the overall link synchronization uncertainty should be increased accordingly.

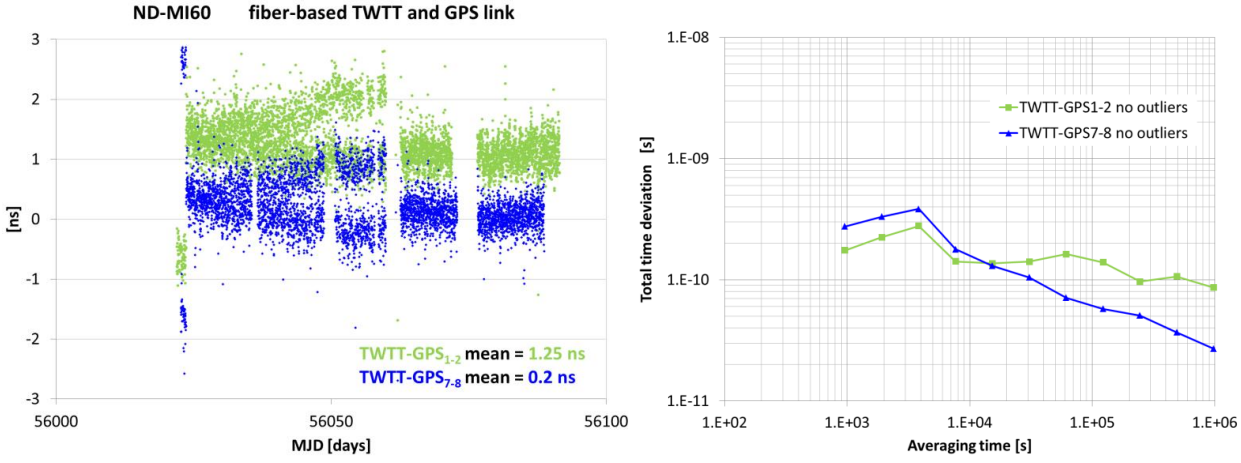

(a) (b)

Figure 12. Double difference between each on the two calibrated GPS links and the fiber-base TWTT for the link between ND and MI60. A small number of outliers (> 5 sigma) was removed from the data series before computing the total time deviation.

Table 4. Comparison of GPS link with fiber-based TWTT.

ND-MI60 (~1.2km)	**Uncertainty [ns]**	
	stochastic effects	***systematic effects***
TWTT- GPS1-2	0.1	1.25
TWTT- GPS7-8	0.03	0.2
GPS7-8 – GPS1-2	0.05	1.1

IV. COMPARISON WITH TWSTFT

The TWSTFT link, courtesy of USNO, was available only for a few days and therefore used for a point calibration, performed in April 2012 (MJD 56036-56037).

The setup included a mobile station mounted on a van that was parked just outside the building at ND and a flyaway station installed at the FD location. Both mobile stations performed time transfer with each other and with the base station at USNO in Washington, DC.

The results therefore include two TWSTFT links between ND and FD: one direct and the other one going through USNO. In Figure 13 the double difference between the two TWSTFT links is shown in blue. The double differences between each of the GPS links and the direct TWSTFT link are shown in purple and pink.

The mean value for both double differences in Figure 13 is well within their associated statistical uncertainties, confirming the absence of detectable systematic effects that was already indicated by the double difference between the two GPS links shown in Figure 9. The statistical uncertainty can be estimated from the total time deviation graph, on the right side of Figure 13 and is approximately 0.6ns.

Because the duration of the link was very short (only a couple of days) and the data points for it are more sparse (approximately once an hour) that in the case of a GPS or a fiber based link, the uncertainty analysis described above should be considered with prudence. What can certainly be said is that the GPS link and the TWSTFT link agree within 0.6ns.

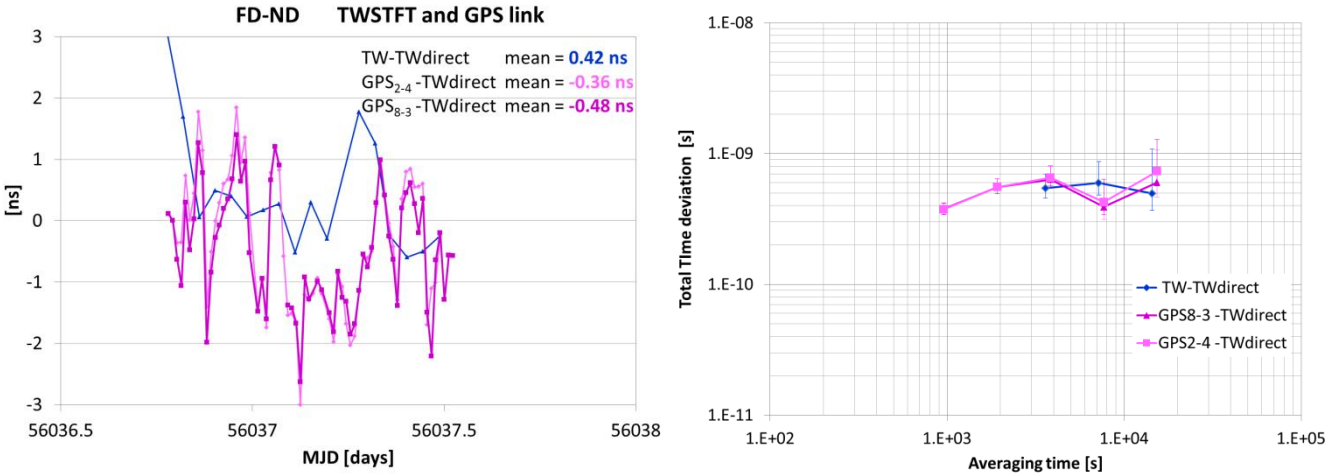

Figure 13. Double difference between each on the two calibrated GPS links and the TWSTFT direct link between ND and FD.

V. SUMMARY AND FUTURE WORK

In the wake of the unexpected results regarding the speed of neutrinos that were published in the fall of 2011, the need to achieve ns-level synchronization between the locations of the MINOS time-of-flight experiment created the opportunity for what we believe is the first known field experiment of this kind outside the timing community.

Two permanent GPS links were deployed between the three locations of the experiment: MI60 (the origin of the particle beam), ND (the Near Detector at approximately 1.2 km) and FD (the far Detector at approximately 735 km). A permanent fiber-based TWTT is available between MI60 and ND, while a point calibration of the longer link ND-FD was performed using TWSTFT (courtesy of USNO).

Two travelling GPS systems were also sent to all locations to periodically calibrate the resident GPS systems.

Different measurements were performed at all locations with the goal of achieving as complete as possible a characterization of the two main links: the short-baseline one MI60-ND and the long-baseline one ND-FD.

The results of the diverse methods used to evaluate the performance of the GSP synchronization links are summarized in Table 5.

The short-baseline link between MI60 and ND shows the presence of systematic effects that increase the overall uncertainty of the link to 1.25 ns. The long-baseline link between ND and FD shows no evidence of detectable systematic effects, showing agreement with the TWSTFT point calibration well within the required ns-level uncertainty. Nonetheless, the reader must keep in mind the limitation of a point calibration with respect to the possibility of a continuous comparison, as it is the case for the short-baseline link.

The large variations between the systematic uncertainties associated with the repeated calibrations are mostly due to the different quality of data collected in the different calibration instances, indicating the possibility for improvements if the environmental parameters are better controlled.

Table 5. Performance of the GPS synchronization link as estimated by a diverse set of measurements.

Measurement technique	**Location/link**	**Uncertainty [ns]**	
		statistical effects	***systematic effects***
Common-clock measurements	***MI60***	0.2	-
	ND	0.2	
	MD	0.04	
Calibrated GPS double difference	***ND-MI60***	0.2	1.1
	ND-FD	0.2	-
Repeated differential calibrations	***ND-MI60***	-	{0.1, 0.8, 0.9, 1.2}
	ND-FD		
Double difference with fiber-based TWTT	***ND-MI60***	0.2	1.25
Agreement with TWSTFT	***ND-FD***	0.6	

The ns-level uncertainty (68% confidence) indicated as a potential goal at the beginning of this paper has been substantially achieved. Nevertheless, he collaboration between NIST, USNO and the MINOS collaboration is continuing and may include a second TWSTFT calibration of the long-baseline link and a third, independent calibration of both links using high-performance Symmetricom 5071A[3] in a travelling configuration often referred to as "clock trip". In addition, a more detailed analysis of the GPS data should allow a more refined estimate of the uncertainties.

[3] This information is provided for technical completeness. As a matter of policy, neither the authors nor their institutions can endorse any commercial product or make generalized evaluations of product performance.

ACKNOWLEDGMENT

The authors would like to thank Judah Levine and Marc Weiss for their useful comments and discussions and Tom O'Brian for his support throughout the collaboration.

REFERENCES

[1] T. Adam, et al., "Measurement of the neutrino velocity with the OPERA detector in the CNGS beam," Preprint submitted to the Journal of High Energy Physics (17 November 2011).
[2] V. Bocean, et al., "Geodetic distance determination between the MINOS detectors for the neutrino time of flight measurements", PTTI 2012 Proceedings.
[3] P. Adamson, et. al., "Improved measurement of muon antineutrino disappearance in MINOS," Phys. Rev. Lett., 108, p. 191801, 2012.
[4] P. Adamson, et al, "Measurement of the velocity of the neutrino with MINOS", PTTI 2012 Proceedings.
[5] P. Adamson, et. al., "Measurement of neutrino velocity with the MINOS detectors and NuMI neutrino beam," Phys. Rev. D, 76, p. 072005, 2007.
[6] *http://webapp.csrs.nrcan.gc.ca/index_e/products_e/services_e/ppp_e.html*
[7] S. Romisch, et al., "Enabling accurate differential calibration of modern GNSS receivers", PTTI 2012 Proceedings.
[8] M.A. Weiss and D. A. Howe, "Total TDEV", Proc. 1998 IEEE Intl. Freq. Cont. Symp.