

Timing on the Fly

Synchronization for Direct Georeferencing on Small UAVs

JOHN PERRY AND JOSHUA CHILDS
UNIVERSITY OF FLORIDA



Researchers have long recognized the importance of synchronization in direct georeferencing – correlating the time of sensor activation with the position of the platform at that instant. Accomplishing this affordably on board lightweight autonomous aerial vehicles, however, has special challenges. This article describes the development of *Burredo*, an inexpensive, lightweight sub-microsecond timing system for synchronizing the position and attitude of a remote sensing system.

Sensor fusion is a predominant feature of modern navigation systems. To integrate navigation systems with other sensors, the spatial and temporal relationship of the sensor systems must be defined and calibrated.

Although system developers have employed a variety of schemes to great success in calibrating the spatial relationship between sensors, temporal integration is often less straightforward. Much of this difficulty arises from the problem of characterizing the timing of sensor measurements. Sensors sampled at varying, out-of-phase, or even irregular frequencies are typical situations encountered in temporal integration and further increase the difficulty.

For the spatial relationship, a strap-down inertial platform often provides a simple approach to calibrating the various sensors using a single framework. However, an analogous approach of fixing the timing of sensor measurements to a single temporal framework is more difficult, if not impossible, to achieve in many scenarios. This is particularly true when attempting to integrate consumer off-the-shelf devices, which often offer limited synchronization capabilities.

For most fused navigation systems, the task of temporal integration is handled within the design of the devices themselves. For example, an integrated inertial navigation system (INS)/GPS might be designed to sample the inertial sensors at a synchronized frequency

division of the GPS one-pulse-per-second (1PPS) signal. However, the proliferation of ad hoc aided navigation systems suggests that such cohesive low-level device design may give way to more modular, higher-level synchronization schemes.

Such was the case when the University of Florida's Unmanned Aerial Vehicle Research Group went about developing a small (one-kilogram) payload for the NOVA II small UAV platform, consisting of a consumer-grade digital single-lens reflex (DSLR) camera with an integrated INS/GPS system (see **Figure 1**). This UAV platform is the result of over a decade of progress by an interdisciplinary group at the University of Florida and was developed principally for state and federal

agencies to conduct infrastructure and environmental monitoring. The principal purpose of the payload is to produce high resolution, directly georeferenced imagery, which requires knowing the position and attitude of the camera at the moment of exposure.

This article describes the development and analysis of a modular synchronization scheme that should elucidate the principle techniques required for such a system. In this case, very little access to low-level sensor circuitry or control logic was available. Instead, we were left to integrate largely black-box sensor systems using a modular approach that relied on simple external timing signals.

Time versus Timing

Historically, time has simply been defined in terms of periodic events observable in nature. The question, “When did this measurement occur?”, might be casually answered by checking the hands from the grandfather clock in the hall, or glancing at the readout on a digital wristwatch, or perhaps examining a recent GPS packet.

All of these reflect the cumulative phase angle of a periodic event, whether from pendular motion, a crystal oscillator, or the resonant frequency of cesium atoms. Such periodic signals establish a basis for measuring time. From these we can establish the “absolute” timing of an event against a given temporal reference (e.g., it occurred at 12:00:00.00 UTC), the “relative” timing of two events (e.g., event A occurred three seconds after event B), or the duration of an event (e.g., it lasted for 3.00 seconds).

Outside of the GNSS community, the fact that the GPS constellation has provided unprecedented access to a globally available timing and synchronization framework is much less recognized than its navigational capabilities. The ability to reference events to GPS time, and in turn to Coordinated Universal Time (UTC) and other timing frameworks, is the best available means by which we can measure time against a standard reference.

To emphasize again, however, under-

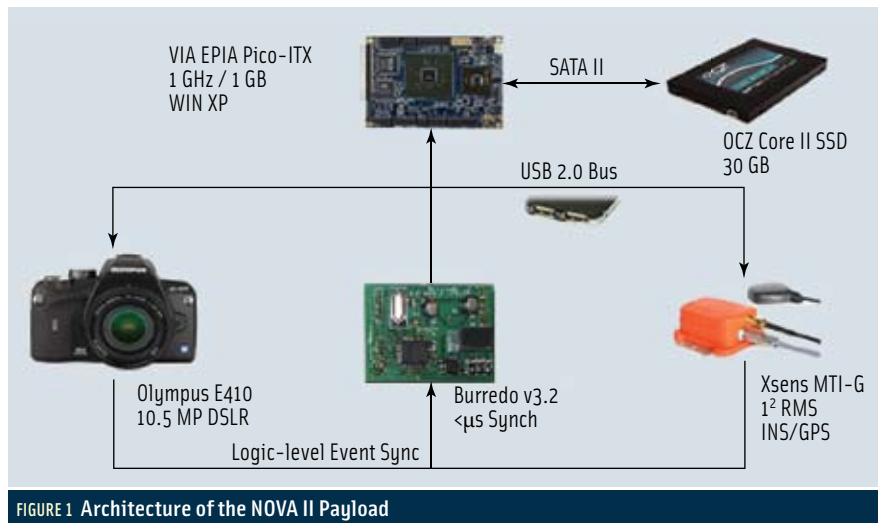


FIGURE 1 Architecture of the NOVA II Payload

lying all timing systems is the simple notion of periodicity and cumulative phase angle as the “measuring stick” of temporal relationships. The idea of synchronization is not limited to the concept of simultaneity, nor does the time need to be indicated with respect to some standard reference. These two concepts are core to modular synchronization.

Synchronization and Georeferencing

Prior research has established the importance of synchronization in direct georeferencing, and given rise to a number of system integration techniques. The previous version of the NOVA II, the NOVA I, carried a payload that employed a synchronization scheme that relied on both on timing the arrival of navigation data packets at the host as well as the camera trigger commands. This design was reported to have a synchronization error of 87 milliseconds. (See the paper by S. Bowman listed in the Additional Resources section near the end of this article.)

A similar implementation found that the synchronization error was as much as 333 milliseconds using the same architecture. (See the paper by H. Chao et alia, Additional Resources.) This method and the resulting wide discrepancy in sensor synchronization error highlight the need for direct measures of the time of sensor sampling. Such timing errors translate into substantial accuracy problems for

meeting NOVA II mission georeferencing requirements, as will be discussed in a later section.

Relying on data packets can be problematic due to variable processing speeds and transmission time of the packet on the sensor end, as well as indeterminate timing behaviors from software-based synchronization on a non-real-time operating system. Particularly for commercial-off-the-shelf (COTS) cameras, features such as white balancing and autofocus can produce stochastic and unquantified delays between trigger and exposure. Without a feedback mechanism, timing such a system is an intractable problem.

A more rigorous approach commonly implemented is the use of a digital acquisition (DAQ) card in the host computer, which directly records timing signals. However, a DAQ card is impractical on a small UAV platform due to its size and lack of an available interface with the host computer. Nonetheless, the use of DAQ cards has been shown to obtain synchronization accuracies ranging from 0.4 milliseconds [Li et al., 2006] to 5 microseconds. (For details, see the articles by B. Li et alia and C. Toth et alia, respectively, listed in Additional Resources section.)

The task of synchronizing the camera to the navigation system does not require that we know the UTC time at which the camera exposure occurs. Rather, the parameters of interest are in fact the navigational state of the system

at that time, e.g., the position and orientation of the camera at the instant of exposure — although this does not preclude knowing the reference time.

Furthermore, we do not need to sample the navigation sensors at the precise moment of the camera exposure, but rather we can estimate the state at the time of exposure from adjacent navigational states given an appropriate model or assumptions. With these two considerations, we proceed to develop the idea of a purely relative timing system with which we can synchronize sensors.

Relative Timing and Synchronization

To begin, some periodic signal must serve as a basis for measurement. Given our rather stringent weight limitations on the UAV, both the grandfather and atomic clocks seemed impractical; so, we settled on a crystal oscillator.

Because the UAV system makes all timing measurements with respect to this oscillator, its frequency will set the maximum obtainable resolution for our synchronization system, denoted f_{clock} . The output of the oscillator is used to drive a digital counter, which increments one integer value for each oscillation of the crystal. A measurement is made by reading off the value of the digital counter at the moment of interest.

Assuming a uniform distribution of measurements with respect to the phase of the oscillator, we expect a quantization error for each timing measurement with a standard deviation of **Equation 1**:

$$\sigma_{quant} = \sqrt{\frac{1}{12}} \cdot \frac{1}{f_{clock}} \quad (1)$$

Suppose we want to synchronize some periodic sensor measurement to another sensor's measurement. The former will be designated the "master" sensor and the latter arbitrary signal, the "slave" sensor. The interval, Δt_{epoch} between any two measurements of the master sensor defines one epoch of time, and is given by **Equation 2**:

$$\Delta t_{epoch} = t_m - t_{m_0} \quad (2)$$

Here, the values of t_m and t_{m_0} refer to the cumulative phase angle of the crystal oscillator recorded at the moment the

synchronization signal arrives at the synchronization device. The synchronization signal is typically characterized by a square-wave electrical signal generated by the sensor, with the time of sampling measured at the leading edge of the signal. Similarly, the timing signal of the slave sensor, t_s , quantifies the moment the slave sensor is sampled.

Using the notion of relative timing, we can use the master epoch as a temporal reference to assign the relative timing of each captured image by calculating the time of exposure, ρ , as a fraction of the master epoch, given by **Equation 3**:

$$\rho = \frac{t_s}{\Delta t_{epoch}} \quad (3)$$

It is useful to consider each epoch separately, so that each occurrence of the master signal resets the counter to zero; hence, $t_{m_0} = 0$, thereby simplifying the relative timing measurement to **Equation 4**:

$$\rho = \frac{t_s}{t_m} \quad (4)$$

Each of these values is given by reading off the digital counter in timing units that are determined by the frequency of the clock. Because the units cancel in ratio, this provides an extremely basic notion of timing: The camera exposure occurs at some fraction of the time between two samples of the navigation sensor, and then, by a simple interpolation of the adjacent parameters, we can solve for the navigation state parameters at the instant of camera exposure. In this case, linear interpolation is appropriate with the assumption that the dynamics are linear between INS measurements.

Synchronization in Action

In practice, this simple synchronization scheme is adaptable to a number of scenarios, and depends largely on appropriately selecting the master and slave sensors. A typical application would be to reference the high-rate inertial sensors to the GPS epoch, allowing each inertial measurement to be given as a fraction of the GPS second. In this synchronization scenario, the INS is the slave and the GPS is the master.

However, we do not need to reference all of the sensors to the GPS epoch. For

example, we can simplify implementation of the interpolation scheme if we reference the camera exposure to the INS samples, because these have the highest temporal resolution of the navigation parameters of interest. In this case, the camera is the slave sensor and it is synchronized to the INS.

In the case of our NOVA II payload, it was also necessary to use the INS as the master sensor because the GPS 1PPS signal itself was not available. This was a convenient circumstance, however, because obtaining the navigation parameters for the camera exposure required simple linear interpolation using the ratio, ρ , between the adjacent navigation packets.

Each sensor must be sampled for a finite amount of time. For example, the camera, as the slave sensor, has a shutter speed of 0.0005 second. This produces a measurement error caused by characterizing an event of finite duration — the exposure time, or period during which the shutter/aperture is open — with a practically infinitesimal timing signal.

We thus face the question: when should the timing signal be generated relative to this finite duration of time? And further, for practical consideration, should the timing signal occur when the shutter begins to open, when it closes, or perhaps at some proportion of this?

This type of error provides some insight into the accuracy required for the clock, and is closely related to the dynamics of the parameters under consideration and their propagation through the sensor system. In analyzing the performance of the synchronization system, we have a critical need to consider not only the timing error but also the effect of those errors on the overall measurement system.

In the case of a camera, the subject of the overall sensor system is the camera's observed image. For direct georeferencing purposes, we are further interested in the change in navigation parameters with respect to the change in the observed image. Although the issue of perspective projective geometry is beyond the scope of this article, we need



to understand that the accuracy of the overall system is necessarily constrained by the aperture error of the sensors.

Clearly, if neither the camera nor the scene is in motion, then whether one defines the moment of exposure at the beginning, middle, or end of the camera's shutter action is irrelevant because the image and navigation parameters are constant. However, for systems (such as UAVs) in motion we must consider the dynamic relationship between the parameters being measured by the sensors.

For example, consider a camera mounted on the UAV, pointed straight down with an exposure time of 0.25 milliseconds and a nominal focal length of 5,000 pixels flying 25 meters per second straight and level at 100 meters above ground. In this scenario, the camera will move a linear distance of 6.25 millimeters while the shutter is open. When motionless, one pixel at the center of the image covers a ground distance of 20 millimeters along each axis. Due to the motion of the UAV during the time of exposure, however, the linear distance that the pixel actually covers is 26.25 millimeters.

Such an image in practice shows no visible blurring, which indicates that the error due to dynamics is negligible for our purposes. That is, almost zero change occurs in the image given the change in the navigation parameters during exposure. Consequently, we can make the loose assumption that the relevant dynamics are constant within the sampling period.

The selection of the parameters in this scenario is not by accident; the shutter speed of the camera is selected based on experiments to eliminate blurring due to typical small UAV dynamics.

Having made the rather weak assumption that dynamics are negligible for an appropriate sampling duration, this means that the accuracy of the synchronization signal must be such that it occurs within the exposure duration. A stronger assumption that the dynamics are linear while the camera shutter is open implies that we can minimize aperture error by applying the average

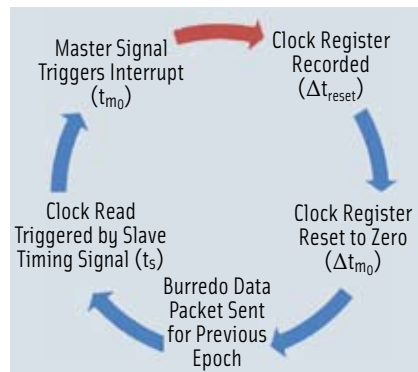


FIGURE 2 Simplified flow chart of the Burredo's operation

parameter value over the period of exposure, which is equivalent to defining the synchronization signal at the center of the duration of exposure.

A more general approach would be to take the time-weighted average of the parameters over the exposure duration, a case that is only practical where the frequency of the navigation parameter data is much greater than the shutter speed. Assuming that the dynamics are linear gives us a straightforward approach: we minimize the aperture error by defining the sampling time at the center of the duration of exposure. The required synchronization accuracy can then be calculated by back-propagating the allowable error in the parameters into a time error given the expected platform dynamics.

The weaker assumption that the parameters are constant over the sampling duration implies that we should again sample at the center of the exposure and ensure the synchronization is accurate to within one half the sampling duration.

The Flying Burredo

We implemented the synchronization scheme outlined here on a custom circuit board, dubbed the "Burredo". The Burredo is based off of an eight-bit microcontroller with an external crystal oscillator to improve the stability of the clock frequency.

We used the microcontroller's built-in timing facilities to implement the relative timing scheme described previously. Another advantage of using a microcontroller is the built-in serial

communication facilities, which allows the synchronization data to be passed to the host computer.

The synchronization signals utilized were the external "Sync Out" signal specified in the INS/GPS device's user manual and the X-sync flash circuitry common to most DSLR cameras. Synchronization signals from most devices can be accommodated by using simple signal conditioning that is included onboard the Burredo.

Total materials cost for the Burredo was less than \$100. It weighs less than 20 grams, and measures just $5 \times 3 \times 1.5$ centimeters. Figure 2 shows a flow chart of the Burredo's operation.

We evaluated the synchronization accuracy of the Burredo by comparing its performance to that of a GPS timing receiver. The receiver outputs the GPS 1PPS signal with an accuracy of ± 15 nanoseconds and has two event inputs that provide an event time stamp in UTC time with a resolution of 488 nanoseconds.

In the experiment, the GPS 1PPS signal was selected as the master signal. Because each master timing epoch was defined by the GPS second, the value of ρ as described in Equation 3 is a fraction of one second. A microcontroller-based circuit generated a pseudorandom trigger signal similar in frequency to the camera's capture rate for both the GPS receiver event input and the Burredo slave input.

The signal propagation delay due to cable length introduced an additional parameter, Δt_{prop} , into the experimental setup, and was calculated to be 120 nanoseconds for the receiver's cable length of 100 feet. All other propagation delays were negligible due to short cable lengths (on the order of <1 ns).

We collected the primary data set over a period of about six hours, with a trigger event occurring on average every 2.5 seconds. Several smaller confirmation data sets were also collected, all showing similar results. We modeled the measurements using Equations 5 and 6 for the GPS receiver and Burredo, respectively. We benchmarked the Burredo using the receiver measurement as

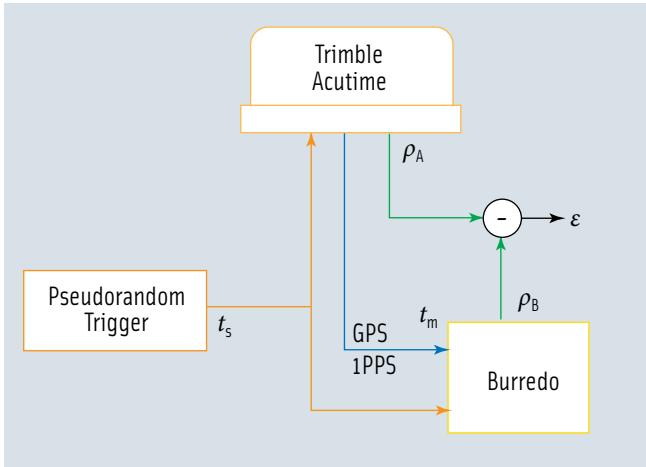


FIGURE 3 Overview of the experimental setup




FIGURE 4 Distribution of the pseudo-random triggers as a fraction of the epoch




FIGURE 5 Evident linear trend in the error of the uncorrected synchronization data attributed to




FIGURE 6 The correlation of the synchronization error and phase angle plotted against values of

the truth value; thus, the error equation was given by Equation 7.

Figure 3 provides an overview of the experimental setup.

$$\rho_{AcuTime} = \rho_A = \frac{t_s - (t_{m_0} + \Delta t_{prop}) - t_m}{t_m - t_{m_0}} \quad (5)$$

$$\rho_{Burredo} = \rho_B = \frac{t_s - (t_{m_0} + \Delta t_{prop})}{(t_m + \Delta t_{prop}) - (t_{m_0} + \Delta t_{prop})} \quad (6)$$

$$\epsilon = \rho_A - \rho_B = \frac{2\Delta t_{prop}}{t_m - t_{m_0}} + \nu_B \quad (7)$$

Experiment Results and Conclusions

Data collected indicate that the trigger signal as a fraction of the epoch was approximately uniform, as shown in Figure 4. Unexpectedly, a strong correlation between the fraction of the epoch and the magnitude of the error

emerged from the data (See Figure 5). Upon further investigation, we found that the error correlation is explained by a constant bias in the measurement of the Burredo's (master sensor) first sample time, t_{m_0} , caused by the finite time required to reset the value of the clock register at the beginning of each epoch. This error does not occur during the slave signal measurement, which is handled asynchronous in hardware.

We addressed this situation by introducing a term for the reset time, denoted Δt_{reset} , that modifies the Burredo's measurement equation, as given in Equation 8. Because the Burredo operates on integer clock cycles, this value was found by minimizing the correlation and rounding to the nearest integer value, so that Δt_{reset} , as shown in Figure 6. Figure 7 pres-

ents the remaining errors after applying this correction.

$$\rho_B = \frac{t_s - (t_{m_0} + \Delta t_{reset} + \Delta t_{prop})}{(t_m + \Delta t_{prop}) - (t_{m_0} + \Delta t_{reset} + \Delta t_{prop})} \quad (8)$$

In Figure 8, the oscillator frequency drift for the Burredo is clearly visible. The synchronization error was found to not be correlated to variations in the crystal frequency. The mean error between the GPS receiver and Burredo measurements was 208.9 nanoseconds. This bias corresponds to approximately three clock cycles of the Burredo.

The actual source of this bias is unclear and may be due to a number of factors, including the internal architecture of the Burredo's microprocessor timing facilities or the timing signal's conditioning circuitry. Without experimental verification of the source of this




FIGURE 7 The value of the synchronization error over time in hundreds of nanoseconds




FIGURE 8 Drift of the Burredo's crystal oscillator over time

bias, whether intrinsic to the device or due to experimental setup, the error may not be excluded from a conservative error analysis and is, therefore, attributed to the Burredo, resulting in an experimental RMS error of 256.4 nanoseconds.

The measurements of both the Burredo and the GPS receiver are subject to quantization error, which we calculated for both using the standard deviation of error introduced by quantization (Equation 1). The expected standard deviations due to quantization were 19.6 and 140.9 nanoseconds for the Burredo and the GPS receiver's event capture, respectively.

On the assumption that these measurements are independent, their difference will have a standard deviation given by the Euclidean norm of their individual standard deviations, or 142.2 nanoseconds.

This calculated value agrees closely with the observed standard deviation of ϵ , 148.7 nanoseconds, indicating that the majority of the observed random error can be explained by quantization alone.

The value calculated by propagating the quantization error does not take into account the uncertainties associated with the 1PPS measurement, stochastic effects within the physical system, or other unmodelled phenomenon, all of which contribute to the greater uncertainty observed in the data. The

presence of these additional uncertainties was confirmed by an F-test of the ratio of calculated to observed variances, which rejected the null hypothesis that the variances are statistically equal at the 95 percent confidence level.

In order to evaluate the effect of errors due to synchronization, we also needed to characterize the flight dynamics. Figure 9 shows the average dynamics calculated from a typical UAV flight line, based on data collected from a NOVA II flight on March 4, 2009 at 9:45 a.m. The temperature during the flight was 63°F

with winds out of the north at approximately 3.5 meters/second. Flying height during the data sample shown here was a commanded 150 meters above ground level.

From these parameters, we calculated the typical flight dynamics, producing the results given in Table 1. The table also shows the magnitude of the change in the parameters for the experimental RMS error in synchronization. Finally, the table lists the change in parameters expected during a camera exposure duration of 1/2000th of a second. It is

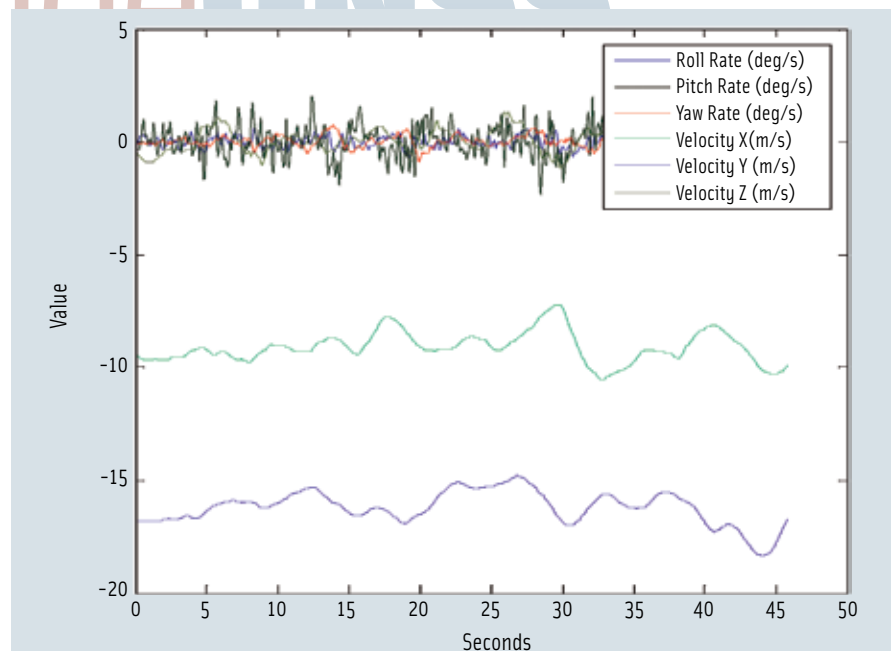


FIGURE 9 Typical NOVA II dynamics while flying straight and level

	Roll (deg)	Pitch (deg)	Yaw (deg)	Horizontal Position (m)	Vertical Position (m)
Dynamics	3.22°/s	7.19°/s	2.53°/s	18.61 m/s	0.79 m/s
Sync Error	$8.26^{\circ} \cdot 10^{-7}$	$18.46^{\circ} \cdot 10^{-7}$	$6.50^{\circ} \cdot 10^{-7}$	$4.77\text{m} \cdot 10^{-6}$	$2.03\text{m} \cdot 10^{-7}$
Camera Exposure	$1.6^{\circ} \cdot 10^{-3}$	$3.6^{\circ} \cdot 10^{-3}$	$1.3^{\circ} \cdot 10^{-3}$	$9.3\text{m} \cdot 10^{-2}$	$4.0\text{m} \cdot 10^{-4}$

TABLE 1. The calculated RMS error in the navigation parameters due to synchronization error and camera exposure duration in the navigation parameters for the given dynamics

evident that the error due to synchronization is orders of magnitude smaller than the potential error due to mischaracterizing the sample of finite duration.

In conclusion, this article has described the basic operation of a simple, flexible synchronization architecture that facilitates the temporal integration of sensors. It does not rely on the availability of the GNSS satellite constellation, although the GPS 1PPS signal does allow ready access to a standard time reference.

The Burredo device is highly modular, interfacing with the sensors using simple logic-level synchronization signals. It then independently calculates and transmits the synchronization information on an independent data bus to the host computer. An accuracy of less than a microsecond was demonstrated in a benchmarking experiment. Using the NOVA II platform and payload as an example, our investigation has shown that the errors due to synchronization were orders of magnitude less than the errors of the relevant sensors, thus confirming its utility in directly georeferencing remotely sensed imagery.

Manufacturers

The integrated GNSS/inertial system is the MTi-G from **Xsens Technologies B.V.**, Enschede, The Netherlands, with an integrated u-blox 5 GPS/Galileo chip. The microcontroller was an ATmega644p from **Atmel Corporation**, San Jose, California, USA. The DSLR camera is the EVOLT E-410 from **Olympus America, Inc.**, Center Valley, Pennsylvania, USA. The GPS receiver was used to evaluate the Burredo was the AcuTime Gold from **Trimble**, Sunnyvale, California.

Additional Resources

[1] Bowman, S., 2008 Design and validation of an autonomous rapid mapping system using a

small UAV, *Master's Thesis*, University of Florida, Gainesville, Florida, USA

[2] Chao, H., and M. Baumann, A. Jensen, Y.Q. Chen, Y. Cao, W. Ren, and M. McKee, "Band-reconfigurable Multi-UAV-based Cooperative Remote Sensing for Real-time Water Management and Distributed Irrigation Control," *Proceedings of the 17th World Congress of the International Federation of Automatic Control*, July 6-11, 2008, Seoul, Korea

[3] Lewandowski, W., and J. Azoubib, W.J. Klepczynski, "GPS: Primary Tool for Time Transfer," *Proceedings of the IEEE*, 87(1):163-172, 1999

[4] Li, B., and C. Rizos, H.K. Lee, and H.K. Lee, "A GPS-slaved Time Synchronization System for Hybrid Navigation," *GPS Solutions*, 10:207-217, 2006

[5] Skaloud, J., "Problems in Direct-Georeferencing by INS/DGPS in the Airborne Environment," *ISPRS Workshop on 'Direct versus indirect methods of sensor orientation', Commission III, Working Group III/1*, Nov. 25-26, 1999, Barcelona, Spain


[6] Toth, C., and S.W. Shin, D.A. Grejner-Brzezinska, and J.H. Kwon, "On Accurate Time Synchronization of Multi-Sensor Mapping Systems," *Journal of Applied Geodesy*, 2(3):159-166, 2008

[7] Xsens MTi-G, 2009. MTi-G User Manual, Xsens, Inc., Enschede, The Netherlands


Authors



John H. Perry is pursuing a Ph.D. in geomatics at the University of Florida. He is a recipient of a National Science Foundation Graduate Research Fellowship. His research focus is the integration of remote sensing and navigation systems using photogrammetric techniques, with an emphasis on UAV platforms.



Joshua Childs is pursuing a Ph.D. in electrical engineering at the University of Florida. His areas of study include machine intelligence and control systems. His interest is

in applying these concepts to UAVs. 

browse the

digital
edition

insidegnss.com