### Overview

The Arbiter Systems® Model 1133A Power Sentinel™ consists of several blocks. They are: GPS receiver and synchronization; voltage and current inputs; programmable-gain amplifiers, multiplexers, and analog-to-digital converter; digital signal processor; host processor; display and keyboard; I/O functions; and power supply. The key features of the instrument and its measurement functions are described in the sections which follow.

The Model 1133A uses two processors to handle the many tasks which are performed each second in such a highly-integrated instrument. The digital signal analysis tasks are performed by a TI TMS320C32 floating-point DSP unit. The I/O functions and interface to the "real world" are handled by an ST 10F167 16-bit microcontroller. Both of these processors have significant additional processing power which is not used in the initial version of the Model 1133A. This will allow for the future expansion of features, which is sure to come, with a minimum amount of upset; indeed, many future enhancements should be possible with only a firmware upgrade.

### GPS Synchronization

The Model 1133A includes an eight-channel global positioning system (GPS) receiver which provides time accurate to a fraction of a microsecond anywhere in the world. Using proprietary technology developed by Arbiter Systems and refined in several generations of GPS timing products, an internal 10 MHz crystal oscillator is slaved to the 1PPS output of the GPS receiver, maintaining its frequency at any time within a few parts in $10^{10}$. All of the internal timing signals are derived from this accurate timebase.

### Current Inputs

The Model 1133A has a three-phase current input which is designed to be accurate to a few ppm over time and temperature (see Error Analysis). This current input section uses a two-stage process, similar to a two-stage current transformer except that the first stage is dc-coupled. Each of these two stages has an accuracy of a few tenths of one percent; together, they have an accuracy of a few ppm.

### Voltage Inputs

The voltage inputs use low-TC voltage divider resistor networks (attenuators). This input may be configured as a three-phase, three-element input, with four connections (A, B, C, and optionally N); or it may be configured as a two-element input, with independent connections to each element (A+, A–, C+, and C–, for example).

### Self Calibration

The input sections are designed to provide exceptional stability over time and temperature through the use of several high-performance (and fairly expensive) components. The rest of the analog circuitry is implemented with lower-cost components. To reduce (essentially eliminate, in fact) the effects of drift and temperature sensitivity in the programmable gain amplifiers (PGAs) and analog-to-digital converter (ADC), an internal self-calibration source is provided. The inputs to the PGAs are multiplexed, each to four different signals: the three current (or voltage) inputs, and the calibration signal.

### Error Analysis

All identified sources of error in the Model 1133A have been quantified using worst-case manufacturers' performance data. These have then been combined using a root-sum-of-squares (RSS) method to yield a performance estimate. Effects due to initial calibration, measurement noise, temperature, and aging are all included. The reason for using RSS analysis is beyond the scope of this paper. However, we have found in our many years' experience building calibration instruments that this method yields the most realistic estimate of actual worst-case performance, provided that numerous errors contribute significantly to the overall performance (i.e., no one error dominates), and provided that actual worst-case data is used in the analysis.

The error analysis for the Model 1133A for power or energy measurements is shown in table 1. Similar analyses have been performed for the other functions of the instrument, and include most of the same factors shown here.

**Table 1 – Power/Energy Error Analysis**

| **Temperature Errors, 0-50° C:** | **Error, ppm** |
|---|---|
| Current input, resistor TCR | 63 |
| Voltage input, resistor ratio | 50 |
| Voltage reference (x2) | 50 |
| Cal. source, resistor ratio (x2) | 25 |
| **Time Stability, 1 year:** | |
| Current input, resistor | 25 |
| Voltage input, resistor ratio | <20 |
| Voltage reference (x2) | 36 |
| Cal. source, resistor ratio (x2) | <40 |
| **Measurement Noise** | 10 |
| **Total RSS Error, Basic** | 117 |
| **Calibration Errors:** | |
| Cal. artifact, Rotek MSB-001A | 50 |
| Traceability to National Standards | 50 |
| **Total Error, RSS** | **136** |
| Specification (0.025%) | 250 |

### *Signal Sampling*

The signals, representing the scaled voltage and current inputs and the calibration signal, are time-multiplexed into a two-channel ADC. Each signal is sampled at a rate of 10240 samples per second. The current and voltage signals for each phase are sampled simultaneously to eliminate errors in power calculation which would occur with non-simultaneous sampling. The sample clock in the Model 1133A is synchronized to UTC-USNO (GPS) within one microsecond, allowing measurements of phase angle across a power grid to be compared directly, and ensuring that revenue is billed at the correct rates.

### *Power and Energy*

Power and energy are determined by making twenty separate measurements per second of the cross-product of voltage and current for each phase. Each measurement uses 1024 samples (i.e., it takes data from a 100 millisecond window), yielding a 50% overlap.

Apparent power (VA) and reactive power (VAR) are determined from the results of the voltage and current magnitude measurements (see next section), using standard identities. The active power measurements and reactive power measurements are then compensated for PT and CT corrections (see below) using a complex multiplication, and corrections are performed for transformer iron and copper losses, if enabled.

At this point, two different things are done with the resulting measurement data. First, a determination is made of the quadrant in which this data should be registered (Wh delivered or received, VARh delivered or received.) The results of each measurement cycle (20/sec) are then added to the proper set of registers. These registers are stored periodically, and accumulation restarted from zero. The user may configure the unit for different intervals to register energy.

The second thing that is done with the data is to determine the actual power level. This number will be displayed on the front panel (as watts or VARs), and it will be returned via the serial interface if a simple request for "power" is made. This result is calculated by averaging the 20 power measurements made during each second. Therefore, the update rate for this quantity is once per second. This data is not registered separately depending on quadrant, as the energy data is; therefore, it is theoretically possible, if the direction of power flow changes periodically, that the sum of measurements reported over the serial interface may gradually depart from the registered energy data. This is due to the loss of information in the averaging process; the registered data is the most accurate. The averaged data is provided primarily as a convenience or for system control purposes; it is not intended for billing purposes.

### *Voltage and Current*

Voltage and current are measured in a similar fashion to power, using overlapping 1024-point measurements. In this case, however, the cross product is replaced with the square of the voltage or current samples. The square root of the resulting sum is proportional to the rms voltage or current value during the measurement interval. This value is corrected for the CT and/or PT correction factors before further use.

The resulting data is used to correct the energy measurements, as described above, and averaged over a one-second interval to provide data for display and reporting to a host system.

### *CT and PT Compensation*

CT and PT compensation may be enabled to correct for the inaccuracies of the CTs and PTs used in the metering setup. The system voltage is relatively constant, so the PT compensation factor is a single, complex (i.e., magnitude and phase, or real and imaginary) correction factor.

CT compensation is more complicated. Due to the fact that magnetizing currents in CTs are not exactly proportional to the load current, a matrix is used, allowing the entry of several different compensation factors measured at different current levels. The Model 1133A interpolates between the numbers in this table (also complex) to determine the correction factor to be used.

Correction for energy is performed using the (complex) product of the PT and CT factors. Correction for voltage or current is performed using the magnitude of the appropriate factor. Correction for phase angle is made using the phase of the appropriate component, i.e. the arctangent of the complex value. The actual calculation performed may be different than this description, due to computational considerations (a complex multiplication is far faster than a trigonometric operation such as an arctangent, for example); however, the end result will be as described.

### *Transformer Compensation*

There are two different types of transformer compensation. They are used to correct for the losses in a transformer when primary-side metering is used to meter the energy delivered to a customer at the secondary of the transformer.

Copper compensation is used to correct for the $I^2R$ losses in the transformer windings due primarily to their (non-zero) resistance. As you would expect, this effect is primarily active (resistive), although there may be minor reactive effects, and it is proportional to the current squared. This factor allows the user to correct for these losses. It is a complex factor, providing both watts and VARs correction, and is proportional to current squared; i.e. so many watts and VARs are to be subtracted from the registered amounts per ampere squared of load current.

Iron losses (also called core losses) are due to magnetizing currents (the small amount of current required to generate the flux in the core, which is unrelated to the load current) and eddy current losses in the core material. These are approximately proportional to the square of the voltage, and the compensation is performed using the same basic method as described above for copper loss.

### *DC Offsets*

DC offsets may be present in the signals applied to the input of the Model 1133A, although this is unusual. More commonly, small dc errors in the measurement circuits of the unit itself result in non-zero average of the samples.

Since the Model 1133A makes wideband measurements of power, voltage, and current (i.e., components at any frequency within the measurement bandwidth – which includes dc – will affect the measurement), this is a potential source of error which must be corrected to obtain maximum accuracy. Therefore, part of the measurement process is to average the (windowed) data, measuring the dc component. The effects of this component are then subtracted from the results.

### *Phase and Frequency*

As a part of its measurement process, the Model 1133A performs a fast Fourier transform (FFT) of the windowed voltage and current samples. In accordance with IEC 1000-4-7, this process is performed twenty times per second, using overlapping 1024-sample hanning-windowed data. This yields new FFT results twenty times per second for each voltage and current input, for a total of 120 FFTs per second. Phase angle may be determined from the relationship between the real and imaginary component of the fundamental-frequency bin of the FFT. (Since the window is 100 ms wide, each bin is 10 Hz apart; therefore, this is bin 5 for 50 Hz and bin 6 for 60 Hz.)

So long as there is significant measured energy in the bin, frequency offsets do not affect the measured phase angle. This is true as long as the signal being measured is the main source of energy in the bin, i.e. there is minimal leakage from adjacent bins, and minimal noise. Provided that the frequency is anywhere near nominal (within 10 Hz or so), the phase measurement is perfectly usable.

The phase measurements may be compared to determine phase angle between voltages and currents, or between any two voltages or currents. Because the sampling process is synchronized via GPS to UTC, absolute phase angle measurements may be made and compared between two units located at some distance from one another.

Frequency is measured by taking the difference in phase angle between subsequent measurements, based on the identity f = dØ/dt. Frequency is averaged over one second prior to being displayed or made available for output.

### Harmonics

Harmonics are measured using overlapping hanning-windowed FFTs of 1024 samples and 100 ms window length. Based on the instantaneous frequency, the location of the bins containing significant energy for each harmonic are determined. This is a total of three bins, one approximately centered on the harmonic and those two adjacent to it. Then, the energy in those three bins is totaled, resulting in the energy for that harmonic. This can then easily be expressed as a percent of the rms signal level, or in whatever form is required. While there is a closed-form correction which can be employed to find the harmonic magnitude in the presence of frequency errors, this approximation was chosen because it is much faster and gives adequate performance.

There is an error in this approximation, due to the fact that there will be a small amount of energy leakage into nearby bins, which will not be included in the three which are measured. This is generally of little consequence, as it turns out. First of all, when the frequency is close to accurate (which is most of the time), then the amount of energy outside of the three bins summed is very low. As a matter of fact, with the hanning window, there is no signal at all outside of these three bins if the frequency is exact.

In the real world, however, the frequency will be off somewhat, and it is reasonable to ask how large the error can be. For small frequency errors, say 0.01 Hz, the $50^{th}$ harmonic will be 0.5 Hz from the center of the nominal bin. This results in an error of about 0.005% – insignificant. The worst-case error will occur when a harmonic is very nearly centered between two bins. In this case, the algorithm described above will 'miss' a bin containing a signal with an amplitude of about 17% of the actual harmonic amplitude. The energy contained in this bin is then $(0.17)^2$ or about 2.9% of the total energy, resulting in a measured energy 0.971x what it should be. The measured harmonic amplitude will then be $(0.971)^{0.5}$ or about –1.5% in error. This is well within the specification limits (5%) of IEC 1000-4-7. This worst-case error would occur for the $50^{th}$ harmonic with a fundamental frequency error of 0.1Hz. At lower harmonics, the frequency error must be progressively greater; for example, to result in a –1.5% error in measuring the 9th harmonic would require a fundamental frequency offset of 0.556 Hz.

The phase angle of the harmonics can be determined by taking the arctangent of the real and imaginary components of the bin closest to each harmonic. This information cannot be used, however, in the averaging process described in IEC 1000-4-7, because this specification requires the rms average of a series of measurements. This, by definition, requires magnitude data only. Therefore, the harmonic phase is not normally calculated, but it can be requested or displayed. In this case, the result shown will be the phase angle determined by a single measurement, and it will be calculated 'on-demand,' i.e. when it is requested only.

### Flicker

Flicker is measured in accordance with IEC 1000-4-15, the successor standard to IEC 868. Unlike the other measurements described above, flicker measurement is a continuous process. This process is performed using a sample rate of 640 samples per second (sps). Anti-alias (decimation) filtering is performed on the 10240 sps data stream, and the resulting samples are further processed following the block diagram suggested in IEC 1000-4-15. The resulting measurements of flicker perceptibility are classified using a 256-level logarithmic classifier at the full 640 sps rate. $P_{st}$ is then determined each ten minutes as described by the standards.

Although no standards currently require it, the Model 1133A also measures flicker on the current inputs. This information can be useful in determining whether a customer's load is causing flicker on the power system, or whether the customer is being subjected to flicker from other sources. It is unrealistic to penalize a utility for 'poor power quality' at a customer's load when the cause of the problem is the load itself.

### Interruptions

Interruptions are monitored on the voltage inputs by comparing the 20/second voltage measurements with a user-supplied threshold. Events where the voltage dips below the threshold trigger an event, which can cause the logging of pre- and post-fault data, contact closure, or any of the other actions described under 'Event Logging.'

### Voltage Fluctuations

Voltage fluctuations are monitored by classifying the 20/second voltage data, per phase, with a 256-step linear classifier covering a range of ±20% of nominal voltage. These data are then summarized as a cumulative probability table over a specified interval, typically 15 minutes. In addition, the minimum, maximum, mean and standard deviation are calculated. The data may be recorded in flash memory either continuously or on demand.

### System Time Deviation

System time deviation, which is the accumulated error of a clock using the system frequency as its reference, compared to an absolute reference such as UTC, is determined from the 20/second phase data described earlier. System time deviation is accumulated as integer cycles of error plus fractional phase, and is converted to seconds as needed. The A-phase voltage phase angle is used for this measurement. Since this is an integrated value, the constant of integration (initial time offset value) must be specified by the user.

### Phasor Measurements

Phasor data are formatted and output in accordance with IEEE Standard 1344-1995. Phasors consist of the real and imaginary component of magnitude for the voltages and currents at a particular point in a power distribution system, along with suitable time synchronization fields and other information. This information is available in real time, and is based on the measured fundamental voltage, current, and phase angle described above, at a 20 records per second rate. There is a measurement delay due to the data acquisition delay of 50 ms, signal processing time of approximately 15 ms, and data transfer time which depends on the data rate.

### Phase Balance

The Model 1133A measures phase balance by calculating the symmetrical sequence components (positive, negative, and zero sequence) for the three-phase voltage input. Normally, if the unit is connected properly, the positive-sequence voltage will be equal to the line voltage and the negative- and zero-sequence voltages will be approximately zero. In the event that a user-specified limit on the imbalance (as a percentage of nominal voltage) is exceeded, an event will be recognized. In addition, the sequence components are averaged over a user-specified interval (typically 10 or 15 minutes) and may be logged if desired. These calculations are performed using the voltage magnitude and phase information, 20/second.

### *Load Balance*

Load balance is calculated in much the same way as phase balance, except on the customer's load current. Both of these measurements may be used to identify serious power system problems, such as a dropped phase, which could cause serious damage to both the utility's and the customer's equipment.

### *Flash Memory and Event Logging*

The Model 1133A includes 2 megabytes standard of flash memory for data and event logging. Up to 8 megabytes total may be optionally installed. This memory may be used to record two basic types of information: registered quantities, which are recorded on a fixed schedule (and may include many more items than a typical energy meter, including just about any function the instrument can measure); and event data. Event data is logged on the occurrence of an event designated to initiate logging.

These events may consist of a measured quantity exceeding a user-specified threshold, an external trigger, or an internal state of the Model 1133A (low flash memory, for example). In addition to causing the logging of specified data, an event may trigger a dial-up modem call to be initiated, reporting the event; and/or a contact may be closed (or opened).

There is a great deal of flexibility as to what may be recorded in flash memory at the time an event is recognized. The time, type of event, state of any number of measured quantities at the time of the measurement, pre-and post-event data, which may be measured quantities (voltage, power, etc.) or waveform data, may all be recorded. The Model 1133A must be configured in advance to specify the events being recognized and actions to be taken.

There is a fundamental difference between registered quantities and event log data. The amount of memory needed to record a certain number of registered items for a certain period can be determined exactly, whereas the amount of memory required to log events depends on the number and type of events, and the number of items to be recorded for each type. Since the number and type of events cannot be known *a priori*, the amount of memory required also cannot be stated with certainty. The Model 1133A handles this by allocating sufficient memory as required for the registered quantities over the specified period of time, and then making whatever memory is left available for event logging. You can configure the Model 1133A to initiate an auto-dial call or contact closure when the remaining event memory (or register memory) drops below a specified limit. Memory is allocated in 64Kbyte blocks, which is the block size of the flash memory.

Flash memory must be erased in blocks, and data cannot be over-written until its block is completely erased. Therefore, the normal process will be to first read out the desired data, and then erase the blocks, making them available for re-use. Each block of data is password-protected, having two levels of security: one to access the data, and a second to clear it from memory. The registered data may be separated into two blocks, each with its own security. This would normally be used to separate billing data from operational data.

In addition to the password security, the flash memory is located on a separate module, mounted internally to the Model 1133A. In the event of a failure or sabotage to the instrument itself, the memory module may be removed and read on a separate Model 1133A mainframe, preventing the loss of critical billing information and other data.

### *Contact Outputs*

Four contact outputs may be used to report events recognized by the Model 1133A, or they may be controlled remotely, by command. They may also be operated on a schedule, which may be downloaded for up to 30 days in advance. In addition to reporting events, these contacts may be used to synchronize external equipment or to operate load-control switchgear.

### *Event Inputs*

Four external event inputs are also provided. These are optically-isolated and accept dc signals at levels of 24 to 240 volts. Upon the application of a signal to one of these inputs, it will be time-tagged to one microsecond resolution and recognized as an event, which can cause any of the actions described above under 'Event Logging.'

### *Serial Channels*

The Model 1133A includes two serial channels. Each of these can be configured at the time of order with RS-232 or RS-485 drivers, or with a V.34bis 33.6k modem. Each channel can have access to all functions of the Model 1133A; alternatively, certain functions (such as the ability to clear the revenue registers) may be enabled or disabled independently. Each channel has an RJ11 (6-position) modular connector with 2 meter (7 ft.) modular interconnect cable , and the RS-232 and RS-485 functions are provided with an adapter to DB-9 (male) plug. A DB-9 female adapter, and longer cables, can be supplied on special order. These adapters may be rewired in the field to match interface requirements.

Both serial channels, and the Ethernet channel, may be operated simultaneously, each serving different hosts with whatever information each needs and is authorized to access.

### *Ethernet MMS Interface*

The Model 1133A contains a provision for an Ethernet interface supporting the emerging MMS standard for substation automation. This interface will have a 10BaseT interface via an RJ45 (8-position) modular connector. This interface will not be available initially, but the hardware is included, to allow future expandability when this capability is added. Contact the factory for more information on this feature.