



# EAGLE HARBOR TECHNOLOGIES

## User's Manual for Integrator Long Pulse ILP8

22AUG2016



International Distributors

**sky blue**  
microsystems

Sky Blue Microsystems GmbH  
Geisenhausenerstr. 18  
81379 Munich, Germany  
+49 89 780 2970, info@skyblue.de  
www.skyblue.de

**ZERIF**  
TECHNOLOGIES LTD.  
A SKY BLUE COMPANY, FOUNDED 1999

In Great Britain:  
**Zerif Technologies Ltd.**  
Winnington House, 2 Woodberry Grove  
Finchley, London N12 0DR  
+44 115 855 7883, info@zerif.co.uk  
www.zerif.co.uk

## Contents

Specifications .....	3
Packing List.....	4
System Description .....	5
RJ45 Channel Mapping.....	8
Customization .....	9
Channel-by-Channel Custom RC Times.....	9
Long Pulse Integrator Stitching Algorithm.....	10
Removing the DC Offset Error.....	10
Removing the Slope Error .....	11
Stitch the Data Segments Together .....	12
Dynamic Range Enhancement .....	14
Performance Data .....	17
Typical Performance .....	17
Frequency Response .....	19
Droop Characteristics.....	20
Drift Characteristics .....	21
Operating Instructions .....	24
Operation.....	24
Interfacing and Shielding .....	24
Calibration Testing .....	24
Grounding .....	26

## **Specifications**

Exceeding any of the Maximum Ratings and/or failing to follow any of the Warnings and/or Operating Instructions may result in damage to or failure of the integrator, or incorrect output signals.

Maximum continuous input voltage: 10 V<sup>a</sup>

Maximum pulse input voltage: 175 V<sup>a</sup>

Maximum output voltage: +/- 8 V<sup>b</sup>

Minimum output load: 50  $\Omega$

Input resistance to ground: 50 k $\Omega$ <sup>b</sup>

Provided external AC-DC converters input power: 110-240VAC, 50-60 Hz, 1.5 A each

RC Time: 1  $\mu$ s – 100 ms, custom set for each channel<sup>b</sup>

Droop: RC dependent, see figures

Drift: RC dependent, see figures

Calibration: +/- 10% for each channel<sup>c</sup>

Offset adjustment: < 5 mV

Common Mode Noise Rejection: < 1 mV/s (no effect detected) with 1 MHz +/- 20V

Channel - Channel Crosstalk: < 1 mV effect (no effect detected) when adjacent channel undergoes 5V excursion

Input impedance: 100 Ohms<sup>b</sup>

Input filtering to ground: none<sup>b</sup>

Input voltage clamp: none<sup>b</sup>

Input Connectors: RJ45, BNC<sup>b</sup>

Output Connectors: RJ45, BNC<sup>b</sup>

<sup>a</sup> The input has a 1 W 100  $\Omega$  termination resistor. The input voltage can exceed 10 V for short periods or at a low enough duty cycle so that the 1 W power rating of the resistor is not exceeded. The RJ45 input connector is limited to +/- 175V, so never exceed this input voltage.

<sup>b</sup> Integrators with different values, ranges, connectors, options are available.

<sup>c</sup> Precision factory calibration available. These can be corrected for in post processing.

### **Related Equipment:**

If longer operation and/or higher gain operation and/or higher effective bit depth is required, then EHT has the long pulse integrator (ILP8) that can run continuously, and meets ITER stability requirements

### **Warnings:**

Integrator Chassis does not tie to earth ground through input power connectors. See discussion below.

## **Packing List**

In the box, please find the following:

- 1) One ILP8 Long Pulse Integrator
- 2) Two 12V DC Adapters
- 3) One ILP8 User's Manual

## **System Description**

The EHT long pulse integrator (ILP8) uses an innovative patent protected technique to prevent drift error from growing exponentially in time as is common with other analog integrators, instead limiting it to a drift error that accumulates only linearly with time. The ILP8 can be used at extremely high gain (10  $\mu$ s RC or lower) and remain stable indefinitely, having been tested for up to 72 hours of continuous operation while maintaining drift < 1 mV/s over that entire time period.

The integrator can be operated continuously, but must be periodically “re-zeroed”. This re-zeroing process takes approximately 1  $\mu$ s, during which data is lost. Because the integrator is periodically re-zeroed, the data that is output is a sequential set of data segments, each starting at 0V, which must then be “stitched” back together in software by the user. The re-zeroing of the integrator is controlled by the Clock input and the frequency is user-controlled. The re-zeroing frequency can be set depending on the application, but a frequency of 1 Hz is sufficient to keep the integrator stable.

The ILP8 has two inputs to control the function of the integrator. When the Enable input is in the low state, all channels are in reset mode, which automatically zeroes their outputs. When the Enable input is in the high state, all channels are in the active mode, and their output will correspond to the time-integral of the input. Specifically, the output of an ideal integrator is given by:

$$V_{OUT} = \frac{1}{RC} \int_0^T V(t) dt$$

RC is the integration time constant of that channel. The RC time constant can be independently set for each channel with a physical resistor. While the ENABLE input is high, the integrator can be rapidly re-zeroed by applying a clock signal with 50% duty cycle to the CLK IN. The integrator re-zeros on the rising and falling edge of the clock signal.

Integrators have several types of error on the output that must be considered: drift, droop, offset, and calibration, and slope.

Drift is an output voltage that strays away from zero over time, even when the input voltage is held at zero. It results from phenomena such as contact potentials, thermal gradients, inherent instabilities of operational amplifiers, etc. All analog integrators suffer from drift error to some extent and the goal with integrator design is to ensure that drift error is small compared to the signal of interest over the relevant timescale. Generally, drift error in analog integrators results from a runaway instability of some sort, meaning that drift becomes exponentially worse over time, limiting the maximum time that an integrator can run before the drift becomes excessive. In the case of the ILP8, the drift process is restarted each time that the integrator is re-zeroed, meaning that it never enters the exponential “run-away” phase and so grows only linearly over time over many integration periods.

Droop is a tendency of the output to droop back towards zero after integrating a real input signal. An ideal integrator would hold its output voltage constant forever after integrating a specified input signal, whereas a real integrator’s output will droop. The droop of the ILP8 is dependent on the

RC value, presented in a graph in the data section. In general, droop can be corrected for in post-processing.

Offset error is simply a DC offset in the output of an integrator. In the case of the ILP8, this has been adjusted to be less than 2 mV for each channel, through use of an internal pot. This offset may change slightly with temperature. The offset error can be fully corrected for in post processing.

Calibration error is a scaling factor in the output voltage equation above. This can be considered simply an error in the RC time constant. For example, a 10 ms RC integrator will have an output of 1 V after applying a 1 V input to it for 10 ms. If the calibration is only good to within +/-10%, then the output could be anywhere from 0.9 V to 1.1 V. The calibration error can be fully corrected for in post processing, as long as the user has a reliable calibration signal.

Slope error is an underlying baseline slope that is output from the integrator even if the input signal is zero. This error is not always present in all integrators, however, it is present in the ILP8 due to the nature of the re-zeroing circuitry. This slope can be fully corrected for in post-processing.



Figure 1. PLACEHOLDER

The front panel (Figure 1) contains the BNC and RJ45 outputs, the EN and CLK inputs, and the status LEDs.

- 1) PWR ON LED  
The AC ON LED is on whenever the ILP8 is plugged into power and the switch on the back panel is on.
- 2) EN LED  
The EN (enable) LED is on whenever a high signal (5V) is given to the ENABLE input. If the LED is on, all integrator channels are in active mode.
- 3) CLK LED  
The CLK LED cycles on and off with the CLK IN.
- 4) ENABLE  
The ENABLE input is used to trigger the integrator channels to active mode. The device used to control the ENABLE input must be able to drive a 50  $\Omega$  load to 5 VDC.
- 5) CLK IN  
The CLK IN is the clock signal that rapidly re-zeros the integrator. The clock signal should have 50% duty cycle since the re-zeroing occurs on the rising and falling edge.

6) BNC Outputs 1-8  
There is a separate BNC output for each channel.

7) RJ45 Outputs  
There are two RJ45 outputs, one for channels 1-4, and one for channels 5-8.




Figure 2. ILP8 Back Panel PLACEHOLDER

The back panel (Figure 2) contains the power inputs and power switch as well as the signal inputs.

1) Power On/Off Switch

2) Input Power Connectors

Both power connectors must always be connected to the integrator. Never turn the power switch on without first connecting both connectors. Only use the power supplies included with the ILP8.

3) BNC Inputs 1-8

4) RJ45 Inputs

There are two RJ45 inputs, one for channels 1-4, and one for channels 5-8.

## RJ45 Channel Mapping

Each RJ45 input on the back panel and output on the front panel provides connectivity to 4 channels, as labeled on the front and back panels. The channel mapping matches the T568A standard. Each differential input signal comes in on one set of individually shielded twisted pair. The channel mapping is illustrated in [Figure 3](#).

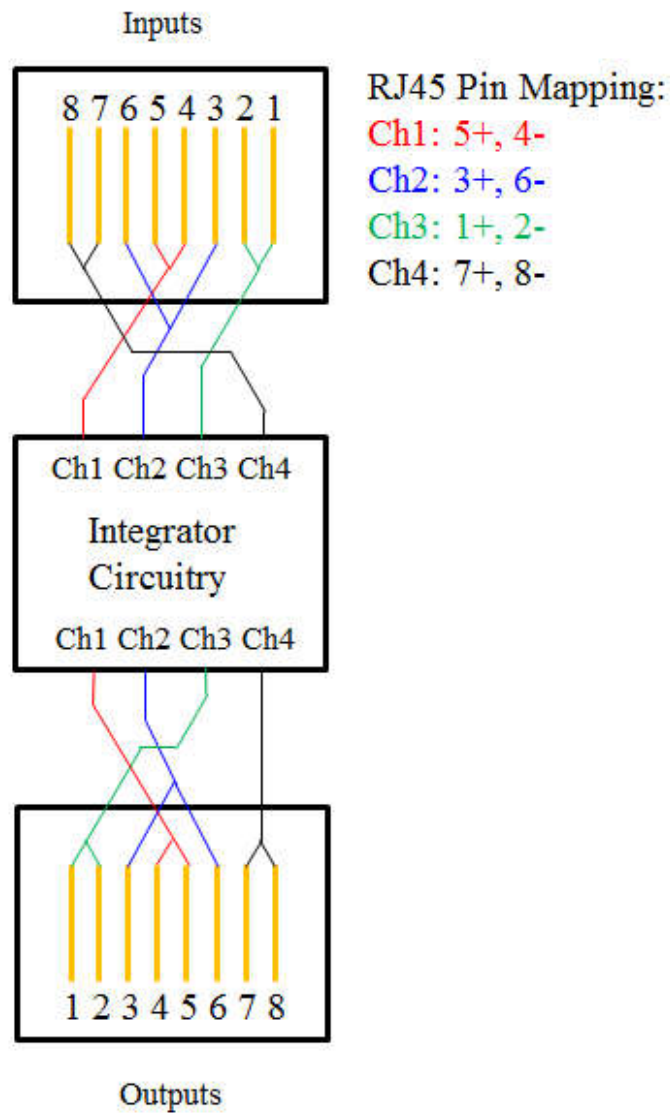


Figure 3: RJ45 Channel Mapping

This repeats for each RJ45 connector.



## **Customization**

The ILP8 integrator is available with a wide range of RC time constants, where each channel can be individually set. Typical RC times range from 1  $\mu$ s to 100 ms, though the integrators have been operated with RC values outside this range as well. Selection of a small RC value is typically done when a very high gain is desired, and for short pulse experiments. Many other features of the ILP8 can be customized as well, including input and output connectors and input terminations.

### **Channel-by-Channel Custom RC Times**

This unit was customized with the following RC times:

**Placeholder: Specify on a case by case basis**

## Long Pulse Integrator Stitching Algorithm

The output of the long pulse integrator consists of individual segments of integrated data, which are re-zeroed at a specified frequency. To reconstruct the overall waveform, the user must “stitch” this data back together in software. This section will describe the suggested algorithm for reconstructing the data. This algorithm consists of several steps, as follows:

- Remove any DC offset error from the data
- Remove any systematic slope error (“slanted baseline”) from the data
- Stitch the segments of the data end to end to form a continuous data set

The integrator re-zeroing is done by means of an external clock input. Each time the clock transitions from high to low or vice-versa, the integrator is re-zeroed and begins a new integration period. Due to the internal workings of the integrator, the DC offset error and slope error may each be different when the clock is high compared to when it is low. We will refer to these two states as “HI” and “LOW”.

To accurately calculate and remove the above listed errors, the user should turn on the integrator for many integration periods prior to the signal of interest. This sample, prior to the signal of interest, will be referred to as the “pre-trigger”.

We will consider the example waveform below, where 9 individual integration periods are shown. Assume that there are additional integration periods preceding this data, where no real signal is present.




Figure 4: Sample raw data from integrator

### **Removing the DC Offset Error**

For the data in the pre-trigger, calculate the average of the first points of each integration period for both the HI state and the LOW state. So for example, if it starts in the LOW state, then take the average of the first data point in integration period 1, 3, 5, 7; this will be the offset error for the LOW state. Similarly, do the same for the HI state using periods 2, 4, 6 and 8.

Then, for each data point for all segments in the LOW state, subtract the LOW state offset error. And for each data point for all segments in the HI state, subtract the HI state offset error. This should result in the data below, where the offsets have been eliminated.




Figure 5: Sample data with DC offsets removed

### Removing the Slope Error

As with the offset error, any slope error (“slanted baseline”) may vary between the HI and LOW state. For the pre-trigger data, where no external signal is being fed into the integrator, calculate the slope of each segment. Group these into the slopes for all segments in the HI state and all segments in the LOW state, and take the average for each, so that there is an average slope for the LOW state and an average slope for the HI state.

Then, going through all the segments of data, subtract away the slope from each data point. This should result in the slope-corrected data below:




Figure 6: Sample data with slope errors removed

### Stitch the Data Segments Together

The data segments must now be joined end to end to construct a continuous data set. The simplest way to do this would be to take the value of the final data point in a segment, and add that value to each data point for the next segment, repeating through all the segments from the beginning to the end. However, this can introduce substantial error, since if there is a real signal across the transition, then the last data point of one segment and the first data point of the next segment should not necessarily be at the same value.

To account for this, the best way is to take a linear (or higher order, if desired) fit through the last  $N$  data points of set  $n$  (call this fit1), and also take a linear fit through the first  $N$  data points of set  $n+1$  (call this fit2). Then, calculate the offset that should be added to all data points in set  $n+1$  so that the fit1 and fit2 intersect at the point in time half way in between the last data point of set  $n$  and the first data point of set  $n+1$ .

$N$  should be chosen with respect to the typical timescale of the signals the user is expecting, it should include enough data points to capture the trend of the signal before and after each transition and not be thrown off by DAQ noise from point to point, but few enough data points that it is also not thrown off by features of the signal far away from the transition.

Stitching the data together segment by segment, the result should be the reconstructed data set below:




Figure 7: Sample data stitched back together

## **Dynamic Range Enhancement**

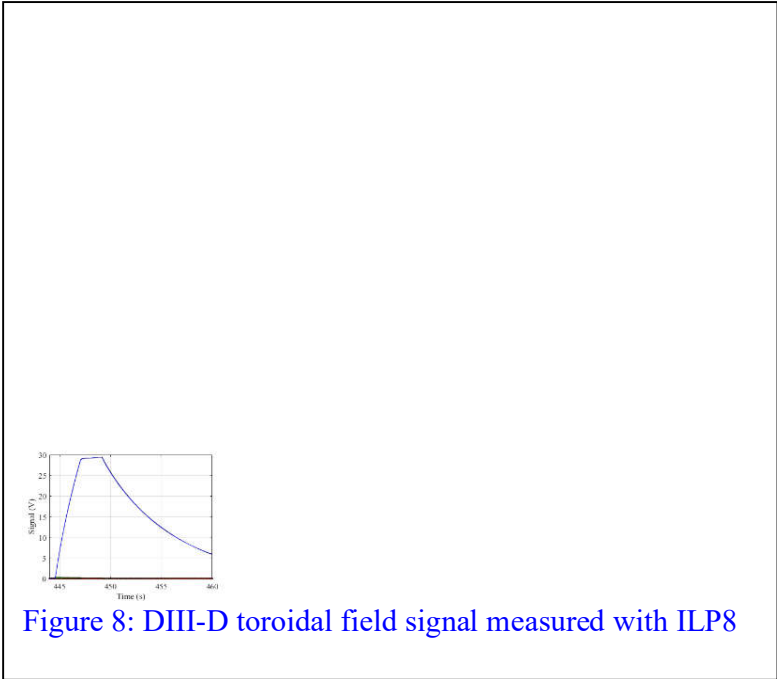
The ILP8 must be used together with a user provided DAQ of some specified bit depth. For example, 8-14 bit DAQs are typical depending on the application. If the user was using the ISP16 or any other standard analog integrator, then this DAQ's bit depth would be the absolute limit for signal resolution in all cases. However, using the ILP8, the dynamic range (bit depth) can actually be enhanced for most typical signals.

The RC time of the ILP8 should be selected so that it fills the DAQ input during one integration period (between re-zeroing transitions) during the time of the signals maximum  $dV/dt$ . If the signal spans multiple re-zeroing periods, then the DAQ is filled multiple times. For example, consider a typical toroidal field signal measured in a fusion experiment, which ramps up over many seconds, then comes to a flat top, then ramps down again over many seconds. If the ILP8 is switched at a frequency of 1 Hz and the RC is selected so that the DAQ is filled during the ramp-up in each integration period, then the dynamic range can be enhanced by a factor of several times compared to an integrator which would have to capture the entire waveform in a single fill of the DAQ.

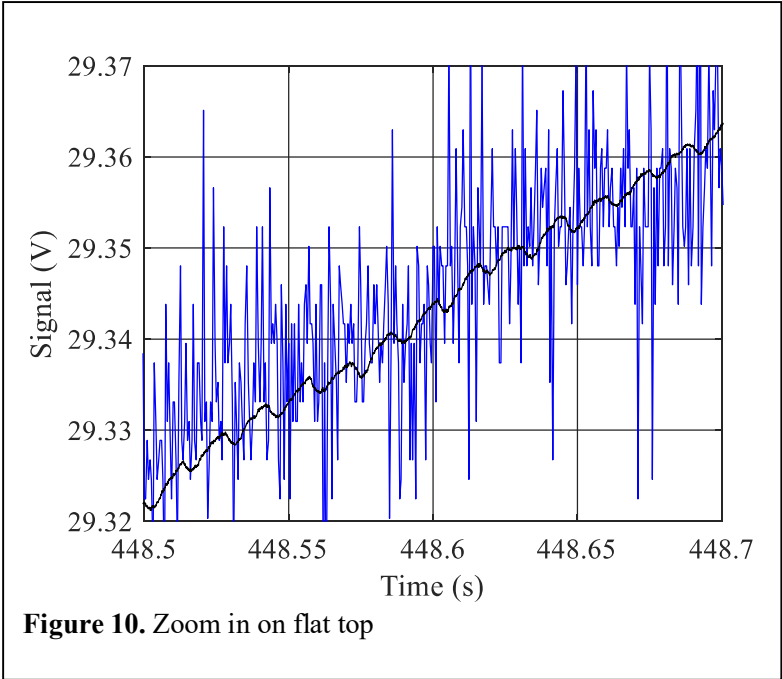
The ILP8 can be re-zeroed at frequencies of up to approximately 100 Hz, meaning that for signals longer than several tens of milliseconds, significant improvements in dynamic range can be achieved. However, this improvement relies on the signal the property that fast signal spikes do not saturate the integrator/DAQ. So signals with long smooth ramps are ideal candidates for dynamic range enhancement, where the higher dynamic range could be used to resolve smaller features of interest riding on the ramps. However, signals dominated by large delta functions would get no benefit.

An ideal example of a signal that benefits from dynamic range enhancement is the toroidal field measurement in a large tokamak. The ILP8 was tested on the DIII-D Tokamak, where the toroidal field ramps up and down over the course of several seconds, meaning that several hundred re-zeroing periods occur during the ramp up and ramp down, allowing a dynamic range enhancement of several hundred times.

Figure 8 shows the stitched, integrated toroidal field signal from a shot of the DIII-D tokamak. In Figure 9, we show the individual segments of output from the ILP8 integrator during the 2 seconds of the toroidal field ramp-up. As we can see, stitching all of these together results in a signal that is of a much higher amplitude than could have otherwise been recorded by the  $\pm 1V$  DAQ.



Zooming in on the flat top of the signal compared to a conventional reference integrator in Figure 10, we can see the significantly enhanced dynamic range.





## Performance Data

### Typical Performance

Testing for various inputs was conducted to make sure the integrator output was as expected. In [Figures 4 and 5](#), short ( $\sim 10 \mu\text{s}$ ) and long ( $\sim 50 \text{ ms}$ ) square wave signals (blue) were input into the integrator, respectively, resulting in the linear ramps shown (purple).




Figure 4. 10  $\mu\text{s}$  Ramp




Figure 5. 50 ms Ramp

The next example waveform in [Figure 6](#) was generated by connecting a magnetic pickup coil into the integrator input, and simply placing a magnet into the coil, waiting a brief time, and then pulling the magnet back out. As expected, the signal returns back to zero.

[Figure 7](#) demonstrates both low and higher frequency performance of the integrator. The 6 s long main signal was generated with a magnetic pickup loop (as in [Figure 11](#)), with the high frequency signal shown in the blow-up view injected through a transformer-coupled function generator. Note

that the integrator simultaneously fully resolves both signals. The higher frequency signal shown is for demonstration only and is still not at the maximum frequency of the integrator.




Figure 6. Magnetic Pickup Loop Signal




Figure 7. Frequency Range Demonstration

**ADD LONG PULSE AND HDR SAMPLE DATA**

## Frequency Response

The maximum frequency of an input signal that can be accurately integrated in terms of both gain and phase on the output is approximately 1 MHz.

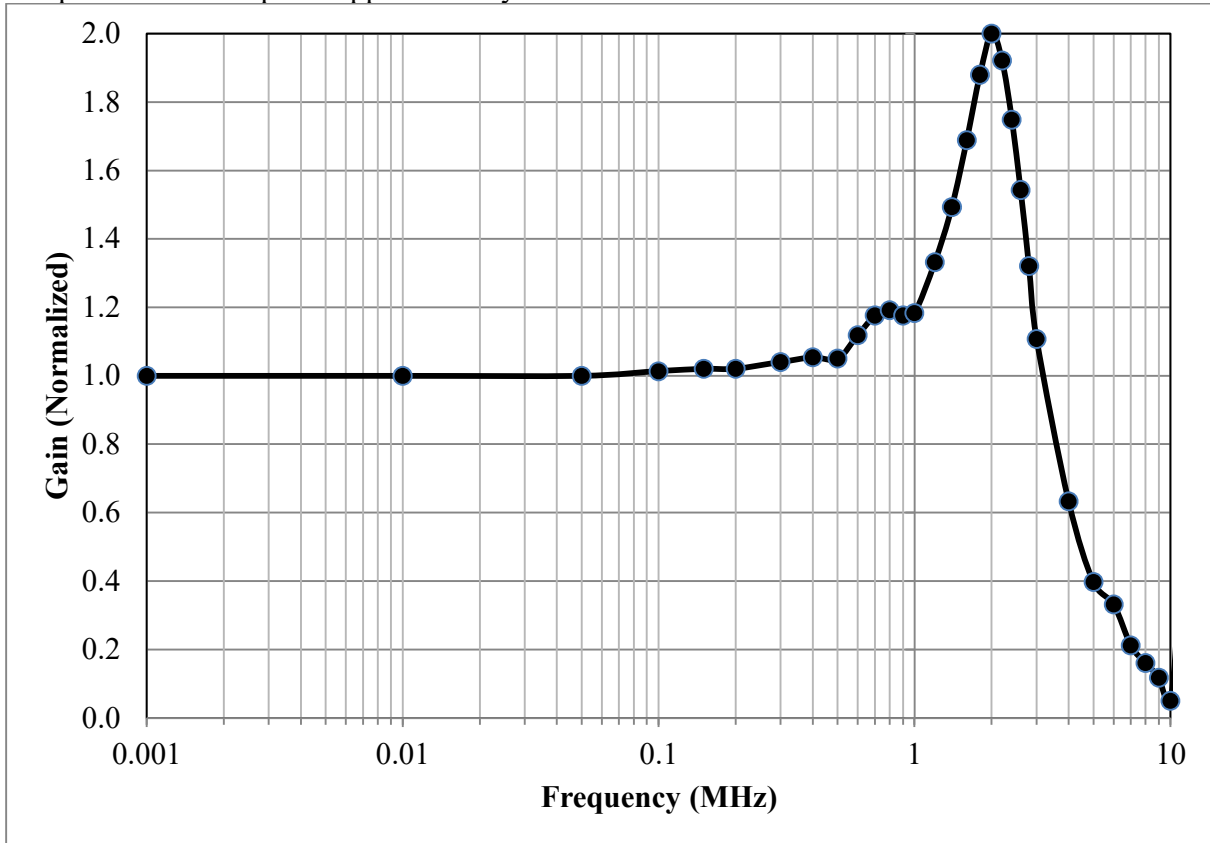


Figure 8. Frequency Response

Faster signals can still be integrated but may include a phase lag as well as incorrect gain and settling/overshoot features. For example, the waveform below is of a 200 ns input square wave with a fast rise and fall (frequency components up to ~15 MHz) into a 500 ns RC time integrator:




Figure 9. Fast Waveform Response

## Droop Characteristics

The droop within a single integration period is dependent on the RC time. The droop rate is approximately inversely proportional to the RC time. It should be noted that the droop of a signal only applies within that integration period. Because the integrator is re-zeroed and starts integrating again after each re-zeroing transition, signal gathered prior to that period can no longer droop. The droop data observed within individual integration periods is presented below.




Figure 10. Droop vs RC Time

The droop was measured by integrating a large signal and then letting the integrator remain on for a long time until the droop became measurable. For example, here is the droop waveform of the 1 ms integrator, gated on for 100s:

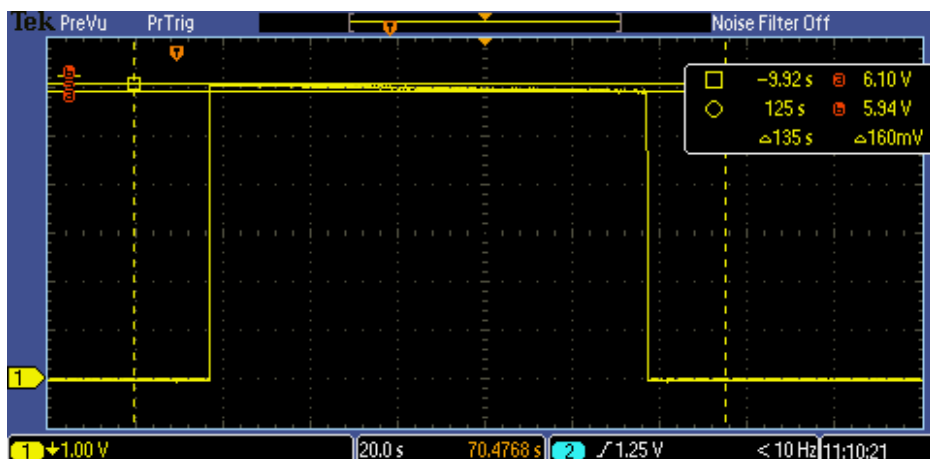


Figure 11. Droop Test Waveform

## Drift Characteristics

Drifts were measured by looking at the typical random drift of all channels at any particular RC time and picking a characteristic upper bound on their drift. For example, here are typical drifts from a set of 4 different  $10\mu\text{s}$  RC channels, being gated on for 1 second and then being allowed to reset for 1 second, and repeating several times. The upper bound in this case is characterized as 5 mV of drift although most instances of the drift fall well below that. This drift data is for individual integration periods.


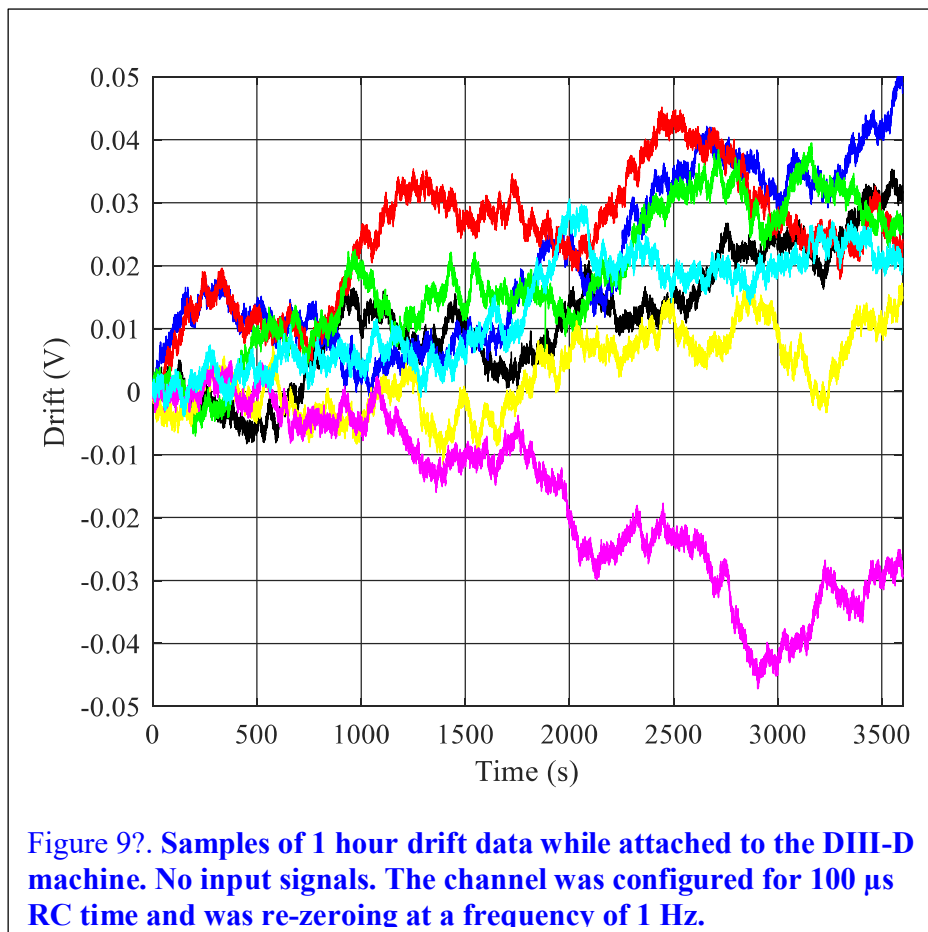
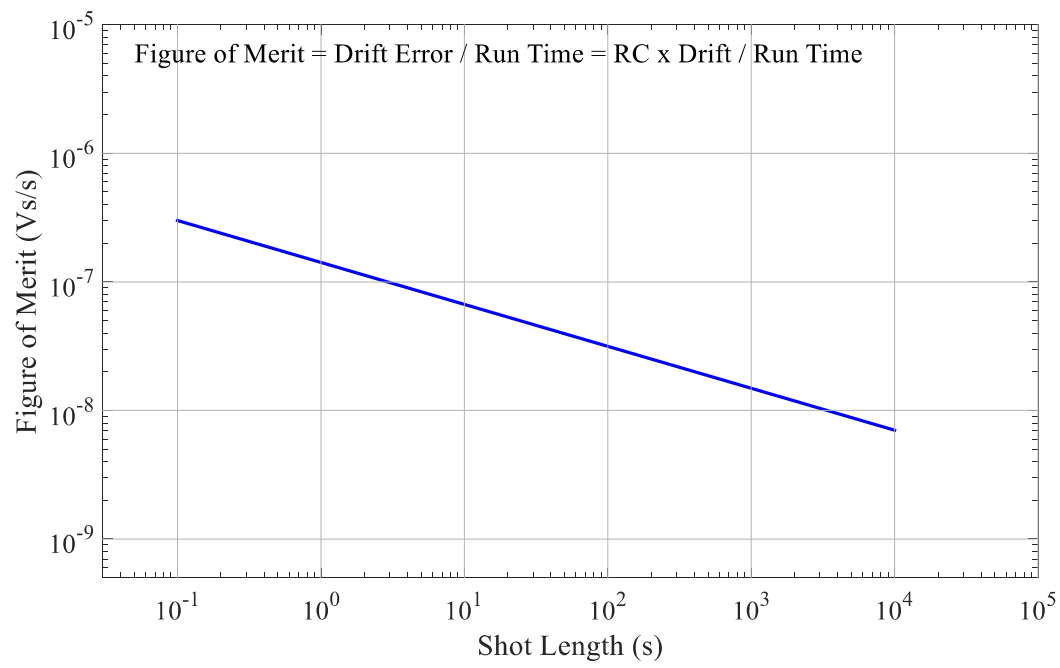


Figure 13. Drift Waveforms

The drift performance over a 1-hour period when stitching together multiple integration periods is shown in the next figure. This was gathered with an ILP8 channel configured for  $100\mu\text{s}$  RC time.



The performance of an integrator can be described in terms of a figure of merit,  $F = \Delta V \cdot RC / T$ , where  $\Delta V$  is the drift,  $RC$  is the RC time of the integrator, and  $T$  is the time period over which the drift occurs. The drift in each integration period is random, and so when many of these periods are added up, the drift grows with the number of periods  $N$  as a random walk problem, so scaling as  $N^{1/2}$ . This means that the figure of merit of the ILP8 is actually better the longer it is operated.



## **Operating Instructions**

### **Operation**

Operation of the ILP8 is relatively straightforward. Attach inputs to the input connectors on the back. Attach outputs to the output connectors on the front, going into the user's DAQ system or oscilloscope.

Gate the integrator on (at the EN input) prior to the signal of interest to be integrated. This should be done far enough in advance to have several tens of integration periods prior to the signal of interest, which will generate data to be used in removing offset and slope errors as described in the Stitching Algorithm section above. With the integrator turned on at the EN input, use a 5 V clock signal with 50% duty cycle capable of driving 50  $\Omega$  into CLK IN BNC. The integrator will re-zero on the rising and falling edge of the clock signal.

Note:

- 1) It is important that no external signals enter the integrator during the pre-trigger period. Because the data from this period is used to compute the slope and offset errors present in the integrator, if there are any real signals superimposed on these then the results will be off.
- 2) Do not exceed 5 VDC into the EN or CLK gate input.
- 3) If not gated on, the integrator will still appear to integrate large/fast signals, though these signals will exhibit large droop. This is the result of internal circuitry trying to zero the output while the integrator is in reset mode.

### **Interfacing and Shielding**

The ILP8 has 8 separate individual inputs, one for each channel. The channels are bundled into groups of 4, each associated with an RJ45 input and output. The signals should be brought in through Cat 6e/7 cables, which have independent electrostatic shielding and where the two signal leads for each channel are individually shielded twisted pairs. This style of cable maintains the differential input of the integrator, which is important for optimal results. The BNC inputs can also be used, however, due to the non-differential nature of BNC cables, this may cause a degradation in integrator performance.

Care should be taken to maintain the differential nature and electrostatic shielding of the input signal as far as possible, ideally right up to (and over) the magnetic pickup loop being used with that channel. Sections of un-shielded and/or un-twisted wire on the input will allow noise to enter the system.

### **Calibration Testing**



Prior to use in an experimental setting, the user should verify the calibration of the integrator, in regards to both DC offset and calibration error.

To check for DC offset, simply power the integrator on, do not connect any inputs, do not turn on the ENABLE input, and look at the output voltage of each channel with an oscilloscope or DAQ. The output voltage of each channel should be 0 V +/- 5 mV. If the DC offset voltage exceeds +/- 20 mV, internal pots should be adjusted to bring the offset back under +/- 5 mV. Contact EHT prior to making this adjustment for detailed instructions.

Next, the user should look at the calibration of each channel relative to the input signal. The user should input a specified voltage signal for a specific amount of time into the integrator and look at the output into the DAQ. An example of such an input square wave is shown in the blue in [Figure 15](#). The purple trace is the resulting output. The ILP8 integrates the square wave input and outputs a linear ramp.

An input square wave of voltage  $V_{in}$  for duration  $T$  should produce a linearly ramping output voltage, ending with the final value of  $V_{out} = V_{in}T/RC$ .  $V_{out}$  may vary by a few % from this value due to variation in precise values of resistors and capacitors used in the integrator, etc. Production units can be precisely calibrated at the factory if desired.

The user should record the overall calibration value for each channel and use it to correct the integrator's output data.

When performing calibration testing, the user should take care to ensure that whatever system is used to input signals into the integrator does not have a slight DC offset, as many standardly used signal/pulse generators typically do. A small DC offset in the input voltage will be integrated as a real input signal by the integrator. Getting rid of a DC offset coming out of a signal generator can often be achieved by transformer coupling the signal, though than transformer droop can become an issue. Other solutions involve using a signal generator that allows the DC offset to be precisely tuned so it can be dialed to zero, or using a photodiode to optically couple in a signal. These integrators are very sensitive, and readily integrate input DC offset signals.




Figure 15. Input Square Wave (Blue), Integrator Output Ramp (Purple)

## Grounding

It is a feature of the ILP8 that the chassis is not tied to the input power ground. The user may tie the chassis to the rack that the ISP is mounted in, or to the input of a DAQ that the ILP8 is being used with. Always avoid creating ground loops when using the ILP8. Ideally, the chassis would only be grounded at the DAQ, and the DAQ would have differential inputs. In addition, all input signals would be differential with no ground connections, and have isolated electrostatic shields that only ground at the integrator chassis. If the input signals are single sided, there is a real possibility that the integrator will respond to induced ground currents.

One effect of ground loops between the ILP8 and the DAQ can be a perceived crosstalk between the channels. Consider a voltage  $V_{\text{sig}}$  being input into the DAQ along a BNC cable from the ILP8, as shown in Figure 16. If the BNC cable is terminated with a resistance  $R_T = 50 \Omega$ , then a fairly high current  $I = V_{\text{sig}} / 50 \Omega$  will flow, and must then return back to the integrator through the ground shield of the BNC cable. The BNC cable has some real resistance  $R$ . If a second channel is also connected through a BNC cable, then half the current will flow back through that cable's ground shield, and will produce an offset voltage  $V_{\text{off}} = I/2 \times R$ .

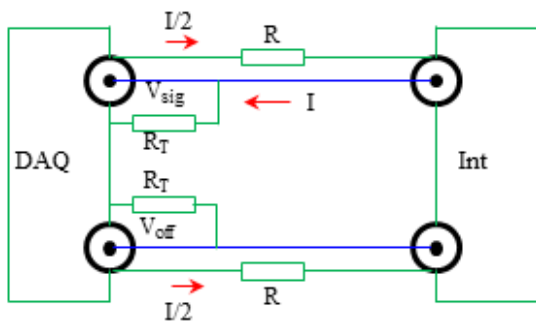


Figure 16. Ground Loop Resistance Issues

This effect is shown in Figure 17, where a 7 V signal on one channel (purple) causes a  $\sim 2$  mV error on another channel (yellow). The perceived crosstalk can be reduced by increasing the value of the termination resistor  $R_T$ , which will reduce the current that flows and therefore the error voltage. The perceived crosstalk can also be reduced by minimizing the resistance of the cables connecting the integrator to the DAQ (i.e. by keeping the cables short).

However, the optimal solution is to use a DAQ with differential inputs that do not tie to ground, thus keeping each channel isolated from all others, avoiding the issue of currents returning through the ground shields of other channels. In Figure 18, the results using a differential probe are shown. Notice that the perceived crosstalk has been eliminated. No real crosstalk is detectable.




Figure 17. Channel – Channel Perceived Response due to Ground Loop




Figure 18. No Channel – Channel Perceived Crosstalk with Differential Probe

International Distributors

**sky blue**  
microsystems

Sky Blue Microsystems GmbH  
Geisenhausenerstr. 18  
81379 Munich, Germany  
+49 89 780 2970, info@skyblue.de  
www.skyblue.de

**ZERIF**  
TECHNOLOGIES LTD.  
A SKY BLUE COMPANY, FOUNDED 1999

In Great Britain:  
Zerif Technologies Ltd.  
Winnington House, 2 Woodberry Grove  
Finchley, London N12 0DR  
+44 115 855 7883, info@zerif.co.uk  
www.zerif.co.uk