AN001: PID intensity-stabilisation with the XRF

This Application Note details the operation of an acousto-optic modulator (AOM) for intensity stabilisation using a MOGLabs *Agile RF Synthesizer* (XRF). The Note discusses practical implementation details and limitations and provides advice on control loop optimisation. It is divided into the following sections,

1. Apparatus configuration
2. Fundamental limitations on loop bandwidth
3. Analog signal-processing
4. Initial loop configuration
5. Loop optimisation
6. Additional details regarding noise suppression

The instructions presented recommend the XRF to be updated to **firmware v1.11 or newer**, which is available from www.moglabs.com.

1. Apparatus

A typical apparatus design is shown in Fig. 1. The XRF drives an AOM, and the diffracted beam is blocked, for example with an iris. An optical wedge (or beam-splitter) is used to pick off a monitor beam which is measured on photodetector PD1. The measured optical power is processed by signal conditioning electronics and fed into the AMP modulation input. Optionally a second photodiode (PD2) is used to independently measure the noise level on the output beam (see §7).

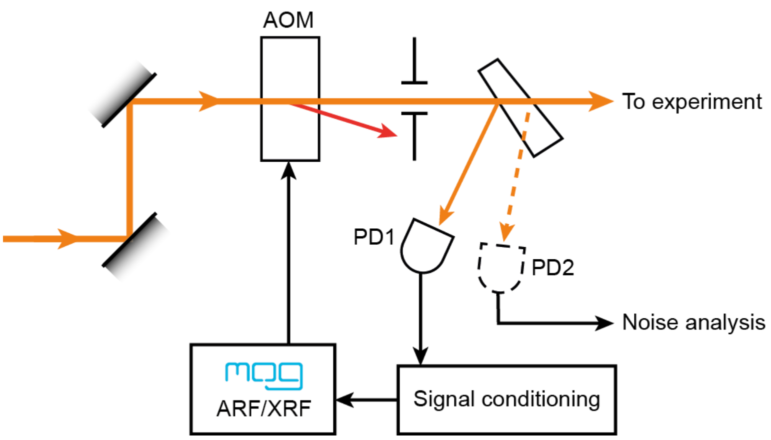

Figure 1: Typical apparatus diagram for PID intensity control with an XRF.

MOG Laboratories Pty Ltd
Suites 34-35, 49 University St
Carlton VIC 3053
Australia
Tel: +61 3 9939 0677
info@moglabs.com

Equivalently, the diffracted beam could be used instead of the undiffracted beam, if the frequency-shifted output is desirable. It is also possible to use the discarded beam as the basis of the measurement (Fig. 2), where an increase in the photodiode measurement is used to infer a decrease in the experiment beam power. This approach requires fewer optical components, but is more susceptible to additional scattering effects within the AOM and is generally less successful with high-power laser beams.

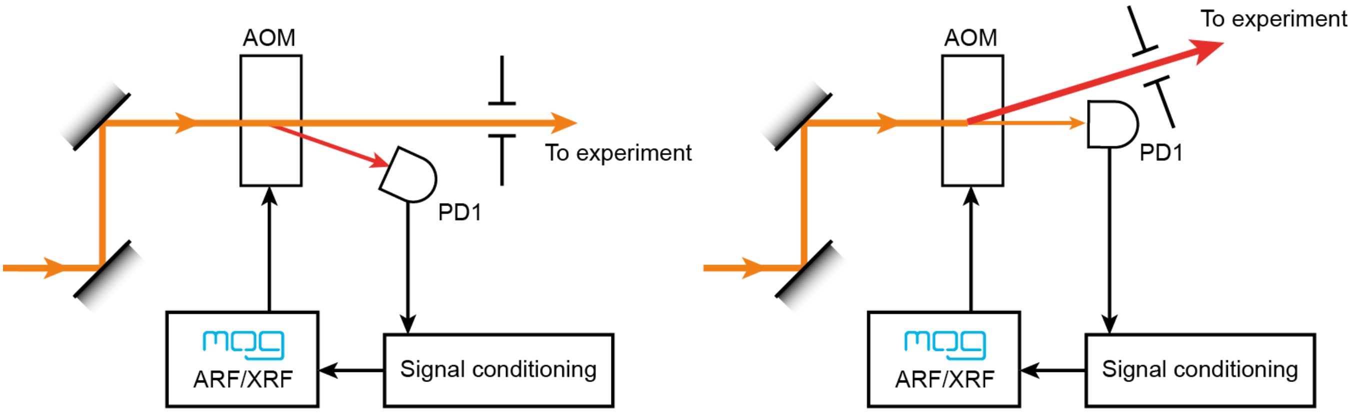

Figure 2: Alternate apparatus configurations for PID intensity control with an “inverted” error signal.

2. Limitations on closed-loop bandwidth

The primary limitation of a closed-loop servo is the “lag”, “group delay”, or “propagation time” between detection of a fluctuation and the actuator’s response. Fundamentally, the servo **cannot respond to any changes that occur on a faster timescale than this delay**, resulting in an inherent bandwidth limitation. For AOM noise-eaters, the main contributions are signal processing delay and AOM response time.

The AOM response time is not often quoted by manufacturers and typically needs to be measured in-situ. Often the response time is **5-10x the rise time** (Fig. 3). These values are geometric considerations related to the speed of sound in the crystal and are generally **unrelated to the AOM modulation bandwidth**, which does not take lag into account. Typically, the AOM aperture is in the centre of the crystal to minimize edge-effects, at the expense of increased response time. In some cases, it may be possible to move the beam closer to the transducer to minimize this time.

The other contributing factor is the time taken to process the photodetector signal into an rf output. This signal processing chain in the XRF is comprised of photodetector → signal conditioning → ADC → FPGA → DDS → amplifier → AOM. Each step in the chain adds a nonzero delay, which is unrelated to the quoted bandwidth of each stage. For an XRF in *fast* modulation mode, the net delay between the modulation input and rf output is about 500ns, whereas in *slow* modulation mode it is about 3us.

In practice, the action of the servo at a given frequency is related to the phase shift caused by the propagation delay at that frequency. The effective servo gain is reduced as the phase shift increases, with the servo failing to suppress noise once the phase reaches 90°, and *increasing* the noise up to 180°, which is typically referred to as the “servo bump” or “Bode bump” in the noise spectrum. Therefore, **in the ideal case**, the effective noise-suppression bandwidth is $1/4\tau$, where τ is the net delay of the system including the response time of the AOM.

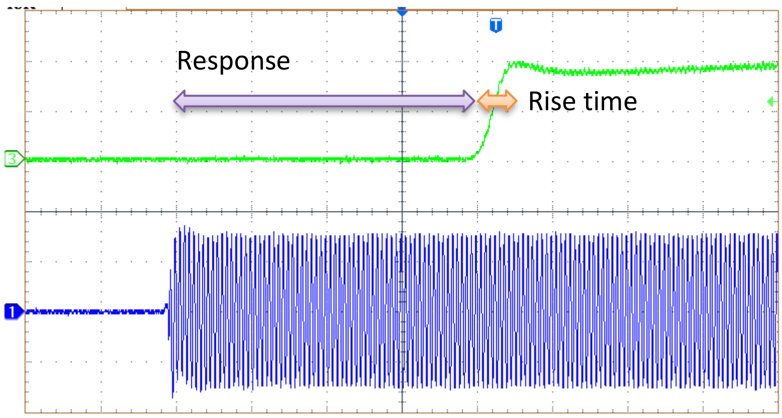

Figure 3: Comparison of diffracted beam power (green) in response to a step change in driving power (blue) for a typical AOM. The response time (purple) is 800ns even though the rise time (orange) is only 80ns.

In practice, the achievable PID bandwidth is **typically ~100kHz**. The individual bandwidths of the photodetector, signal-processing board, AOM, and XRF modulation inputs are typically orders of magnitude higher than this, which can give users an unrealistic expectation of the practical achievable bandwidth.

3. Signal conditioning

To achieve optimum intensity stabilisation, it is necessary to *condition* the photodiode signal by shifting the DC level to zero and then applying gain to amplify the residual fluctuations. The goal of the servo is usually to be able to suppress weak fluctuations on a large DC offset, so this conditioning provides the best matching of the input signal to the resolution of the ADCs.

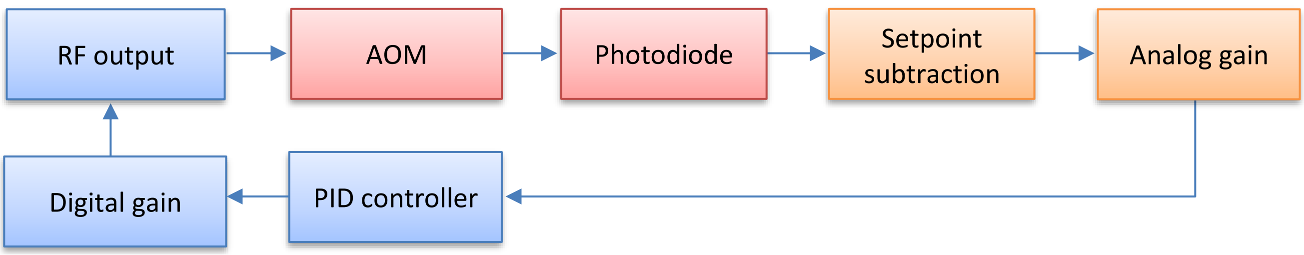

It is important to distinguish between analog gain and digital gain for the servo controller. The analog gain can be thought of as adjusting the *input sensitivity* of the control loop whereas the digital gain adjusts the *output range* of the servo. The digital gain should be adjusted to represent the largest correction that the PID controller is permitted to make, to prevent the controller from drifting too far or overreacting to perturbations.

Starting with serial numbers A09xxx, the XRF contains integrated analog signal processing which can be adjusted from software. Older units require the signal processing to be done externally (see Appendix A).

4. Initial configuration for intensity stabilisation

The following procedure is recommended for configuring intensity-stabilisation using the XRF's built-in signal processing feature. Equivalent steps can be carried out for an external signal processing board, or consult an earlier revision of this application note.

1. Connect the photodetector to the ***AM*** SMA connector of the XRF.
2. Adjust the RF power on the AOM to achieve the desired diffracted power.
3. Measure the *Setpoint voltage,* which is the DC photodetector voltage measured at this power.

   It is recommended to adjust the optical power on the photodetector until the voltage is in the range 0.1 - 2V.
4. Connect to the XRF using ***mogrf*** and configure modulation via *Settings > Modulation*
   1. Set the Amplitude modulation gain to 20%
   2. Set the Proportional gain, KP to 0.0
   3. Set the Integral gain, KI to 0.1
   4. Set the Derivative gain, KD to 0.0
   5. Set the Setpoint to the desired DC photodiode level (measured above)
   6. Set the Analog Gain to 0 dB
   7. Set the Sample rate[1] to 7.81 MHz
   8. Ensure PID stabilisation is set to "Amplitude"
5. Click the *PID debug* button to display the integrated diagnostic feature
6. Click the *Single* button to capture the behaviour of the PID controller as the loop is engaged. If the *PID input* diverges away from zero then *Invert PID action* must be toggled (Fig. 6)
7. If the *PID input* starts converging towards zero but saturates at a nonzero value, confirm that the setpoint voltage can be reached by adjusting the RF power on the AOM (Fig. 7).

   For example, if optical alignment drifts or coupling efficiency has dropped, it may no longer be possible to achieve the desired setpoint voltage. Remeasure the setpoint voltage above.
8. Increase the analog gain until the desired noise suppression is achieved, ensuring the gain is not increased so high that the controller becomes unstable and oscillates.
9. If required, further optimise the loop parameters to maximise the noise suppression using the *PID debug* feature as discussed in the next section, or by measuring the performance on a DC-coupled spectrum analyser.

[1] Increasing the sample rate increases the responsivity of the loop but also increases the integrator action. The time constant associated with the sample rate must not exceed the overall loop group delay.

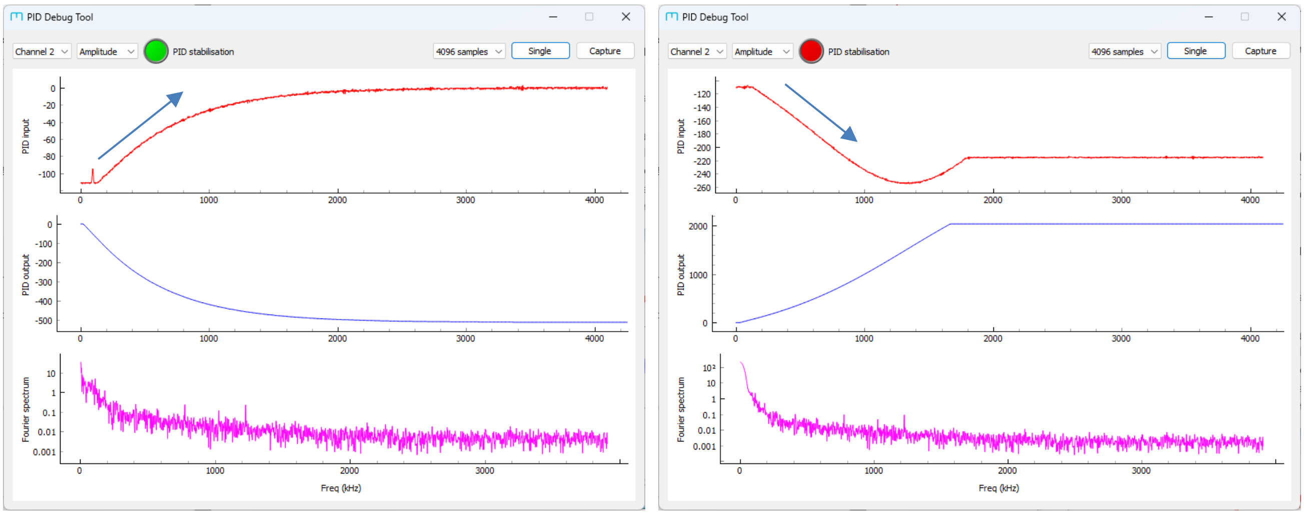

Figure 4: When the PID sign is correct (left) the *PID input* converges to zero. When the sign is incorrect (right) the *PID input* does not converge to zero, the *PID output* diverges to the limit value ±2048, and the indicator turns red.

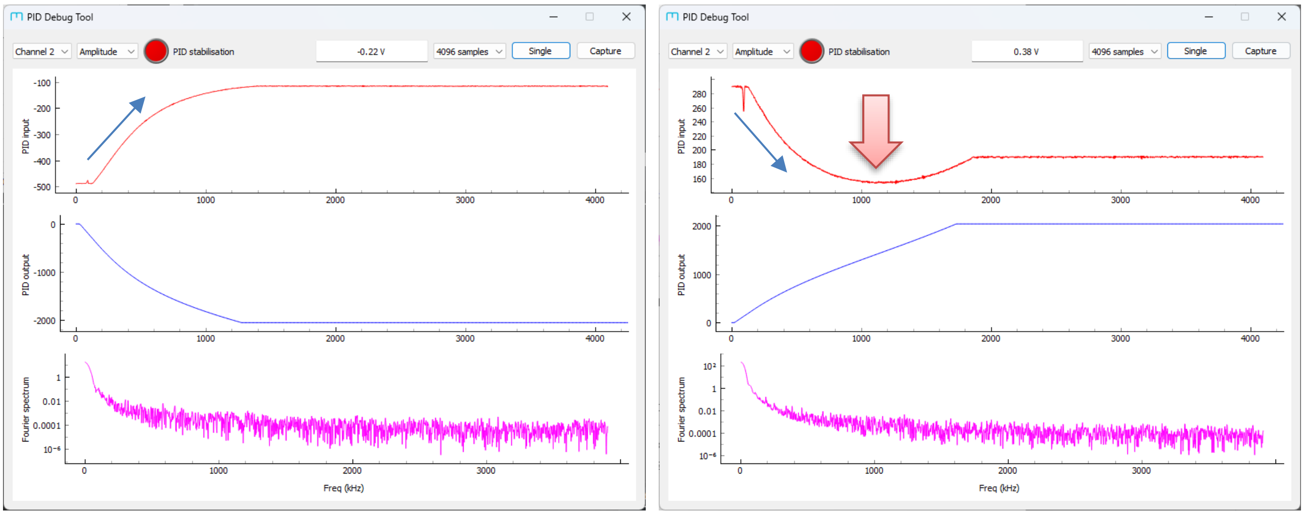

Figure 5: When the PID is unable to achieve the setpoint, the *PID input* <u>initially</u> heads towards zero but converges to some nonzero value (left). If the *PID input* graph shows an extremum despite *PID output* smoothly increasing (right), this implies the AOM cannot achieve the desired photodiode voltage at <u>any</u> RF power. This is likely due to drift in the apparatus alignment and/or input beam power.

5. Loop optimisation and analysis using *mogrf*

The two most common problems in optimising the loop are saturation and oscillation. Saturation occurs when the integrator reaches its limit and cannot continue to correct for drift in the apparatus. This typically occurs when the **digital gain is too low**, the response of the system is too slow, or the diffraction efficiency has changed too much. This can be seen in the debug window by the *PID output* reaching a value of ±2048 and is indicated by the red indicator LED (Fig. 7). The saturation can also be checked with the command **PID,VALUE,<ch>**, which returns a value in the range ±1.0 and may be beneficial to monitor in custom control systems.

Conversely, oscillation occurs when the overall gain is too high and the PID controller over-compensates for fluctuations. Note that the overall gain is described by the product of the digital gain and analog gain. Typically increasing the analog gain improves noise suppression but also makes the loop vulnerable to oscillation. The Fourier spectrum display of the debug window makes the onset of oscillation easy to identify (Fig. 6).

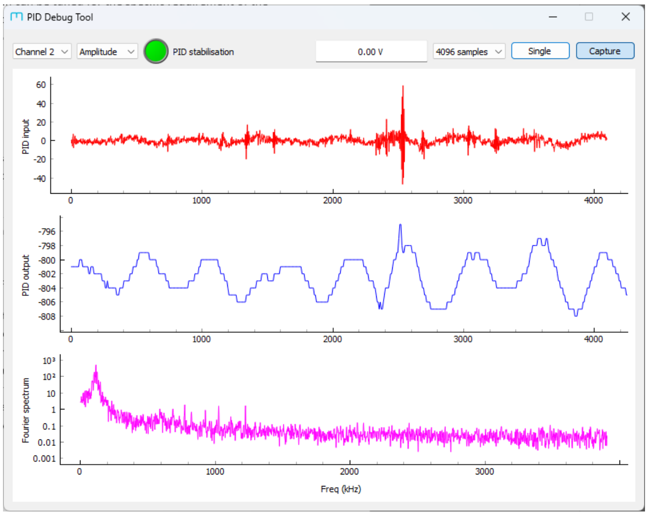

Figure 6: The *PID debug* window of ***mogrf*** showing a stable closed loop, where the *PID input* is centred on zero and the output has stabilised without saturating. The Fourier spectrum indicates the closed loop bandwidth is around 120 kHz, and the presence of clear oscillation indicates the loop gain is too high.

Once the maximum gain without oscillation is obtained, the analog gain can typically be increased further by reducing the digital gain. However, there is a trade-off between oscillation and saturation since reducing the digital gain reduces the maximum perturbation (typically thermal drift) that can be corrected by the PID servo. So one approach is to estimate the maximum controller deviation for normal conditions and set the digital gain appropriate for that, then adjust the analog gain to achieve optimum noise suppression.

The control behaviour described thus far is the "pure integrator" configuration, where only KI is nonzero. This is the simplest system to optimise because there are fewer free parameters, but the control system may exhibit suboptimal convergence when initially locking or when changing the setpoint voltage by overshooting (Fig 7). This is generally only of concern for "pulse shaping" applications, as discussed later.

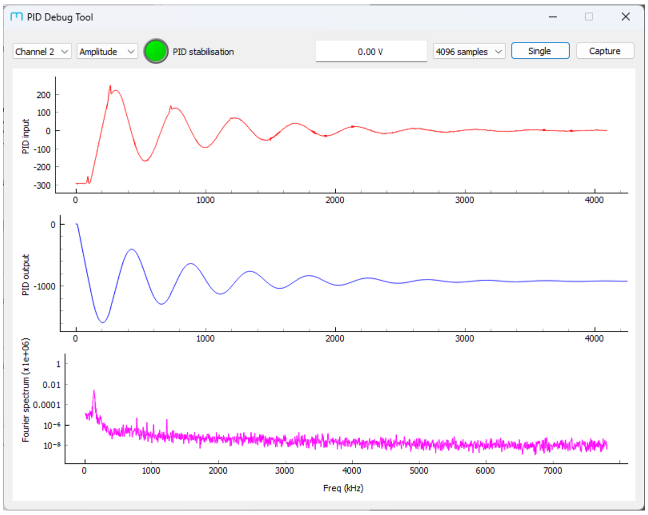

Figure 7: When engaging the lock, the PID controller overshoots the setpoint before settling down, suggesting that the gain should be reduced, despite not oscillating in steady-state.

The XRF implements a complete PID algorithm which can be tuned for the specific requirement of the experiment. In particular, when the PID is expected to follow changes to the setpoint voltage it may be preferable to switch to a "strong proportional" regime with weak integrator that corrects for droop. It is strongly recommended to consistently check both the steady-state performance as well as the convergence behaviour (using the *Single* capture option) when optimising the PID gains.

6. Further notes on intensity stabilisation

The underlying assumption of the control loop is that the measured photodetector signal is proportional to the power in the primary beam. It is therefore important to ensure that the photodetector does not detect scatter or diffuse light from elsewhere in the experiment. For example, an unshielded photodetector will measure fluctuations from fluorescent room lighting and apply an opposite modulation to the laser beam. Similarly, small amounts of scatter from other surfaces and other laser beams will incorrectly modulate the beam. This is especially prevalent with high power lasers, where optical attenuators (such as neutral density filters) and irises on the photodetector are strongly recommended. Ground loops can also result in uncorrelated noise in the error signal at the mains frequency, effectively ***adding*** noise to the laser.

It is also recommended that noise-sensitive applications use an identical independent ("out of loop") photodetector to quantify noise-levels. A well-tuned intensity stabilisation loop may be able to stabilise the beam down to the electronic noise floor, in apparent violation of the shot-noise limit. Although this seems impossible, the PID loop only ***suppresses the shot-noise in the measurement arm***. The beam used for the experiment is split off elsewhere and therefore contains ***uncorrelated shot-noise*** in accordance with the shot-noise limit. Use of a pick-off mirror and independent photodetector in the experiment arm will provide a more accurate measurement of the true noise-level of the beam (Fig. 8). Typically 20dB of suppression can be achieved on the out-of-loop detector.

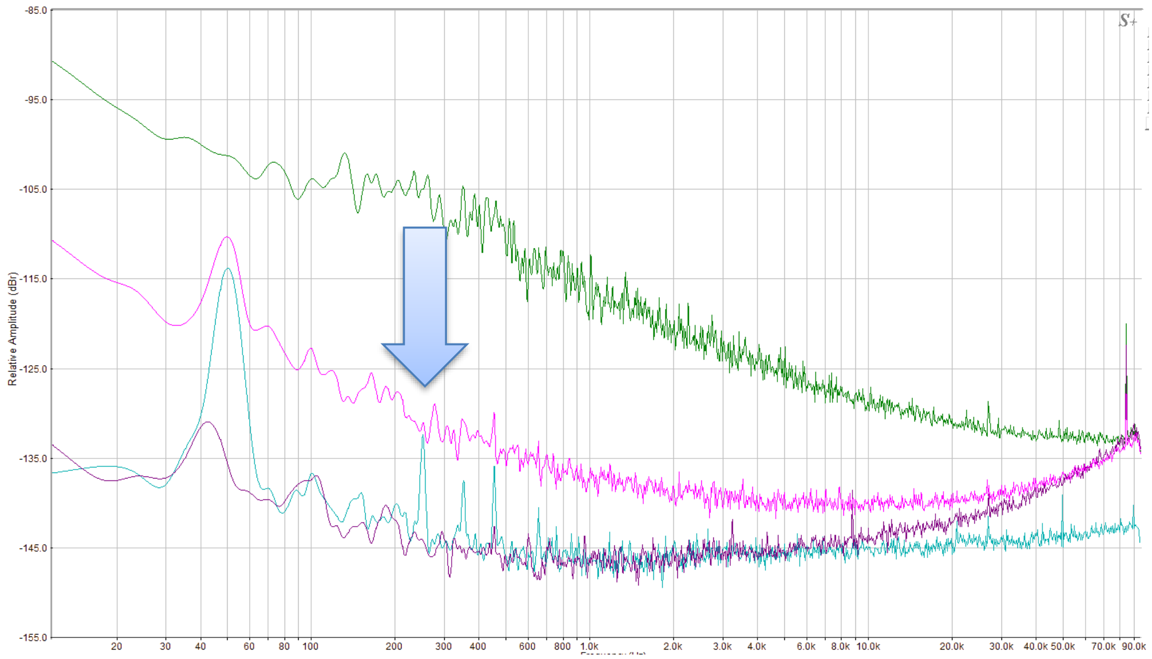

Figure 8: Typical measured noise spectra in an intensity-stabilisation experiment. Upon engaging the control loop, the incident beam noise (dark green) becomes suppressed (dark purple) into the dark noise (cyan), in apparent violation of the shot-noise limit. An independent (“out of loop”) detector more accurately describes the intensity noise on the output beam and obeys the shot-noise limit (magenta). Noise suppression is typically observed up to 100kHz.

The direct relationship between rf power and diffracted beam optical power also means that any noise features present in the rf signal get demodulated by the AOM directly into intensity noise in the output beam. Compared to photon shot-noise, even noise levels that are typically not characterised by most rf amplifiers can be observed in the diffracted beam intensity noise spectrum. The XRF contains custom low-noise power-amplifiers that were developed to minimise such noise coupling, but it is still important to shield the XRF and AOM from external noise sources (especially nearby high-voltage drivers).

Appendix A: Optional external signal-processing (B3122)

Older model ARF/XRF units require an external signal processing circuit to match the photodiode signal to the modulation input. This may also be necessary for scenarios where the desired laser intensity is varying with time - such as slow envelope shaping using an external analog signal.

The B3122 is a PCB that performs signal conditioning for use in intensity-stabilisation (noise-eating) applications. It is powered by a single +12V DC jack input (centre-positive), and multiple boards can be stacked together to make use of the same supply connector[2].

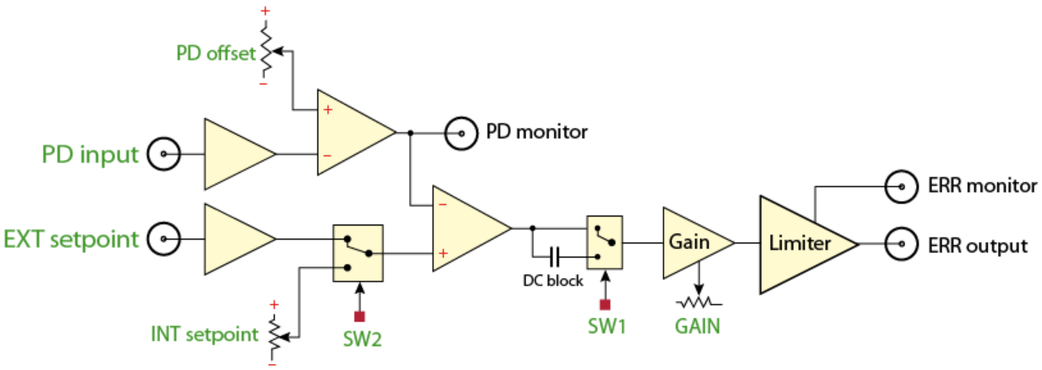

Figure 9: Simplified schematic of the B3122 signal processing chain. The inputs are ±6V tolerant, the PD offset is 0-5V and the setpoint has a ±1.25V range. The output limiter prevents the output voltage damaging the XRF's modulation input. The PD monitor can be used for noise analysis.

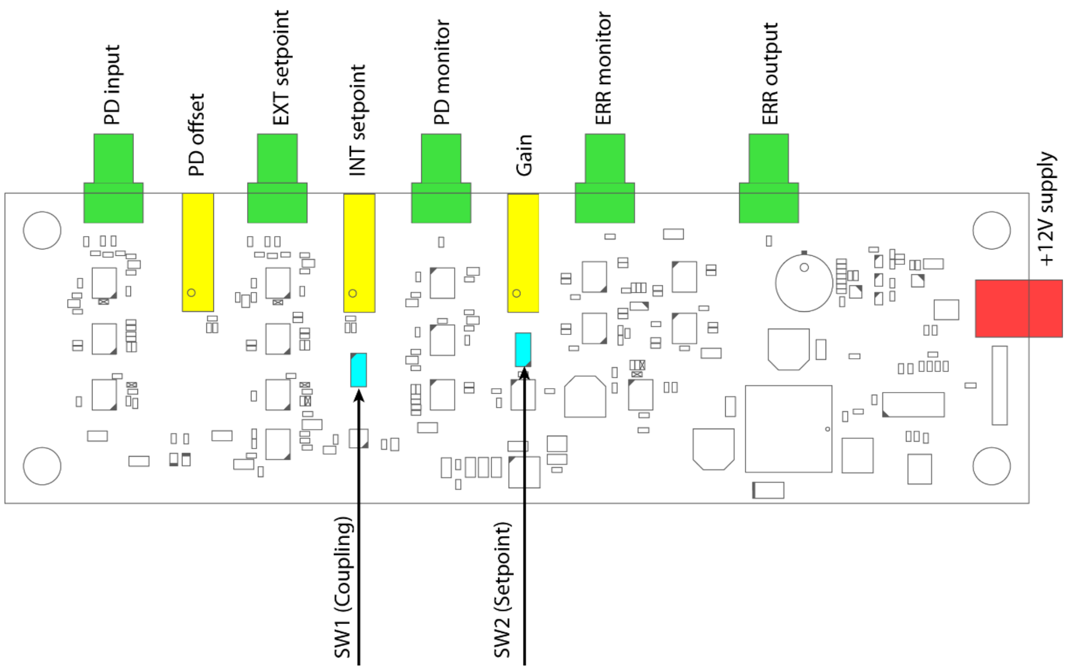

Figure 10: The B3122 signal processing board showing the SMA connectors (green), switches (aqua) and adjustment trimpots (yellow).

[2] Contact MOGLabs support for more information.

The board is configured for either *INTERNAL* or *EXTERNAL* setpoint mode using ***SW2***. When set to *INT* mode the setpoint is controlled by the integrated trimpot, whereas in *EXT* mode the setpoint is supplied as an analog voltage to the associated BNC input. The ***ERR output*** BNC should be connected to the relevant modulation input of the XRF (typically the AM input).

For most applications the coupling switch ***SW1*** should be set to DC mode. The ***PD offset*** trimpot is provided to correct for small DC offsets in the photodiode voltage, and the ***PD monitor*** output is provided for performance evaluation via noise spectrum analysis.

Appendix B: Pulse-shaping application

In some applications, it is desirable to control the laser intensity in a time-dependent way, for example to generate a specific pulse shape. The non-linear response of the AOM means that performing direct AM of the rf waveform generally results in undesirable pulse shapes (Fig. 11). In principle the AOM response can be calibrated and the modulation signal corrected for this ("feed-forward" control). However, this calibration is typically sensitive to heating and alignment perturbations, often leading to poor pulse shapes. An alternate approach is to use the PID controller to compensate for the non-linearities on the fly ("feedback" control) by using a time-varying setpoint voltage[3] (Fig. 12).

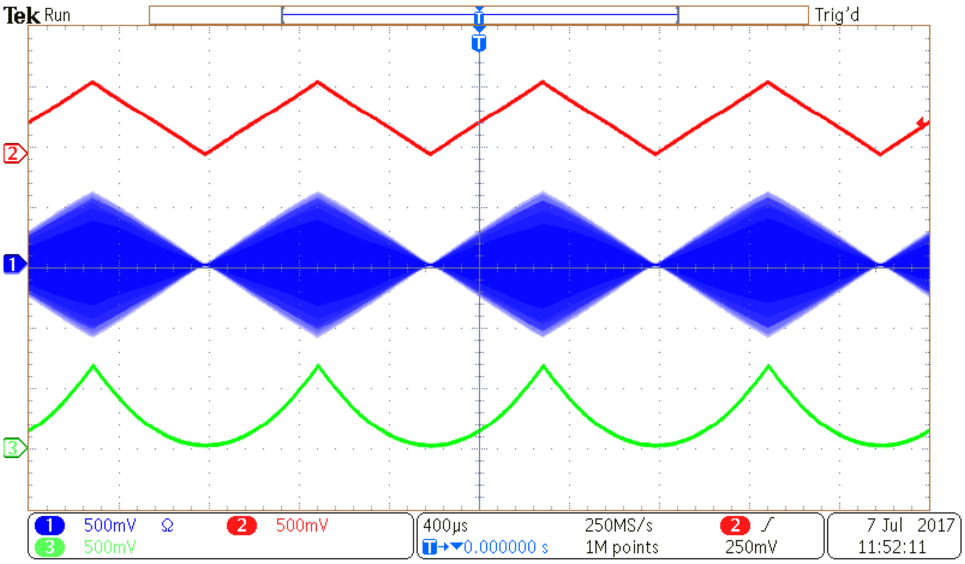

Figure 11: The output of a function generator (red) is used to amplitude-modulate the rf output (blue), which results in non-linear change in the beam intensity measured on a photodiode (green).

[3] This requires an external signal processing circuit, such as the B3122 described in Appendix A

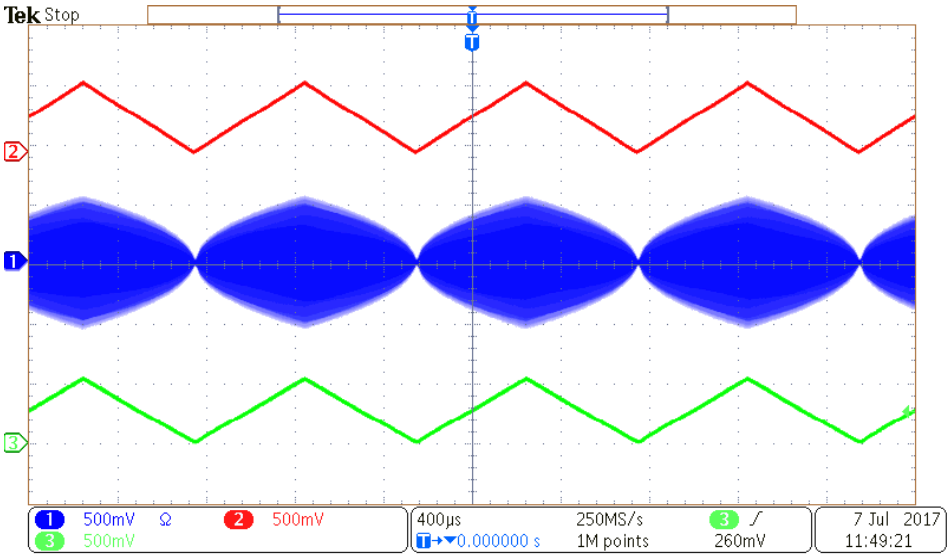

Figure 12: Using a closed-loop PID servo, the diffracted beam intensity (green) follows the desired form (red). The PID servo accounts for non-linearity of the AOM as seen in the rf envelope (blue).

In comparison to the noise-eater application, the AM gain typically needs to be higher as the controller needs to have wide enough action to follow the changes in setpoint. Furthermore, the combination of changing setpoint and signal processing delay means the integrator is especially susceptible to saturation. In this scenario it is recommended to lower both the PID sample rate and the integrator gain KI such that the proportional action gives rapid response to the changes in setpoint, and the residual integrator compensates for droop and non-linearity.

Furthermore, it is important that if the intensity needs to be driven to zero, that this is done in combination with the XRF's CHx-OFF TTL control feature. The PID controller can be used to smoothly reduce the intensity in accordance with the desired pulse shape, but the integrator will continue to accumulate small DC offsets while the beam is "off". Switching the beam off entirely using the TTL input eliminates any residual diffracted light and also holds the integrator in reset. This also improves the response when bringing the beam back up to non-zero intensity.