

AN001: PID intensity-stabilisation with the ARF

This Application Note details the operation of an acousto-optic modulator (AOM) for intensity stabilisation using a MOGLabs *Agile RF Synthesizer* (ARF). 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
7. Application to pulse shaping

The instructions presented recommend the ARF to be updated to **firmware v1.8 or newer**, which is available from www.moglabs.com.

1. Apparatus

A typical apparatus design is shown in Fig. 1. The ARF 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 photodiode 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).

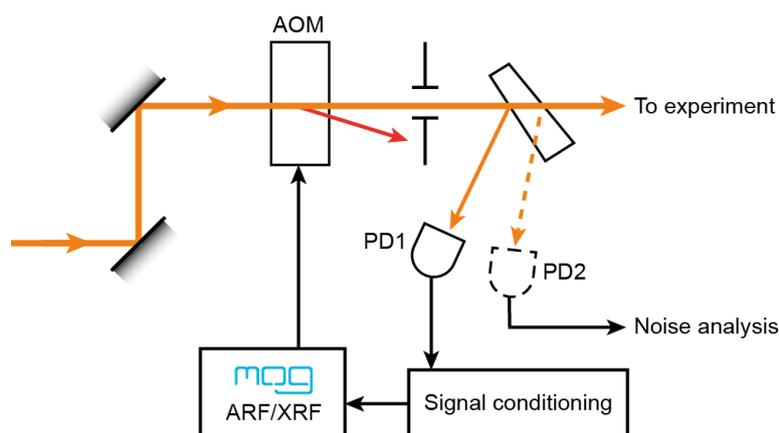


Figure 1: Typical apparatus diagram for PID intensity control with an ARF.

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.

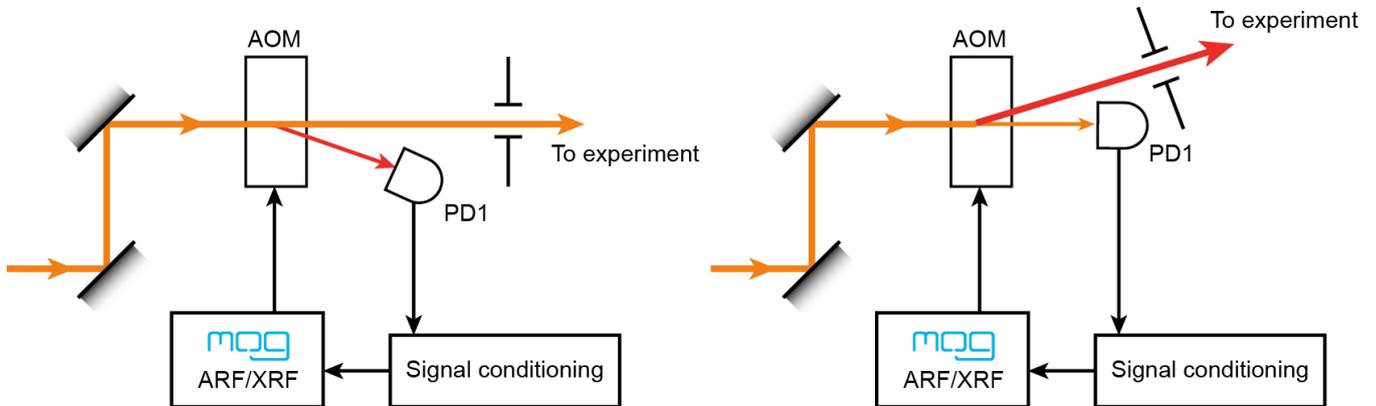


Figure 2: Alternate apparatus configurations for PID intensity control with an “inverted” error signal.

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 ARF 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 ARF 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 3 μ s.

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.

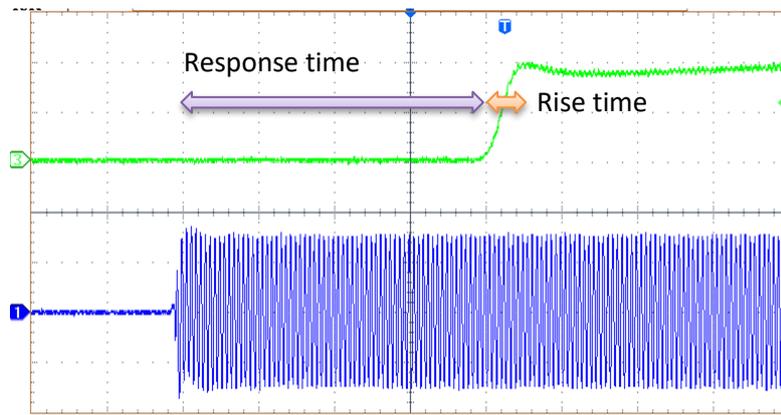


Figure 3: Comparison of diffracted beam power (green) in response to a step change in driving power (blue) for a typical AOM. The response time (purple) is 800ns even though the rise time (orange) is only 80ns.

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

3. Signal processing board (B3121)

The B3121 board (Fig. 4) is used to generate an error signal for PID control of rf amplitude for intensity stabilisation (Fig. 5). 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 connector¹.

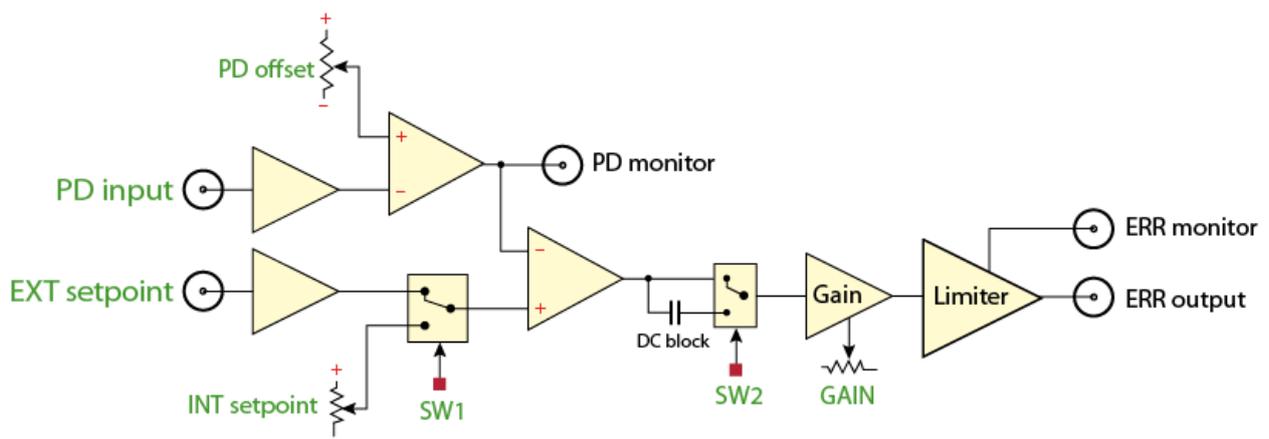


Figure 4: Simplified schematic of the B3121 signal processing chain. The inputs are $\pm 9V$ tolerant, and the offsets have a $\pm 1.25V$ range. The output limiter prevents the output voltage damaging the ARF's modulation input. The PD monitor should be used for noise analysis.

¹ Contact MOGLabs support for more information.

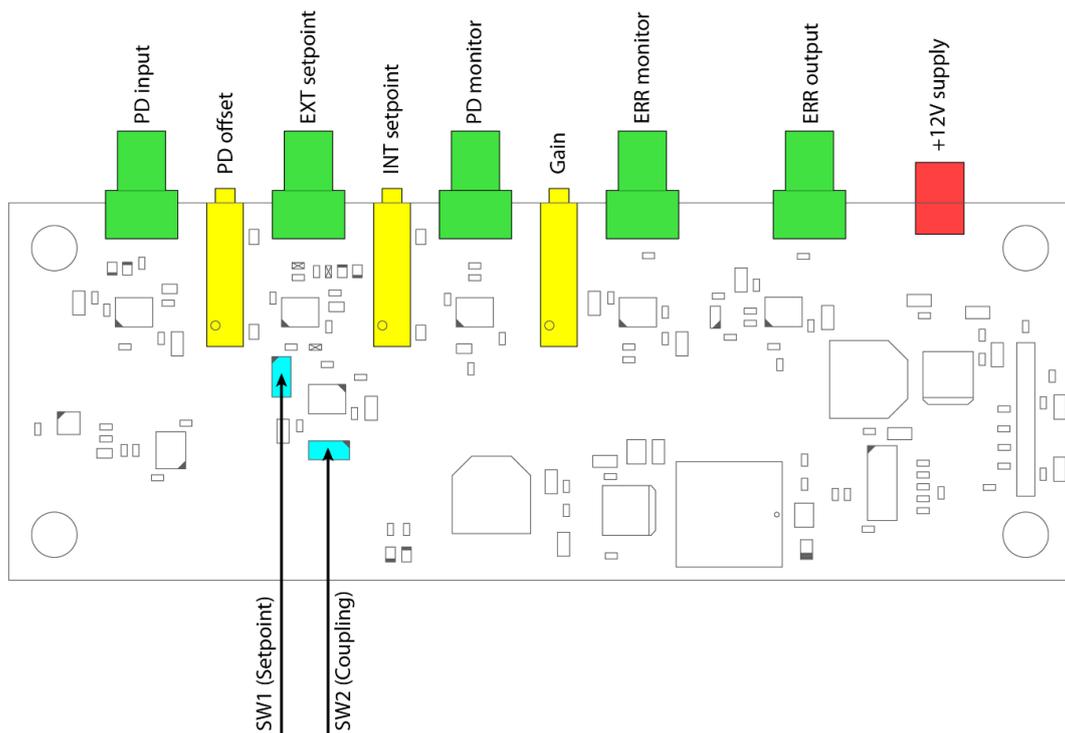


Figure 5: The B3121 signal processing board showing the SMA connectors (green), switches (aqua) and adjustment trimpots (yellow).

4. Configuring an intensity stabilisation apparatus

Fluctuations in the input beam power directly result in fluctuations in the output beam power, which are measured by the photodetector. When maintaining a fixed intensity, the error signal is zero and a large amount of analog gain to be applied, making the PID control loop sensitive to small fluctuations and allowing the servo to suppress them.

The following procedure is recommended for configuring an ARF for intensity-stabilisation using the B3121.

Note: for earlier versions of the signal processing board (B3120), please refer to the R1 version of this application note.

1. Connect the photodetector to the **PD input** SMA connector.
2. Set **SW1** to **INT** and adjust the **Offset** and/or **Setpoint** trimpots until the **ERR monitor** output is zero at the desired intensity.
3. Ensure **SW2** is set to DC.
4. Reduce the **Gain** to minimum.
5. Connect **ERR out** to the AM input of the corresponding channel on the ARF.
6. Connect to the ARF using *mogrf* and configure modulation as follows²:

² Alternatively the command language can be used to specify the configuration options; consult the ARF manual for more details.

In *mogrf*, open the Settings > Modulation window

- i. Set the Amplitude modulation gain to 1000
 - ii. Set the Proportional gain, KP to 1.0
 - iii. Set the Integral gain³, KI to 0.01
 - iv. Set the Derivative gain, KD to 0.0
 - v. Set the Sample rate⁴ to 7.81 MHz
 - vi. Set PID stabilisation to “Amplitude” and enable
7. Slowly increase the gain and watch the **ERR monitor** output. If the value diverges away from zero, then the action must be inverted using the **Invert PID action** checkbox.
 8. Increasing the analog error gain will improve the noise suppression, until the control loop becomes unstable at high gains. This is clear by observing one of the monitor outputs on an oscilloscope or spectrum analyser; large amplitude oscillations will be apparent when the gain is too high.
 9. Reduce the error gain slightly and fine-tune the PID constants using *mogrf*, or the command language. Since v2.8, *mogrf* includes some diagnostic tools for improving loop performance (see §6). Alternatively, a DC-coupled spectrum analyser (such as a computer sound card) can be used. Note that **ERR monitor** should not be used for noise spectrum analysis because it scales with gain.
 10. If necessary, the analog gain should be reduced such that the error signal is within $\pm 1V$ before engaging the PID.

5. Loop optimisation and analysis using *mogrf*

Since v2.8, the *mogrf* software suite provides a `PID debug` window for analysing and optimising the action of the PID controller. This can be accessed through the *Modulation Settings* window by clicking on `PID debug`. This window displays the raw input to the PID loop from the device ADC, the resulting PID output, and a Fourier transform (see Fig. 6). These graphs provide a convenient way to observe the action of the PID controller without requiring an independent spectrum analyser. In particular, it is easy to see whether the input or output has hit a limit, or whether the control loop has entered oscillation.

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 compensate loop behaviour. Typically this occurs when the overall gain is too low, or the response of the system is too slow. This can be seen in the debug window by the PID output reaching a value of ± 2048 . If the integrator becomes saturated, then the PID loop must be disengaged and re-engaged. The saturation can also be checked with the command **PID,VALUE,<ch>**, which may be beneficial to monitor in custom control systems. It is recommended that stable loops operate with a PID output within ± 1000 , to reduce the likelihood that the integrator saturates as a result of compensating for long-term fluctuations (e.g. thermal effects).

³ If it is not possible to achieve a stable lock initially, reduce KI to 0 until the other parameters (analog gain, AM gain and sign) are approximately determined, then introduce KI to improve noise suppression.

⁴ 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.

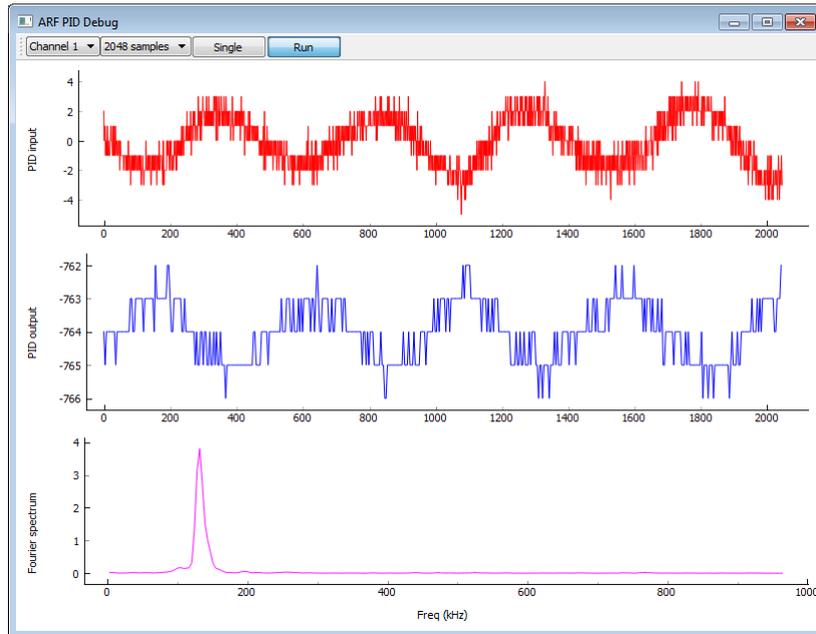


Figure 6: The PID analysis window of *mogrf* showing a stable closed loop, where the PID input is centred on zero and the output has stabilised without saturating. The Fourier spectrum indicates the closed loop bandwidth is around 120 kHz, and the presence of clear oscillation indicates the loop gain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

Conversely, oscillation typically occurs when the gain is too high and the PID controller over-compensates for fluctuations. Generally this is because the integrator action is too strong, which gives the best noise suppression at low frequencies but also makes the loop vulnerable to oscillation. The Fourier spectrum display of the debug window makes the onset of oscillation easy to observe. Trading off between the overall gain and the integrator KI allows oscillation to be suppressed without sacrificing performance.

When optimising the action of the PID controller, it is important to look at the step response of the system. In particular, when the loop is stable it is possible to greatly increase the gain without the loop becoming unlocked but the gain might become too high for the loop to lock from initial conditions.

The analog gain therefore must be chosen to ensure that for the expected range of initial conditions that the error signal is not initially saturated (i.e. within $\pm 1V$). The PID debug window provides a button labelled `Single` which performs a single-shot measurement of the PID action upon disengaging and re-engaging the control loop, which can be used to diagnose this behaviour (Fig. 7). When the loop is unstable, the behaviour can indicate how the PID parameters should be adjusted (Fig. 8).

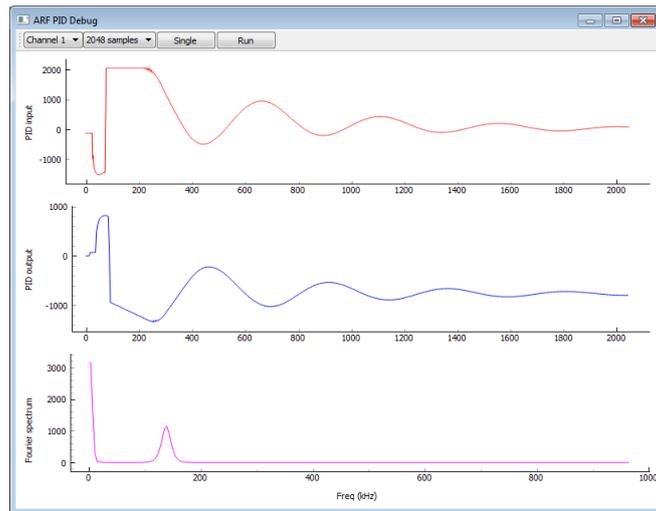


Figure 7: Demonstration of single-shot response. Initially when the PID is engaged the input becomes saturated, but the controller action successfully converges the input to zero without the PID output ever approaching saturation. Some oscillations occur initially but are exponentially dampened, indicating this loop is stable.

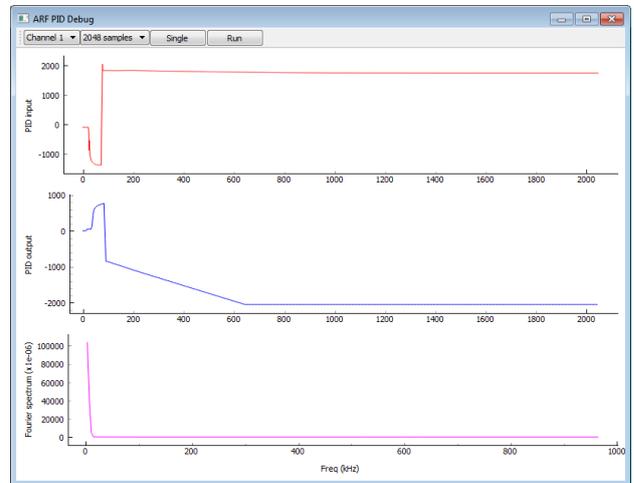
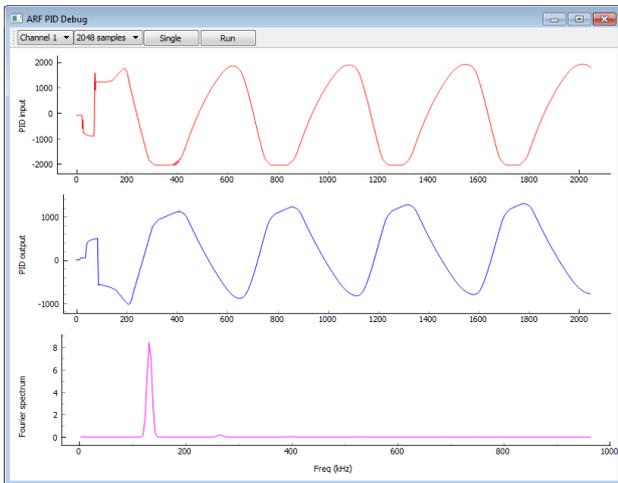


Figure 8: Demonstrations of the single-shot response for unstable control loops. When the gain is too high (left) the controller immediately begins oscillating without dampening. When the gain is too low (right) the controller action is insufficient to bring the error to zero, causing the integrator to saturate.

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 (Figure 9). Typically 20dB of suppression can be achieved out-of-loop.

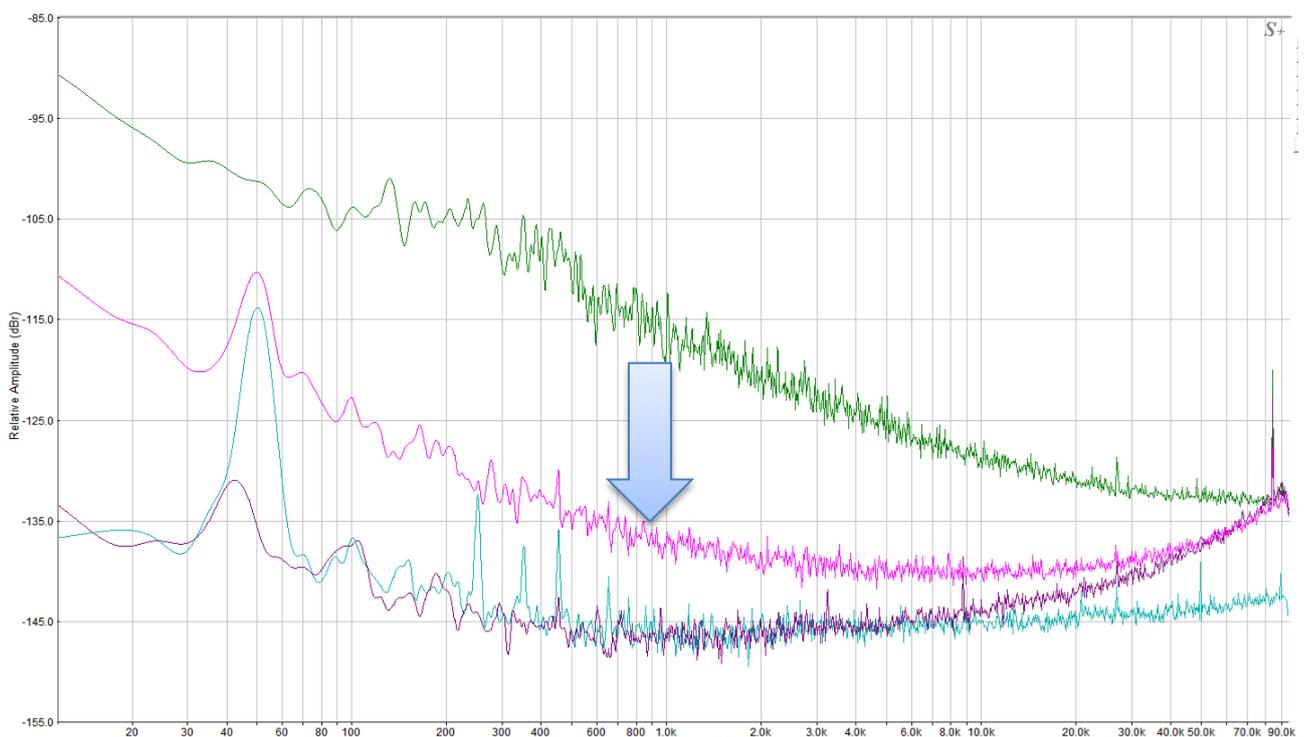


Figure 9: Typical measured noise spectra in an intensity-stabilisation experiment. Upon engaging the control loop, the incident beam noise (green) becomes suppressed (purple) into the dark noise (cyan), in apparent violation of the shot-noise limit. An independent (“out of loop”) detector more accurately describes the noise on the output beam and displays the shot-noise contribution (magenta). Noise suppression is typically observed up to 100kHz.

The relationship between rf power and diffracted beam optical power also means that any noise in the driving rf is demodulated 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 RFHIC amplifiers included in early 421-series ARFs (Rev4 and earlier) generate weak asymmetric 13.5kHz sidebands at <math><60\text{dBc}</math>. The AOM acts to demodulate the sidebands into intensity noise at harmonics of the sideband frequency (Fig. 10). This noise cannot be suppressed by the PID controller, reducing the effective bandwidth of the noise-eater in applications sensitive to noise at the shot-noise limit. **This noise is not present in Rev5 devices and newer**, as custom low-noise power amplifiers were developed to eliminate this issue.

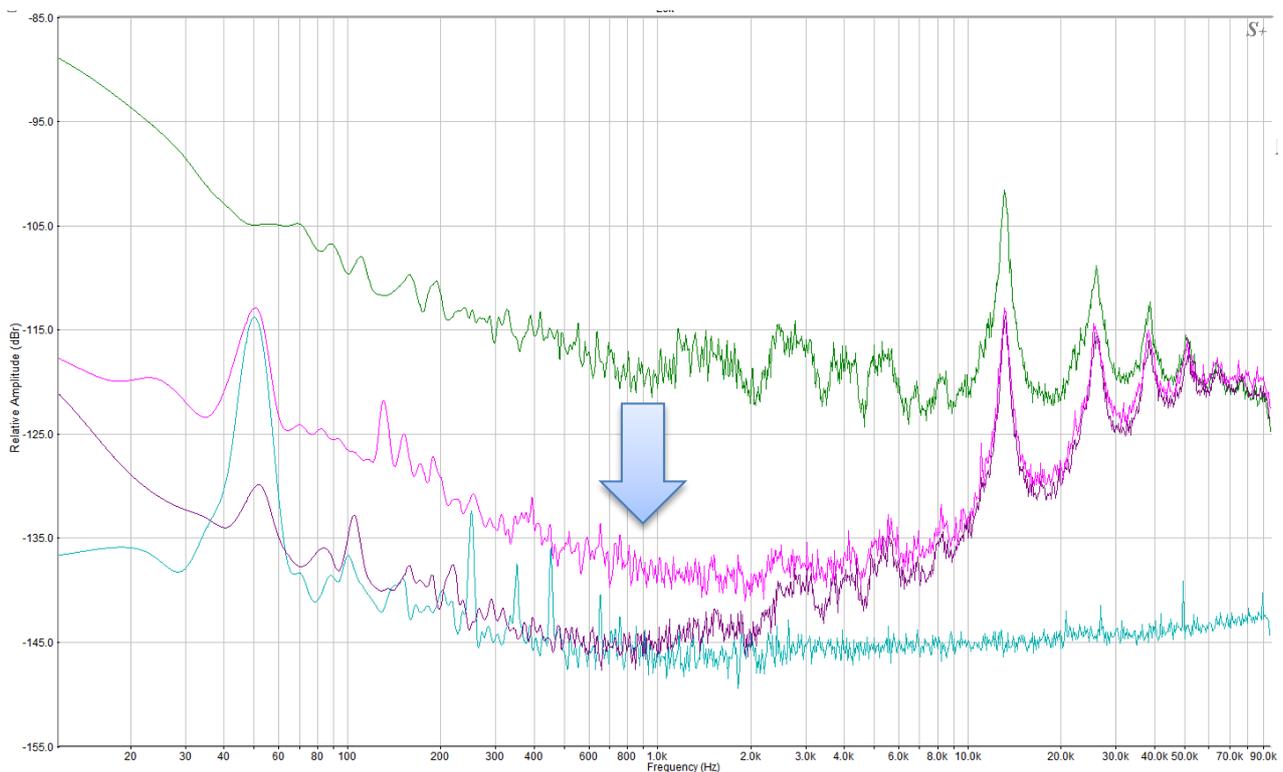


Figure 10: Typical measured noise spectra using a Rev2 ARF421 without PID (green), with PID locked (purple) and out-of-loop measurement (magenta). The noise features between 1-5kHz are residual power supply noise **only present in Rev2 units**, and the harmonics of 13.5kHz are caused by the RFHIC amplifiers **present in Rev4 and earlier devices**.

7. 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 to changes in both rf power and frequency mean that directly modulating the rf itself result in undesirable pulse shapes (Fig. 11). In principle the AOM response can be calibrated and the modulation signal corrected for this behaviour. However this calibration is sensitive to heating effects and alignment perturbations, often leading to poor pulse shapes. One approach is to use the PID controller to compensate for the non-linearities on the fly by feeding the signal processing board a time-varying setpoint (Fig. 12).

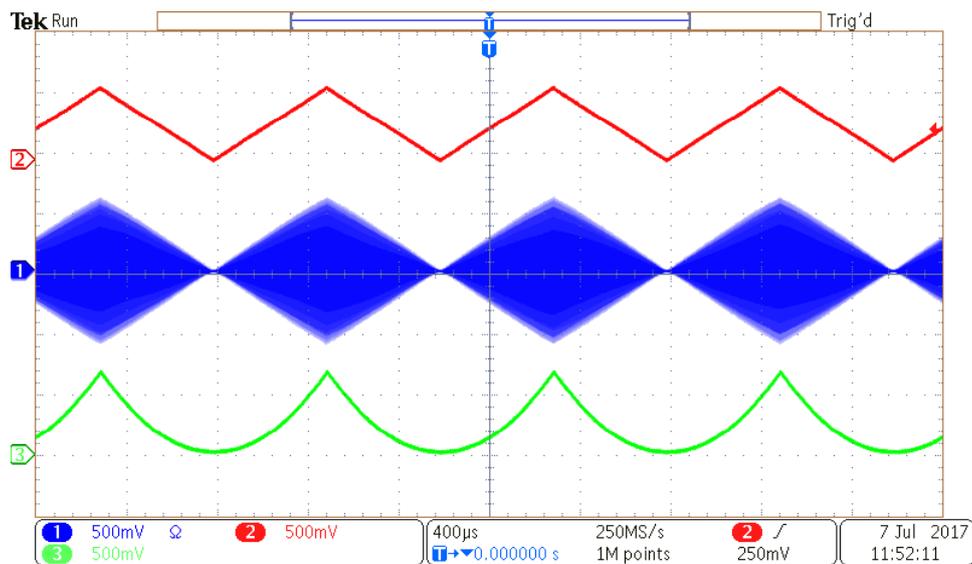


Figure 11: The output of a function generator (red) is used to amplitude-modulate the rf output (blue), which results in non-linear change in the beam intensity measured on a photodiode (green).

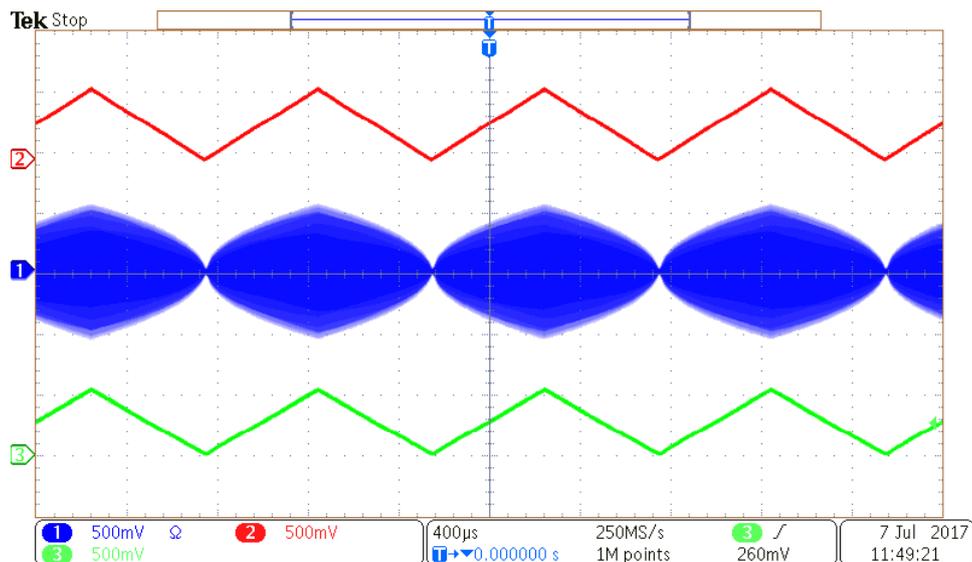


Figure 12: Using a closed-loop PID servo, the diffracted beam intensity (green) follows the desired form (red). The PID servo accounts for non-linearity of the AOM as seen in the rf envelope (blue).

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, as opposed to just drift. Furthermore, the combination of changing setpoint and signal processing delay means the integrator is especially susceptible to saturation. Therefore it is recommended to lower both the PID sample rate and integrator gain KI. The combination of large AM gain and proportional PID gain KP ensures rapid response to the changes in setpoint, and the residual integrator compensators 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 ARF'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.