AN001: PID intensity-stabilisation with the ARF

Acousto-optic modulators (AOMs) are devices that use an rf signal to diffract a laser beam. The frequency and power of the driving rf control the efficiency of the diffraction process, allowing the AOM to be used as a variable-attenuator by changing the amount of light in the diffracted beam. AOMs typically have a rapid response to changes in the driving rf, enabling amplitude modulation to be used to generate short pulses with a particular envelope. This modulation can also be used to shape the amplitude of the pulse for a specific purpose, or even counteract intensity fluctuations present in the incident laser beam, thereby reducing the classic intensity noise (“noise-eating”).

This Note describes using a MOGLabs Agile RF Synthesiser (ARF) as closed-loop servo that stabilises the diffracted beam intensity through an AOM. The instructions presented require **firmware v1.4.5 or newer**, which is available from www.moglabs.com.

1. Amplitude modulation and pulse-shaping

The MOGLabs ARF provides for amplitude modulation of the rf based on an external analog input, which allows for custom rf pulses. However, the AOM responds non-linearly to changes in both rf power (Figure 1) and drive frequency, causing the intensity of the resulting diffracted light to vary non-linearly with the modulation signal (Figure 2). Correcting this non-linearity requires either a “feed-forward” or “feedback” technique to generate a specific pulse shape in the intensity of the diffracted light.

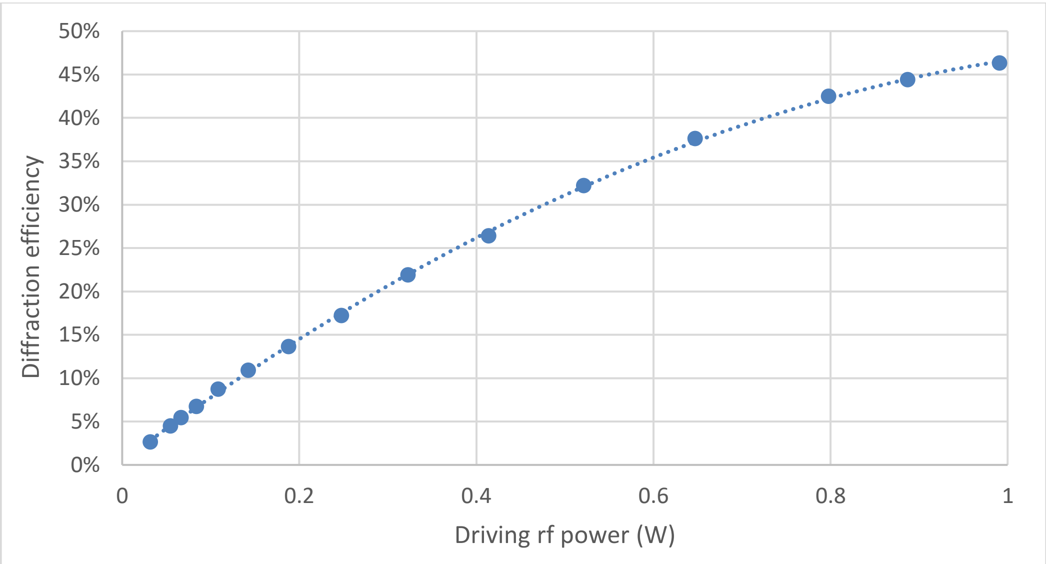

Figure 1: Typical saturation curve showing AOM diffraction efficiency as a function of rf power.

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

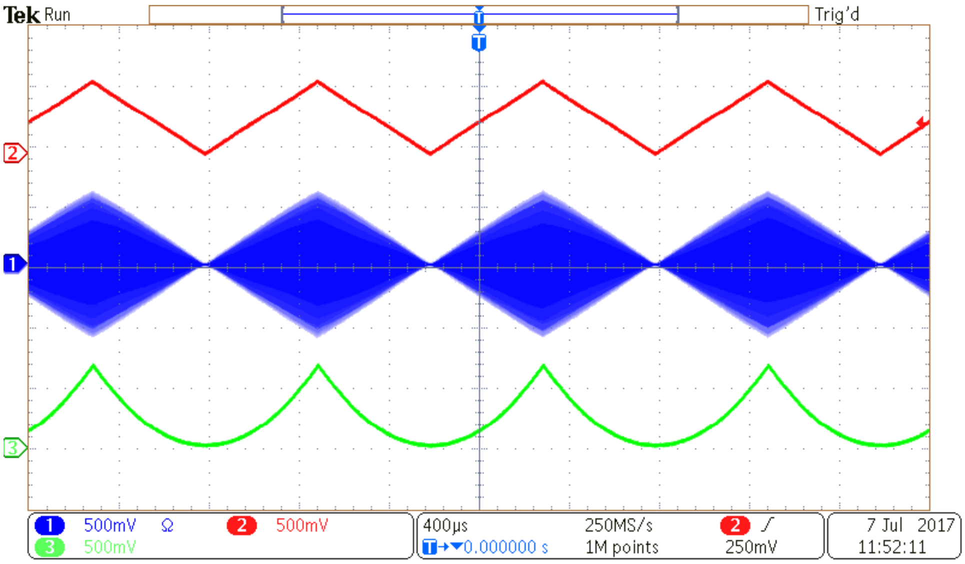

Figure 2: The output of a function generator (red) is used to amplitude-modulate the rf output (blue), which results in non-linear change in the beam intensity measured on a photodiode (green).

The ***feedforward*** approach is to carefully calibrate the response of the AOM as a function of power and frequency, then adjust the modulation signal fed into the ARF to compensate. However, this technique is susceptible to changes in calibration, such as heating effects, which can greatly affect AOM performance. Furthermore it cannot respond to changes in the input beam, and therefore cannot be used to suppress fluctuations in the laser intensity.

Another approach is to apply ***feedback***, whereby the diffracted light intensity is measured on a photodiode. The desired intensity is subtracted to compute a correction factor to amplitude of the rf driving the AOM. The ARF provides a PID control loop that converts this difference ("error signal") into changes in rf amplitude ensuring that the desired signal shape is achieved (Figure 3). This set-point subtraction must be done using analog signal processing, for which MOGLabs provides the B3120 signal conditional board.

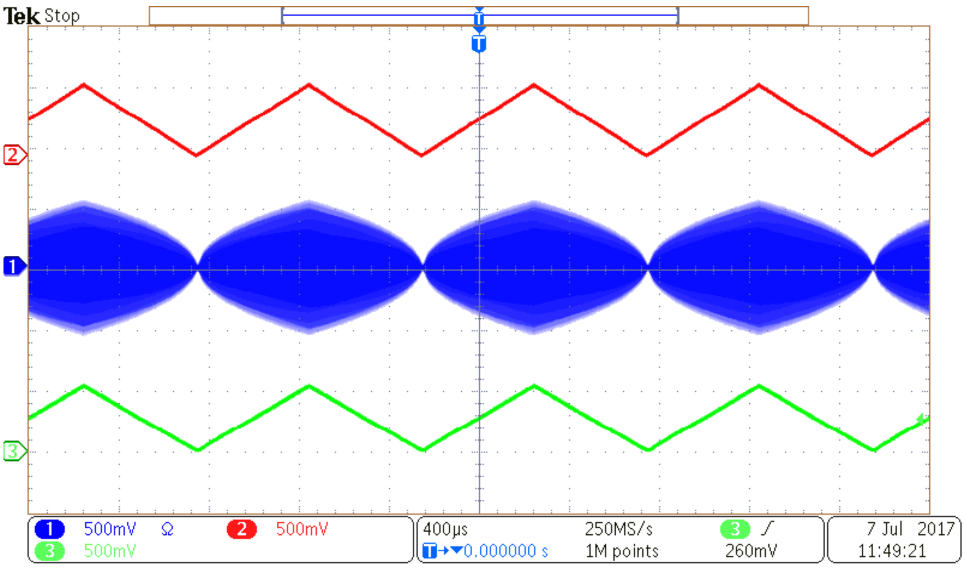

Figure 2: Using a closed-loop PID servo, the diffracted beam intensity (green) follows the desired form (red). The PID servo accounts for non-linearity of the AOM as seen in the rf envelope (blue).

2. B3120 signal processing

The B3120 board is used to generate an error signal for PID control of rf amplitude for intensity stabilisation. The signal processing chain is as follows:

1. The photodiode signal (±10V tolerant) is inverted depending on SW1, and amplified as controlled by RT1. In most configurations the master gain should be kept at minimum (unity gain).
2. The PD offset is subtracted, and a monitor is provided on the SMA output for noise analysis.
3. Depending on whether SW2 is set to LOC or REM, either the "Local" set-point defined by RV2, or the "Remote" set-point provided to the SMA connector (±10V tolerant) is subtracted.
4. If SW3 is set to AC, the result is AC-coupled, which is beneficial to some applications.
5. Analog gain is applied to the error signal, and a monitor is provided on the "Error Monitor" output for observation and debugging, and should not be connected to the ARF.
6. The result is clipped to ±1V and output on the "Error Output" SMA for connection to the **AMP modulation input** on the ARF.

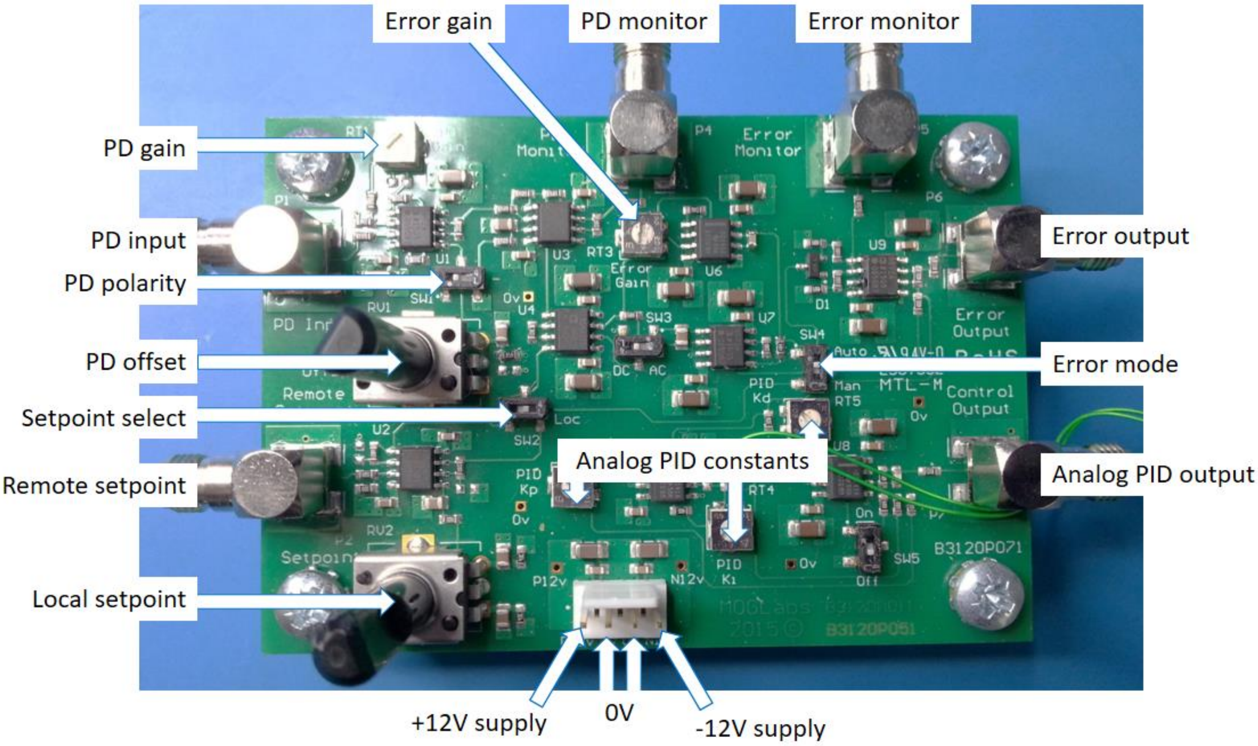

Figure 3: The B3120 signal processing board showing connectors, switches and adjustment potentiometers

The B3120 also contains an analog PID implementation, which can be used by connecting the "Control Output" to the ARF amplitude modulation input. However, due to difficulty in tuning the analog PID constants, it is recommended to use the digital PID implemented within the ARF; the "Control Output" SMA connector is unused in this configuration.

Note that the handles on the large potentiometers RV1 and RV2 can be cut to a more convenient size, as perturbing their value when the PID loop is locked can cause sudden jumps in the output. For intensity stabilisation applications, SW4 should be kept on "Auto".

3. Apparatus

A typical apparatus design is shown in Figure 4. The ARF drives an AOM, and the diffracted beam is discarded by an iris. An optical wedge (or beam-splitter) is used to pickoff 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 §6).

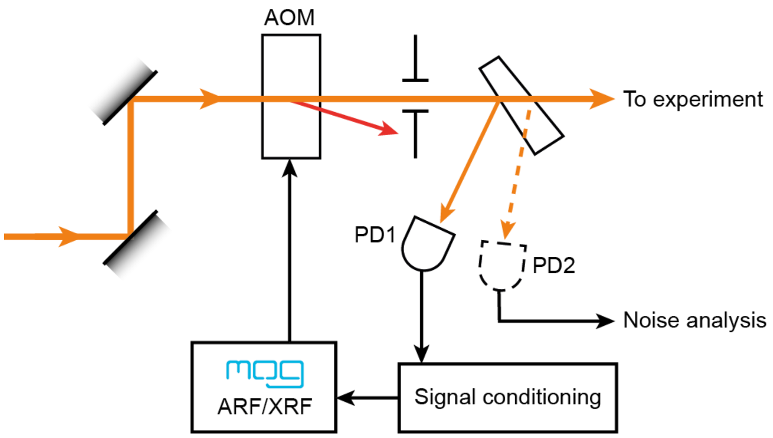

Figure 4: Apparatus diagram for PID intensity control with an ARF.

Equivalently, the apparatus of Figure 4 could be set up on the diffracted beam instead of the undiffracted beam if the frequency shift is desirable. It is also possible to use the discarded beam as the basis of the measurement (Figure 5), where an increase in the photodiode measurement is used to infer a decrease in the experiment beam power. This requires fewer optical components, but is more susceptible to additional scattering effects within the AOM and should not be used with high-power beams.

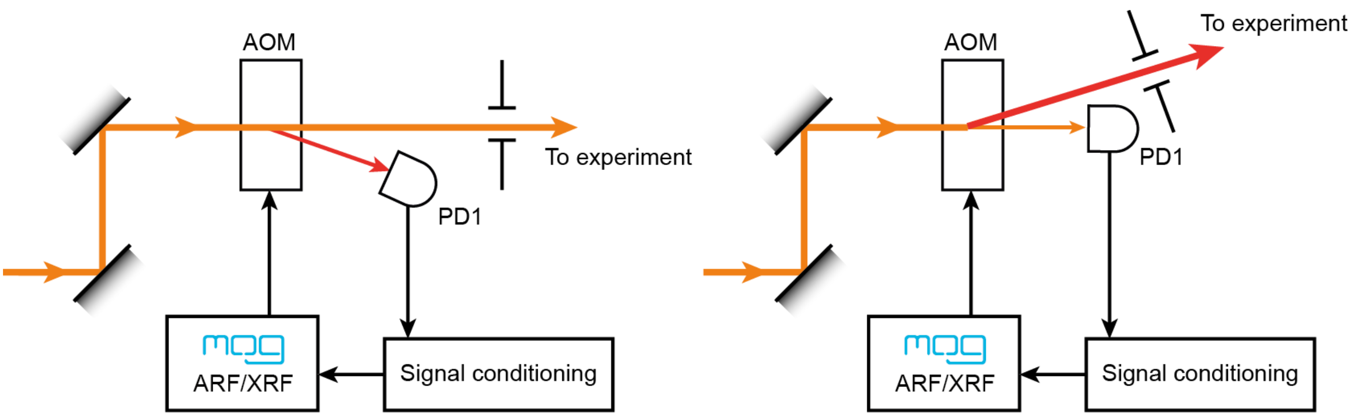

Figure 5: Alternate apparatus configurations for PID intensity control with an "inverted" error signal.

4. Set-point tracking

The following procedure can be used to configure the B3120 and ARF to follow an external setpoint, which is connected to the "Remote setpoint" SMA input.

1. Connect the photodiode to the "PD Input" SMA connector.
2. Block the light falling on the photodiode, and adjust the "PD offset" potentiometer until the output on the "PD Monitor" shows 0V.
3. Turn the "Master PD gain" trimpot (RT1) fully counter-clockwise, corresponding to unity gain.
4. Set SW2 to "Rem" and verify that the signal on the "Error Monitor" output is the negative of the setpoint signal.
5. Ensure that "Error mode" (SW4) is set to "Auto".
6. Adjust the "Error Gain" (RT3) until the "Error Monitor" spans 0V to ±1V.
7. Connect the "Error Output" to the AMP modulation input of the corresponding ARF channel.
8. Unblock the photodiode and activate the PID controller in AMPL mode, using one of the methods below:

   In ***mogrf*** (v2.4.1 or newer), open the Settings > Modulation window
   i. Set the Amplitude modulation gain to 8192
   ii. Set the Proportional, KP to 1.0
   iii. Set the Integral, KI to 0.1
   iv. Set the Derivative, KD to 0.0
   v. Set PID stabilisation to "Amplitude"

   Using Device Commander, issue the following commands (assuming CH1):
   ```
   GAIN,1,AMPL,8192
   PID,GAIN,1,P,1.0
   PID,GAIN,1,I,0.1
   PID,GAIN,1,D,0
   PID,ENABLE,1,AMPL
   ```

9. Observe the "Error monitor" output and increase the analog "Error Gain" (RT3). The amplitude of the error signal should decrease, as the controller follows the changes in set-point. If the error signal amplitude ***increases***, then the PID controller action must be inverted.
10. Continue to increase the "Error Gain" until the "PD monitor" shows the output is following the changes in set-point, but do not increase it past the point the output begins to oscillate.
11. Adjust the ARF amplitude modulation gain until the measured voltage reaches the desired bounds.
12. Tune the PID constants until the photodiode voltage most closely tracks the set-point signal. Different constants will be optimum depending on whether the set-point contains sudden changes or is smooth. Typically derivative gain is only useful for sudden changes in the set-point.
13. If the application is not tolerant to small amounts of light leaked when the input signal is 0V, set the input signal to 0V and adjust the "PD offset" potentiometer until no light is diffracted.

5. Intensity-noise suppression ("Noise-eating")

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 approximately 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 constant intensity-stabilisation:

1. Connect the photodiode to the "PD input" SMA connector.
2. Turn the "Master PD gain" (RT1) to minimum (fully counter-clockwise).
3. With desired level of light incident on the photodiode, adjust the "PD offset" (RV1) until the "PD monitor" is zero.
4. Set SW2 to "Loc" and adjust the "Local setpoint" potentiometer (RV2) until the "Error monitor" is also zero.
5. Reduce the "Error Gain" (RT3) to minimum (fully counter-clockwise).
6. Ensure that "Error mode" (SW4) is set to "Auto".
7. Connect the "Error output" SMA connector to the AMP modulation input of the corresponding channel on the ARF.
8. Activate the PID controller of the ARF, using one of the methods below:

   In ***mogrf*** (v2.4.1 or newer), open the Settings > Modulation window

   i. Set the Amplitude modulation gain to 1000
   ii. Set the Proportional, KP to 1.0
   iii. Set the Integral, KI to 0.1
   iv. Set the Derivative, KD to 0.0
   v. Set PID stabilisation to "Amplitude"

   Using Device Commander, issue the following commands (assuming CH1):

   ```
   GAIN,1,AMPL,1000
   PID,GAIN,1,P,1.0
   PID,GAIN,1,I,0.1
   PID,GAIN,1,D,0
   PID,ENABLE,1,AMPL
   ```

9. Slowly increase the error gain (RT3) and watch the "Error monitor". If the DC average moves away from zero, then the PID loop is amplifying the noise instead of suppressing it, and the action must be inverted. Select "Invert PID action" in ***mogrf***, or use the command `PID,INVERT,1,ON`.
10. Increasing the analog error gain will improve the noise suppression, until the control loop becomes unstable at high gains. This is clear by observing the "Error monitor" output on an oscilloscope, which exhibits large amplitude oscillations when the gain is too high.
11. Reduce the error gain slightly below this maximum, and fine-tune the PID constants using ***mogrf*** or the command language. This is best done by observing the "PD monitor" on a spectrum analyser. Soundcard-based spectrum analysers are a cost-effective solution for noise analysis as the controller action is primarily below 100kHz. Note that the "Error monitor" should not be used for noise spectrum analysis because it scales with gain.

6. Further notes on intensity stabilisation

The underlying assumption of the control loop is that the measured photodiode signal is proportional to the power in the primary beam. It is therefore very important to ensure that the photodiode does not saturate, or detect scatter or diffuse light from elsewhere in the experiment. For example, an unshielded photodetector will measure fluctuations from fluorescent room lighting at harmonics of the mains power frequency (i.e. 100Hz or 120Hz) and therefore 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.

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 suppresses the shot-noise ***in the measurement arm***, which cannot be simultaneously used for another purpose. The beam used for the experiment is split off earlier and contains ***uncorrelated shot-noise***, and therefore remains shot-noise limited. An independent pick-off and photodetector in the experiment arm will provide a more accurate measurement of the true noise-level of the beam (Figure 6).

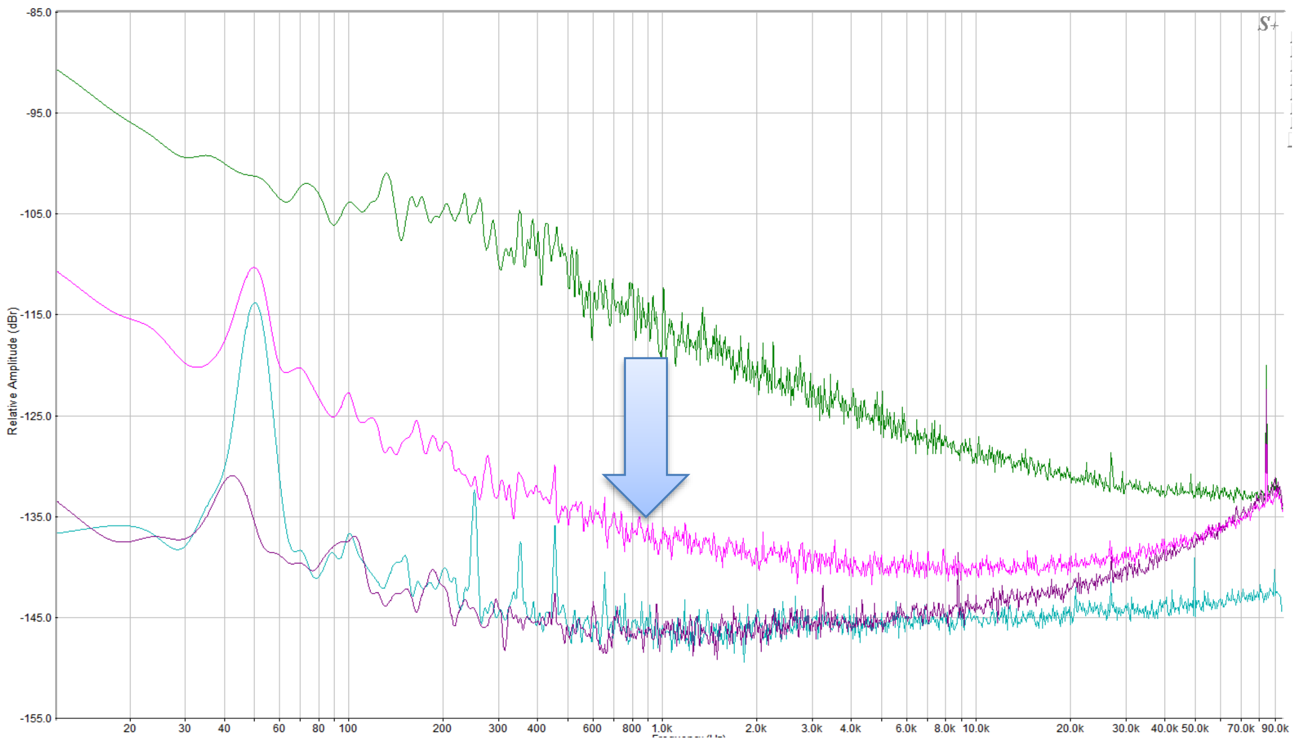

Figure 6: Typical measured noise spectra in an intensity-stabilisation experiment using a Rev4 ARF021. Upon engaging the control loop, the incident beam noise (green) becomes suppressed (purple) into the dark noise (cyan), in apparent violation of the shot-noise limit. An independent measure of the noise on an "out of loop" detector (magenta) more accurately describes the noise on the output beam. Noise suppression is typically observed up to 100kHz.

The primary limitation of any closed-loop servo is the "lag" or "propagation delay" between changes in the control signal causing changes in the error signal. In the ARF PID implementation, this is caused by the signal processing chain FPGA → DDS → rf electronics → AOM → photodetector → signal conditioning → ADC. The net result is that the achievable PID bandwidth is approximately 100kHz. The analog bandwidth of the photodetector, the signal-processing board and ARF modulation inputs limit how rapidly the PID loop is able to respond to step-changes (such as changing the intensity set-point) and is typically much higher than the practically achievable closed-loop control bandwidth.

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 generate weak asymmetric 13.5kHz sidebands at <-60dBc. The AOM acts to demodulate the sidebands into intensity noise at harmonics of this frequency (Figure 7). 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. However, the latest 421-series ARFs (Rev 5+) have low-noise MOGLabs power amplifiers that do not exhibit this limitation.

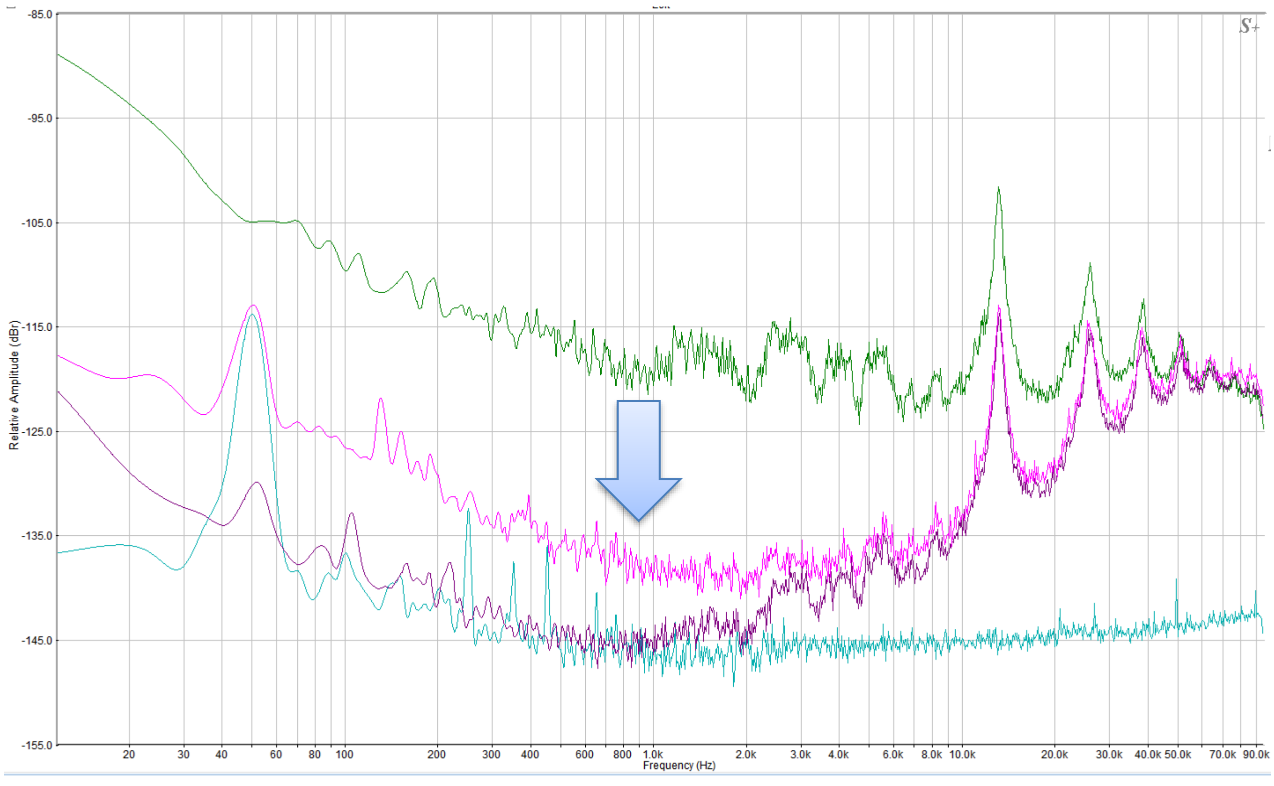

Figure 7: Typical measured noise spectra using a Rev2 ARF421 without PID (green), with PID locked (purple) and out-of-loop measurement (magenta). The additional noise features between 1-5kHz are residual power supply noise in Rev2 units, and the harmonics of 13.5kHz are caused by the RFHIC amplifiers integrated into early 421-series devices.