Balanced Photodetector

Robin Oswald

December 18, 2012

Abstract

A balanced photodetector is built with applications such as saturated absorption spectroscopy in mind. Because noise common to both the signal and the reference beam cancels only completely if the photocurrents of both photodiodes are equal in magnitude, an electronic feedback loop is incorporated that balances the photocurrents from both photodiodes to ensure total noise cancellation, even if the optical powers in the beams differ (excepting shot noise). The circuit, possible modifications, its performance and common issues are discussed. Finally, advice regarding the operation of the photodetector is provided.

Contents

1 Introduction

Light sources such as lasers used for optical measurements or control are inherently noisy. This can be ameliorated by careful elimination of noise sources, frequency and intensity control, as well as by mechanical and thermal stabilization. However, it is often not necessary to eliminate all the noise in order to get precise measurements. By using an appropriate detection scheme one can extract a small signal buried in noise. When one knows the signal being looked for, lock-in, heterodyne or autocorrelation can be used. Otherwise, balanced photodection can be employed. The basic idea behind it is to split the laser beam into two equal parts, called the signal and the reference beam. The signal beam is then sent through the experiment, while the reference beam is left unaltered (or in the case of saturated absorption spectroscopy sent through the gas cell without a counterpropagating pump beam). In the end, both beams are incident on photodiodes, which generate currents proportional to the optical intensities. These respective photocurrents are then subtracted. Ideally, this results in complete cancellation of noise common to both beams, leaving behind only the signal as well as shot noise, which is due to the discreteness of photodetection events [1].

However, this only works if the optical powers of both beams are matched at all times, which may be hard to guarantee in an experimental setting. Otherwise, noise cancellation will be incomplete. This issue can elegantly be circumvented by noticing that while it is not so easy to adjust optical powers dynamically in order to ensure equal magnitudes, the photocurrents generated by the photodiodes can easily be scaled down electronically and matched. A simple solution to this problem is therefore to make the reference beam stronger than the signal beam by a factor of $1.1 - 2$. The resulting photocurrent of the reference beam can then be scaled down electronically to exactly match the signal beam photocurrent at all times and hence achieve full noise cancellation, even for differing and time variant optical intensities. So instead of ensuring identical optical powers at all times, the photocurrents are matched electronically (on a slower time-scale than that of the signal, so that it is not lost), which is much more convenient.

The goal of this project is to build such an autobalancing photodiode circuit for use in our laboratory.

2 Photodiode Frontend Circuit

The purpose of a photodiode fronted circuit is to amplify and process the weak photocurrent generated by a photodiode to a strength useful for further processing. Hobbs describes the common challenges and how to overcome them in [2][3]. There are three main criteria that one tries to optimize: Signal to noise ratio (SNR), absolute signal strength, and speed. As is often the case, one can easily achieve two of these, while the third is counteracted by progress in the others. Using a load resistor as front end for example, the SNR and signal strength increase with a higher load resistance, while the cut-off frequency decreases with increasing resistance due to the large CR time constant.

2.1 Final Design

A cleaned up version of the final schematic is shown in Figure 1. It is based on a circuit proposed by Hobbs in [4], adapted to work in our lab environment. The full schematic is shown in Figure 3 in the Appendix on page 14.

The two photodiodes PD1 and PD2 are reverse biased with plus and minus 15 V respectively. The current I_{Signal} of PD1 flows through transistor Q1, which acts as a common base amplifier. It serves as a current buffer and the output current equals the input current. This arrangement allows for higher speeds compared to directly feeding the current to the transimpedance amplifier, because the photodiode sees only a small signal resistance of $R = \frac{kT}{eI_{\text{Signal}}}$ [2]. The photocurrent of PD2 also flows through a common base amplifier consisting of Q2 and Q3, with the difference that the current is split into two subcurrents: I_{C1}, which is flows to ground, and I_{C2}. The splitting ratio between these currents is determined by the biasing voltage applied to the base of Q3 [4]:

$$\frac{I_{\text{C1}}}{I_{\text{C2}}} = \exp(-\frac{eV_{\text{Q3B}}}{kT}) \qquad (1)$$

The current $I_{\text{Signal}} - I_{\text{C2}}$ is amplified and converted to a voltage signal by the transimpedance amplifier U1.1. U1.2 then serves as a unity-gain buffer in order to drive the usual 50Ω loads that coaxial cables present.

The signal is also fed to the integrating op amp U1.3, which has a RC-time constant in the order of milliseconds. The accumulated voltage upon integration is then used as the bias voltage for the base of Q3, closing the negative feedback loop, which works as follows:

If the integrated voltage $(I_{\text{signal}} - I_{\text{C2}})R_2$ is getting negative, the current I_{C2}

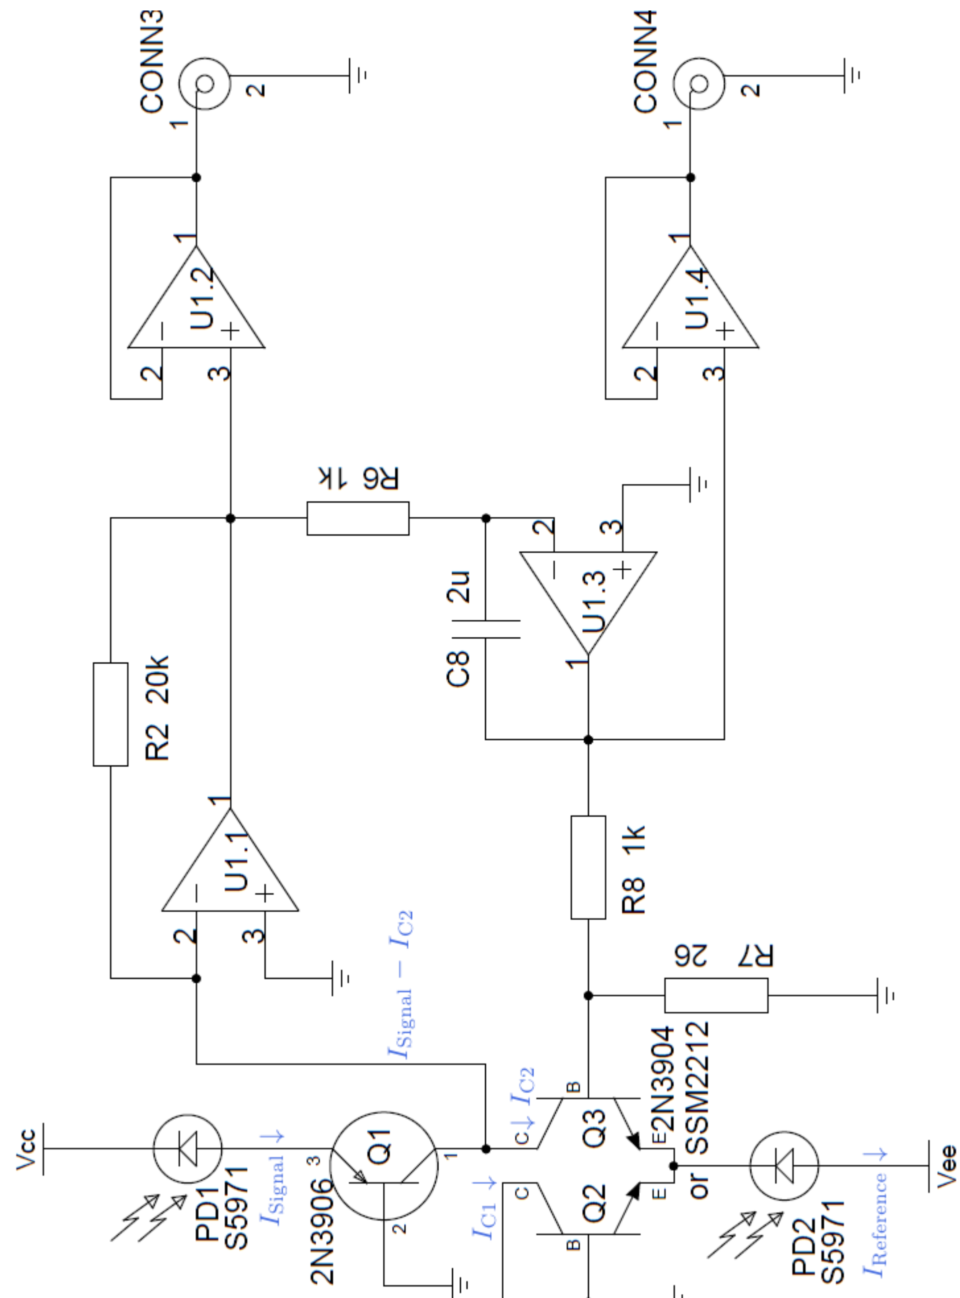

Figure 1: The simplified circuit layout. The reference designations of the components matches those of the complete schematic shown in the Appendix.

is reduced, resulting in a more positive $(I_{\text{signal}} - I_{\text{C2}})R_2$ and vice versa. So this process keeps $(I_{\text{signal}} - I_{\text{C2}})R_{\text{F}}$ close to zero at all times, ensuring equal time averaged photocurrents and hence resulting in full noise cancellation. Please note that the noise cancellation is not limited by the low feedback frequency, but only by the current gain-bandwidth product f_{T} of the transistors. Therefore, the noise cancellation works up to high frequencies and not just in a narrow frequency band.
CONN4 can provide an optional logarithmic output (however, be warned that this has hardly been tested):

$$V_{\text{CONN2}} = -\ln\left(\frac{I_{\text{ref}} - I_{\text{signal}}}{I_{\text{signal}}}\right) \tag{2}$$

Compared to the simplified circuit, the actual circuit contains some additional elements. These are:

- Additional resistors in order to generate static biasing voltages for the base of Q3 to set fixed current splitting ratios. The voltage divider consisting of R3 and R4 and generates +300 mV, forcing all the current to go through Q3 and allowing the photodiode to work without any autobalancing. R3 together with potentiometer R10 allow for manual adjustment of the splitting ratio. Switch S2 allows to choose between autobalancing or one of the static voltages, while switch S1 allows to choose between manually adjusting the splitting ratio with potentiometer R10 or a fixed, ordinary balanced operation where both photocurrents are simply subtracted.
- Supply decoupling capacitors C1, C2, C4, C5, C6, C7, C11 and C12.
- Voltage regulators U2, U3 and associated buffering capacitors C9 and C10.
- Power LED LED1 and current limiting resistor R11.
- Current limiting resistors R1 and R5 for the photodiodes.

The final circuit uses mostly standard components which should be readily available. Only if the logarithmic output is to be used a matched pair npn transistor such as the SSM2212 has to be employed instead of two 2N3904 npn transistors.

2.2 Possible Modifications

There are many possible modifications that can be done to the circuit and the pcb is already prepared for some of them:

- If the autobalancing is too slow (which is the case if there is a dc component left in the output signal), R6 and C8 can be lowered. This increases the $\frac{1}{RC}$ prefactor of the integrator and leads to faster adjustment of the splitting ratio due to a faster voltage build-up. However, if the autobalancing update speed is made too fast, the signal will deteriorate, because the autobalancer will literally balance it out.

- If the final voltage signal is in the wrong range, change the feedback resistor R2. If the update speed of the autobalancing is to be kept the same, R6 and R8 need to be adjusted as well.

- If the autobalance feature is not needed, the matched pair can be replaced by a single 2N3904 and all the components associated with the feedback loop should be left out. In this configuration, the circuit works as an ordinary balanced photodetector. It is strongly advised to use this configuration if the autobalance feature is not needed, because it greatly simplifies the design and hence reduces the amount of possible error sources.

- If speeds higher than 1 MHz are desired, then the op amp serving as the transimpedance amplifier should be replaced with a faster model, as it is the limiting component (the S5971 photodiodes work up to 100 MHz with a reverse bias of 15 Volts). It is important that the application notes in the data sheet of the op amp are followed, as high speed op amps are very prone to oscillations.

- If there is no need for a logarithmic output (which should usually be the case), the matched pair SSM2212 (SOIC8 package) can be replaced by two inexpensive 2N3904 npn transistors. The circuit board already contains additional holes to accommodate the pins of a TO-92 package.

- If there is no need for a feature such as manual balance, one can leave away associated components and greatly simplify the assembly, as the box is quite crammed in the full configuration.

- The circuit can also be stripped down to accommodate a single photodiode. Leave out the matched pair, PD2 and its resistor and capacitors and everything from the feedback path. However, as the circuit was not designed with single photodiode operation in mind, performance will most certainly be suboptimal.

2.3 Performance and Measurements

The photodetector has been tested in a simple setup where a laser beam was incident on PD1, while PD2 was blocked. The detector was operated with autobalancing turned off, so that only the voltage signal due PD1 was measured. The incident laser beam was strobed on and off with the help of a AOM at various frequencies, and the voltage signal of the photodetector was measured on an oscilloscope for a wide range of optical powers, which were also measured with an optical power meter. Three modulation frequencies were used: 100 Hz, 1 kHz and 10 kHz. Furthermore, the noise signal voltage range was also recorded. The results are plotted in Figure 2, while the raw data is stored in `Measurements.ods`.

As the log-log plot shows, the detector is linear for a wide range of optical powers. Only the data point around 25 μW for 100 Hz is troublesome and is suspected to be a erroneuous reading. The magnitude of the noise voltage remained constant at around 700 μV for virtually all measurements. Accordingly, the SNR increases for increasing optical powers.

During these measurements, a serious problem has been discovered, namely that the signal voltage slowly decreases over time, even when it theoretically should stay at a constant level. This indicates that a capacitor is slowly charged. Due to a lack of time, no detailed investigation has been carried out, but it is suspected that the integrating capacitor C8 is the culprit. If this is indeed the case, the problem, which should only occur if autobalancing is turned off, can best be avoided by assembling a separate pcb for use in normal balanced mode without the components required for autobalancing. A more elaborate solution would be to somehow disconnect C8 if autobalancing is turned off. However, that would probably require changes to the circuit board and is hence not so easy to implement.

Unwanted oscillations have also been noticed. Because this problem has been anticipated, it can easily be resolved by soldering in a capacitor in the low picofarad range into the currently empty slot C3 that is present on the pcb.

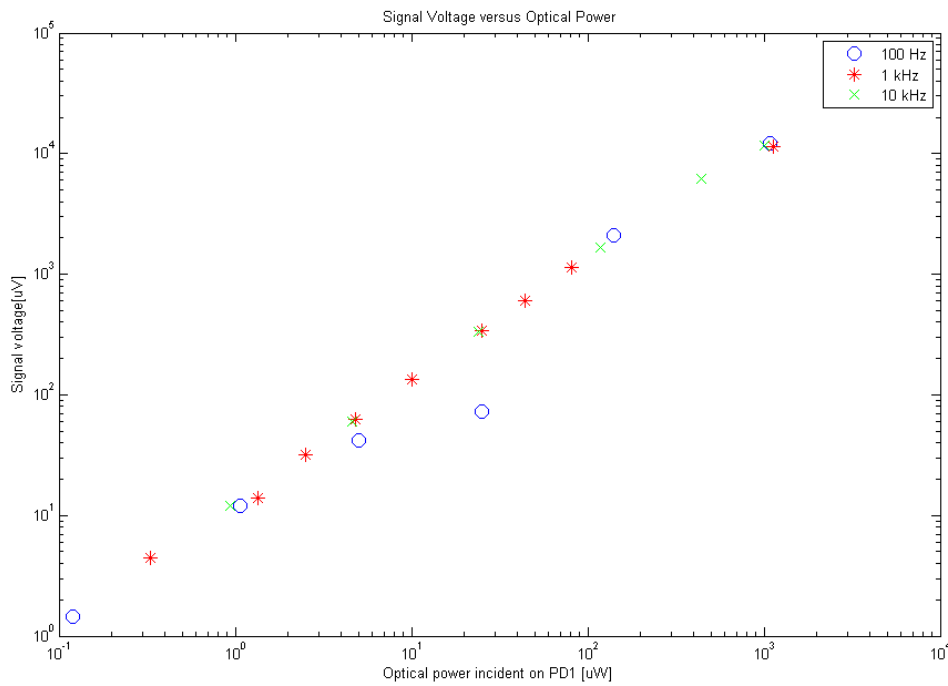

Figure 2: The signal voltage measured at the detector output for different optical powers and modulation frequencies.

3 Operating Manual

In this section, advice regarding the regular operation of the autobalancing photodetector is provided. Please note that the described device is very similar to the Newport Nirvana 2007, which is also a descendant of the design proposed by Hobbs. Therefore, its manual [5] can also be consulted for further guidance.

3.1 Specifications

The photodiodes are model S5971 from Hamamatsu Photonics [6]. They operate between 320 to 1060 nm with a peak sensitivity of 0.64 A/W at 900 nm. Using a bias voltage of 15 V, their cut-off frequency is around 100 MHz, far higher than that of the TL074 op amp, which limits the maximum detection frequency to around 1 MHz.
The photodetector needs to be supplied with $-15, 0, +15$ Volts. It can simply be powered using the LEMO plug and the standard TIQI power supply (PETH-40B3). The output can be approximated by

$$V_{\mathrm{OUT}} \approx 0.5 \frac{A}{W} \cdot 20\mathrm{k}\Omega \cdot (P_{\mathrm{Sig}} - P_{\mathrm{Ref}}) \tag{3}$$

3.2 Detector Operation

The balanced photodetector has three basic modes of operation: Autobalancing (automatic matching of the photocurrents for optimal noise cancellation), Balanced 1:1 (the photocurrents are subtracted and both beams have the same gain) and Manual Balance (the photocurrents are subtracted, but the gain of PD2 can be adjusted manually). The two switches on the side of the box allow choosing the appropriate mode. The lower one toggles autobalancing on or off, while the upper one switches between 1:1 and manual balance should autobalancing be turned off (it has no effect if autobalance is on).

Balanced 1:1 In this mode the circuit operates as a normal balanced photodetector, subtracting the signals of the two beams. The voltage output at the SMA connector is $V_{\mathrm{Out}} = S(P_{\mathrm{Sig}} - P_{\mathrm{Ref}})R_{\mathrm{F}}$, where P_{Sig} stands for the optical power in the Signal beam, S for the photo sensitivity of the photodiodes and R_{F} for the feedback resistor (20kΩ).

Manual Balance In this mode, the output is $V_{\text{Out}} = S(P_{\text{Sig}} - nP_{\text{Ref}})R_{\text{F}}$, where one can set n manually between around 0.5 and 1 by adjusting the potentiometer. Please note that there may be many turns necessary in order to notice a change. Due to time constraints and a lack of direct applications this feature has not been tested extensively.

Autobalancing In this mode, the magnitude of the photocurrent generated by the reference beam is dynamically matched to that of the signal beam with an update rate in the low Hertz range, which is dependent on factors such as the optical powers, power differences between the beams and the values of the resistors and capacitors chosen. If detailed knowledge of the update frequency is desired, one can monitor the biasing voltage that sets the splitting ratio on an oscilloscope with the help of the pin CONN5.
As previously explained, this autobalancing results in better common mode noise cancellation than simple balanced operation. Because the photocurrent of the reference beam is scaled down to match that of the signal beam, **it is essential that the reference beam has more power than the signal beam. Otherwise the detector can not function properly!** A ratio $P_{\text{Ref}}/P_{\text{Sig}}$ between about 1.1 and 2 is recommended.
If a modulated signal is to be measured, make sure that the modulation frequency is well above the cut-off frequency of the autobalancing mechanism. Otherwise, the signal is heavily degraded because the reference photocurrent will erroneously be adjusted due to the autobalancer trying to counteract the modulation.
The output in this mode is ideally only the signal. Note also that due to the matching of the dc powers of the photocurrents, the output signal has no dc component. Therefore, a remaining dc component is a good indicator that the autobalancer does not work as intended.

3.3 Quick Start for Autobalancing Operation

1. Mount the detector in the optical setup. Make sure that the reference beam has a higher optical power than the signal beam.

2. Set Switch 1 to Balanced and Switch 2 to 1:1.

3. Connect a SMA cable to the signal output and monitor it on an oscilloscope.

4. Apply $\pm 15, 0$ Volts to the detector using the LEMO connector.

5. Block one of the photodiodes and maximize the signal output by adjusting the beam coupling. Repeat this procedure for the other beam.

6. Change Switch 1 to AutoBalance.

3.4 Common Error Sources and Debugging

Should the photodetector not work as expected, it is advised to check the basic functionalities one after another:

- Check if the power regulators works. There are two probe places marked `15V` and `-15V` on the pcb that should make it easy to probe these voltages.

- In the balanced 1:1 mode, is there a measurable voltage change at the negative input of U1.1 if the beam is chopped manually? If yes, check for both photodiodes that 0.1 mW of optical power results in approximately 1 Volt at the SMA output. Soldering in the photodiodes the wrong way appears to be a very common mistake.

- Verify that in manual balance mode, the biasing voltage generated by the voltage divider consisting of R3 and R10 controls the prefactor of the reference beam photocurrent.

- When debugging the autobalancer, it is very helpful to look at the voltage at CONN5 which is the biasing voltage that sets the current splitting ratio.

- It is also very useful to look at the signal output on a spectrum analyser. A peak at an unexpected frequency indicates a parasitic oscillation that has to be analysed further.

4 Summary

An autobalancing photodetector has been built after some first experiences with a prototype have been gathered. In the simplest test settings, the detector worked well and as expected, with a signal to noise ratio of up to 70 dBm for frequencies below 1 MHz.
The real performance indicator for the device will of course be real world applications. Unfortunately, a quick test in a saturated absorption spectroscopy setup ended negatively. The detector was able to measure both signals independently, but upon subtraction, only noise was left over. Due to a lack of time, the cause of this failure could not be determined. Additionally, during the performance measurements, some further flaws have been discovered. Suggestions on how to resolve, or at least avoid the issues have been given but not yet tested or implemented.
So in conclusion, a solid platform for an own autobalancing photodetector has been built, but some further efforts are still needed before it can reliably be used for general purpose real world applications. For specific applications, existing problems can be circumvented by assembling the pcb tailored to the specific usage case.

5 Appendix

All the files used in this project are stored in the git repository called `balancedpd` and are stored on the J-drive.

5.1 Assembly Guidance

The assembly of the complete balanced photodetector poses some challenges due to space constraints, because the author recklessly ignored warnings. Therefore, some guidance is provided here in order to facilitate the process.

1. Solder in all the components except of the SMA connector and the LED. It is helpful to solder in the more complicated parts such as the TL074 first when there is still a lot of free space on the pcb.

2. Solder cables to the switches and then solder the other ends to their respective positions on the pcb. Choosing appropriate lengths for the cables is very helpful later on.

3. Solder the 3 pin connector to the LEMO connector.

4. Solder in the LED, paying attention to the right length of the pins. Bend it such that it will fit through the hole reserved for it in the box.

5. Put the LEMO and SMA connectors in their respective positions in the box.

6. Attach the switches to the box. This will confine the pcb to the box.

7. Put in the pcb from below the box and solder it to the legs of the SMA connector. Try to get the pcb in the right orientation from the beginning, as unsoldering and resoldering it is very tedious.

8. Add the insulating material to the bottom cover piece, leaving openings for the photodiodes.

9. Screw everything together and check for short circuits before powering it up.

5.2 Complete Schematic and PCB

Figure 3 shows the final circuit, `balancedv1.2_final.sch` and Figure 4 the final pcb layout, `balancedv1.2_final.pcb`

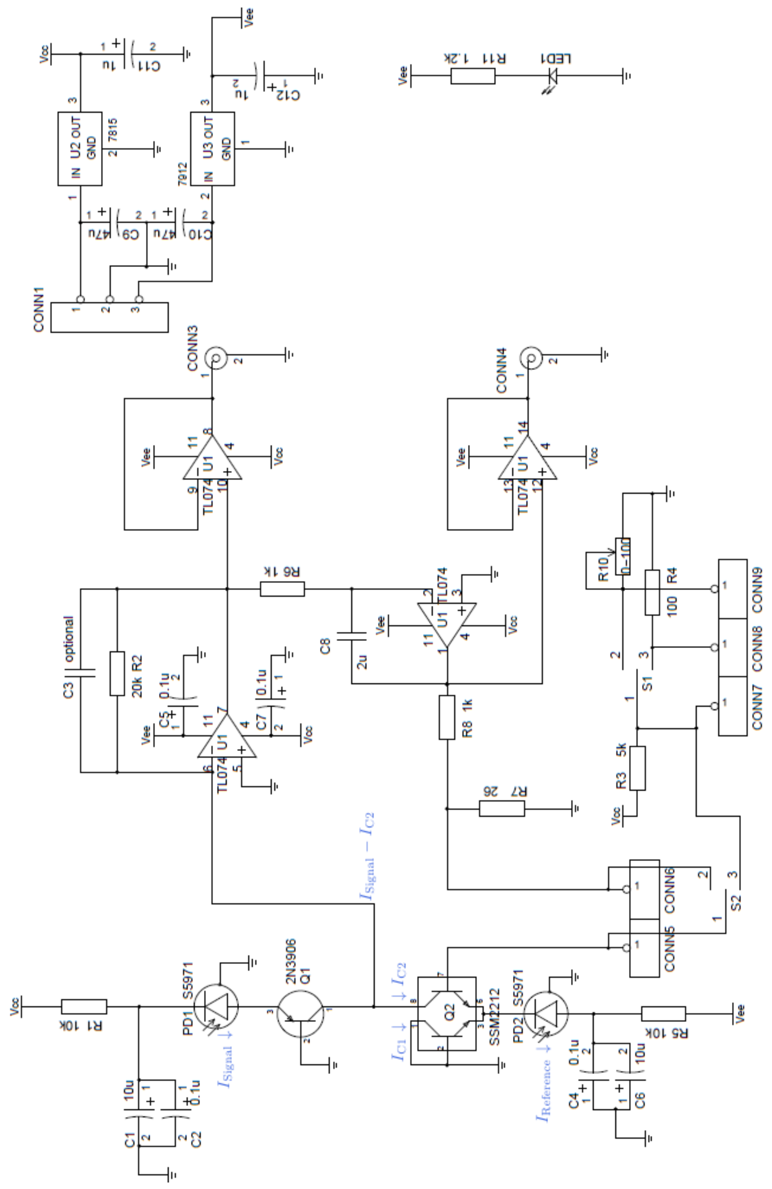

Figure 3: The complete schematic

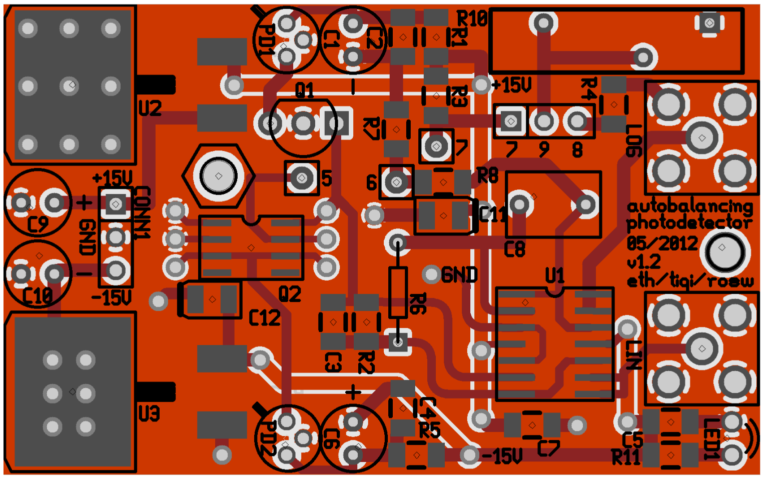

Figure 4: The pcb layout sent off to manufacturing. There are around 15 more bare pcbs lying around, waiting to be assembled.

5.3 Box

Figure 5 shows the box design, `box_assembled.iam`

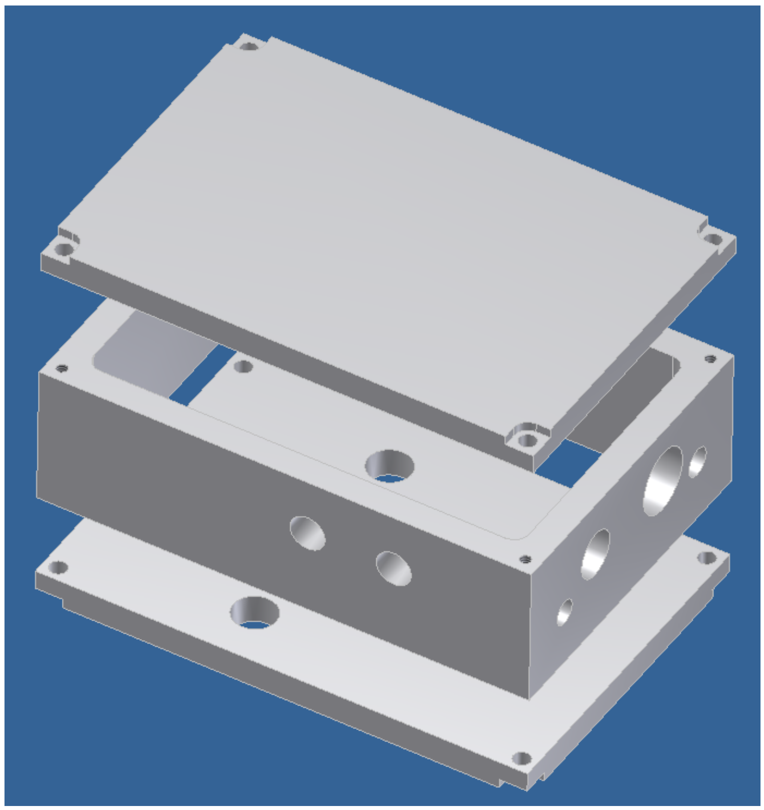

Figure 5: The box design

References

[1] “Quantum Optics: An Introduction”, Mark Fox, Oxford University Press, 2006

[2] “Photodiode Front Ends The REAL Story”, Philip C.D. Hobbs, April 2001, Optics & Photonics News 42-45

[3] “Building Electro-Optical Systems: Making It all Work”, Philip C.D. Hobbs, Wiley, 2009

[4] “Reaching the Shot Noise Limit for $10”, Philip C.D. Hobbs, April 1991, Optics & Photonics News, 17-23

[5] Newport Nirvana Auto-Balanced Photoreceiver 2007 Manual, `http://assets.newport.com/webDocuments-EN/images/15187.pdf` retrieved 2012-09-12

[6] Hamamatsu S5971 Data Sheet `http://sales.hamamatsu.com/assets/pdf/parts_S/S5971_etc.pdf` retrieved 2012-09-12